

IBM Spectrum Protect Plus
Version 10.1.3

Installation and User's Guide



Sixth edition (October 2019)

This edition applies to version 10, release 1, modification 3 of IBM Spectrum Protect Plus (product number 5737-F11) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2017, 2019.**

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

Contents

| | |
|--|------------|
| About this publication..... | vii |
| Who should read this publication..... | vii |
| What's new in V10.1.3..... | ix |
| Chapter 1. Product overview..... | 1 |
| Product components..... | 1 |
| Product dashboard..... | 3 |
| Product deployment roadmap..... | 4 |
| Alerts..... | 5 |
| Role-based access control..... | 5 |
| Offload to secondary backup storage..... | 6 |
| Replicate backup-storage data..... | 8 |
| IBM Spectrum Protect Plus on IBM® Cloud..... | 9 |
| IBM Spectrum Protect Plus on AWS..... | 9 |
| Chapter 2. Installing IBM Spectrum Protect Plus..... | 11 |
| System requirements | 11 |
| Component requirements | 11 |
| Hypervisor requirements | 23 |
| File indexing and restore requirements..... | 23 |
| Microsoft Exchange Server requirements..... | 27 |
| Db2® requirements..... | 29 |
| MongoDB requirements..... | 31 |
| Oracle requirements..... | 34 |
| Microsoft SQL Server requirements..... | 37 |
| Obtaining the IBM Spectrum Protect Plus installation package..... | 41 |
| Installing IBM Spectrum Protect Plus as a VMware virtual appliance..... | 41 |
| Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance..... | 43 |
| Assigning a static IP address..... | 44 |
| Uploading the product key..... | 45 |
| Installing iSCSI initiator utilities..... | 45 |
| Chapter 3. Installing and configuring vSnap servers..... | 47 |
| Installing vSnap servers..... | 47 |
| Installing a physical vSnap server..... | 47 |
| Installing a virtual vSnap server in a VMware environment..... | 48 |
| Installing a virtual vSnap server in a Hyper-V environment..... | 49 |
| Managing vSnap servers..... | 50 |
| Adding a vSnap server as a backup storage provider..... | 50 |
| Initializing the vSnap server..... | 53 |
| Setting vSnap storage options..... | 54 |
| Expanding a vSnap storage pool..... | 54 |
| Establishing a replication partnership for vSnap servers..... | 55 |
| Changing offload throughput rate..... | 55 |
| vSnap server administration reference | 56 |
| Storage management..... | 57 |
| Network management..... | 59 |
| Uninstalling a vSnap server..... | 60 |

| | |
|--|------------|
| Chapter 4. Getting off to a quick start..... | 61 |
| Start IBM Spectrum Protect Plus..... | 62 |
| Create backup policies..... | 62 |
| Create a user account for the application administrator..... | 64 |
| Add resources to protect..... | 66 |
| Add resources to a job definition..... | 68 |
| Start a job..... | 69 |
| Run a report..... | 70 |
| Chapter 5. Updating IBM Spectrum Protect Plus components..... | 73 |
| Updating the IBM Spectrum Protect Plus virtual appliance..... | 73 |
| Updating vSnap servers..... | 75 |
| Updating the operating system for a physical vSnap server..... | 75 |
| Updating the operating system for a virtual vSnap server..... | 75 |
| Updating a vSnap server..... | 76 |
| Updating VADP proxies..... | 76 |
| Applying early availability updates..... | 77 |
| Chapter 6. Managing SLA policies for backup operations..... | 79 |
| Creating an SLA policy..... | 79 |
| Editing an SLA policy..... | 81 |
| Deleting an SLA policy..... | 82 |
| Chapter 7. Protecting hypervisors..... | 83 |
| VMware..... | 83 |
| Adding a vCenter Server instance..... | 83 |
| Backing up VMware data..... | 91 |
| Managing VADP backup proxies..... | 95 |
| Restoring VMware data..... | 99 |
| Hyper-V..... | 107 |
| Adding a Hyper-V server..... | 107 |
| Backing up Hyper-V data..... | 109 |
| Restoring Hyper-V data..... | 112 |
| Restoring files..... | 116 |
| Chapter 8. Protecting applications..... | 119 |
| Db2..... | 119 |
| Prerequisites for Db2..... | 119 |
| Adding a Db2 application server..... | 122 |
| Backing up Db2 data..... | 125 |
| Restoring Db2 data | 130 |
| Exchange Server..... | 144 |
| Prerequisites..... | 144 |
| Privileges | 144 |
| Adding an Exchange application server..... | 145 |
| Defining a Service Level Agreement backup job..... | 147 |
| Incremental forever backup strategy..... | 149 |
| Restoring Microsoft Exchange databases | 150 |
| MongoDB..... | 160 |
| Prerequisites for MongoDB..... | 160 |
| Adding a MongoDB application server..... | 163 |
| Backing up MongoDB data..... | 166 |
| Restoring MongoDB data | 170 |
| SQL Server..... | 182 |
| Adding a SQL Server application server..... | 182 |
| Backing up SQL Server data..... | 184 |

| | |
|--|------------|
| Restoring SQL Server data..... | 187 |
| Oracle..... | 191 |
| Adding an Oracle application server..... | 191 |
| Backing up Oracle data..... | 193 |
| Restoring Oracle data..... | 195 |
| Chapter 9. Protecting IBM Spectrum Protect Plus..... | 201 |
| Backing up the application..... | 201 |
| Restoring the application..... | 201 |
| Managing restore points..... | 202 |
| Deleting IBM Spectrum Protect Plus resources from the catalog..... | 203 |
| Chapter 10. Managing jobs..... | 205 |
| Job types..... | 205 |
| Starting jobs..... | 206 |
| Pausing and resuming jobs..... | 206 |
| Canceling jobs..... | 206 |
| Rerunning partially completed backup jobs..... | 207 |
| Backing up a single resource..... | 207 |
| Configuring scripts for backup and restore operations..... | 208 |
| Uploading a script..... | 208 |
| Adding a script to a server..... | 208 |
| Chapter 11. Configuring and maintaining the IBM Spectrum Protect Plus system environment..... | 211 |
| Managing secondary backup storage..... | 211 |
| Managing cloud storage..... | 211 |
| Managing repository server storage..... | 215 |
| Managing keys and certificates..... | 221 |
| Managing sites..... | 224 |
| Adding a site..... | 224 |
| Editing a site..... | 224 |
| Deleting a site..... | 224 |
| Managing LDAP and SMTP servers..... | 224 |
| Adding an LDAP server..... | 225 |
| Adding an SMTP server..... | 226 |
| Editing settings for an LDAP or SMTP server..... | 227 |
| Deleting an LDAP or SMTP server..... | 227 |
| Logging on to the administrative console..... | 228 |
| Setting the time zone..... | 228 |
| Uploading an SSL certificate from the administrative console..... | 229 |
| Uploading an SSL certificate from the command line..... | 230 |
| Logging on to the virtual appliance..... | 230 |
| Accessing the virtual appliance in VMware..... | 230 |
| Accessing the virtual appliance in Hyper-V..... | 231 |
| Testing network connectivity..... | 231 |
| Running the Service Tool from a command-line interface..... | 231 |
| Running the Service Tool remotely..... | 232 |
| Adding virtual disks..... | 233 |
| Adding a disk to the virtual appliance..... | 233 |
| Adding storage capacity from a new disk to the appliance volume..... | 234 |
| Chapter 12. Managing reports and logs..... | 237 |
| Types of reports..... | 237 |
| Backup storage utilization reports..... | 237 |
| Protection reports..... | 238 |
| System reports..... | 239 |

| | |
|---|------------|
| VM Environment reports..... | 240 |
| Report actions..... | 241 |
| Running a report..... | 241 |
| Creating a custom report..... | 242 |
| Scheduling a report..... | 242 |
| Collecting and reviewing audit logs for actions..... | 243 |
| Enabling the display of alerts in the system log..... | 243 |
| Chapter 13. Managing user access..... | 245 |
| Managing user resource groups..... | 246 |
| Creating a resource group..... | 246 |
| Editing a resource group..... | 248 |
| Deleting a resource group..... | 249 |
| Managing roles..... | 249 |
| Creating a role..... | 250 |
| Editing a role..... | 252 |
| Deleting a role..... | 253 |
| Managing user accounts..... | 253 |
| Creating a user account for an individual user..... | 253 |
| Creating a user account for an LDAP group..... | 253 |
| Editing a user account credentials..... | 254 |
| Deleting a user account..... | 254 |
| Managing identities..... | 255 |
| Adding an identity..... | 255 |
| Editing an identity..... | 255 |
| Deleting an identity..... | 255 |
| Chapter 14. Licensing overview..... | 257 |
| Software License Metric (SLM) tags..... | 257 |
| Integration with IBM License Metric Tool (ILMT)..... | 258 |
| Chapter 15. Troubleshooting..... | 259 |
| Collecting log files for troubleshooting..... | 259 |
| Appendix A. Search guidelines..... | 261 |
| Appendix B. Accessibility..... | 263 |
| Notices..... | 265 |
| Glossary..... | 269 |
| Index..... | 271 |

About this publication

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

Who should read this publication

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

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

What's new in Version 10.1.3

IBM Spectrum Protect Plus Version 10.1.3 introduces new features and updates.

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

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 cloud storage or a repository server such as an IBM Spectrum Protect server 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 Blueprints](#).

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

IBM Spectrum Protect Plus server

This component manages the entire system. The server consists of several catalogs that track various system aspects such as restore points, configuration, permissions, and customizations. Typically, there is one IBM Spectrum Protect Plus 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.

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

Site

This component is an IBM Spectrum Protect Plus policy construct that is used to manage data placement in the environment. A site can be physical (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.

Example 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.

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

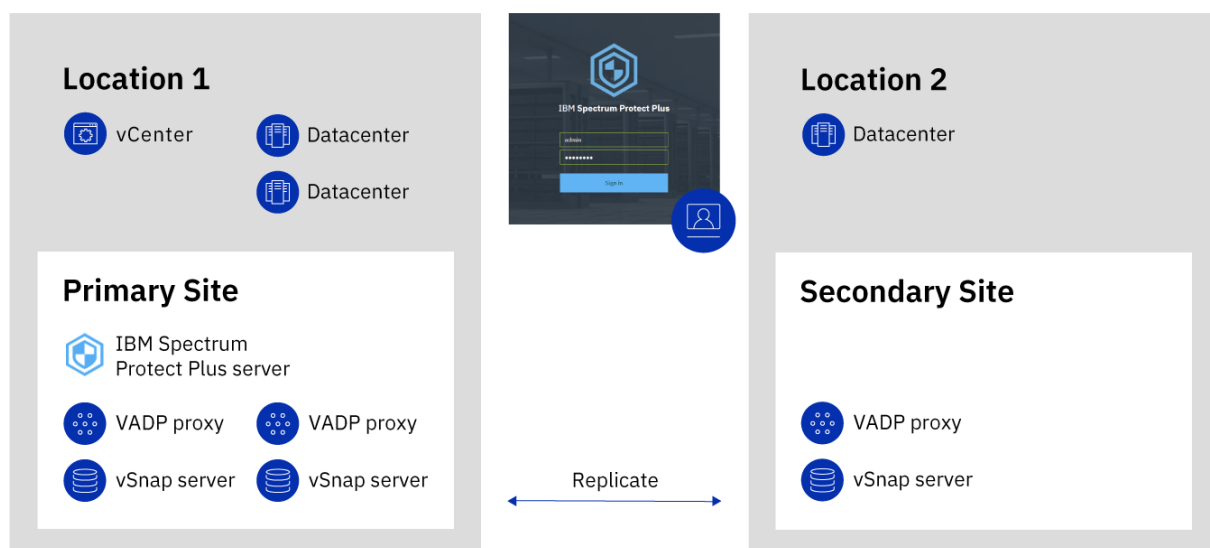


Figure 1. IBM Spectrum Protect Plus deployment across two geographical locations

Product dashboard

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

Jobs and Operations

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

Currently Running

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

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

History

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

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

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

Completed jobs are shown by job status:

Successful

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

Failed

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

Warning

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

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

Destinations

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

Capacity Summary

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

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

Total Devices

The **Total Devices** section shows the total number of devices that are available for use.

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

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

Data Reduction

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

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

Coverage

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

Source Protection

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

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

Policies

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

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

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

Product deployment roadmap

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

| Action | How to |
|--|--|
| Ensure that your system environment meets the hardware and software requirements. | See “System requirements” on page 11. |
| Determine how to size, build, and place the components in your IBM Spectrum Protect Plus environment. | See the IBM Spectrum Protect Plus Blueprints . |
| Install IBM Spectrum Protect Plus. | See Chapter 2, “Installing IBM Spectrum Protect Plus,” on page 11. |
| If additional vSnap servers are required to support your environment, install and configure the servers. | See Chapter 3, “Installing and configuring vSnap servers,” on page 47. |

| Action | How to |
|--|--|
| 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 95. |
| 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 61. |

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 245.

Offload to secondary backup storage

The vSnap server is the primary backup location for snapshots and all IBM Spectrum Protect Plus environments have at least one vSnap server. Optionally, you can offload snapshots, to cloud storage or a repository server for longer-term protection.

The following concepts summarize the IBM Spectrum Protect Plus backup and restore operations for offloading snapshots:

Backup operations

- The following storage targets are available for offload backup operations:
 - IBM Cloud™ Object Storage (including IBM Cloud Object Storage Systems)
 - Simple Storage Service (Amazon S3)
 - Microsoft Azure
 - Repository servers (for the current release of IBM Spectrum Protect Plus, the repository server must be an IBM Spectrum Protect server)
- During the first offload of a backup volume, the snapshot is backed up in full. After the first offload of the base snapshot is completed, subsequent offloads are incremental and capture cumulative changes since the last offload.
- After an offload process is completed, all metadata is written to a blockmap file, which is stored in the cloud. The blockmap allows cloud restore operations to be performed from any available vSnap server. vSnap servers also maintain a manifest file that logs all blockmap files for each offload of a vSnap volume. The manifest file is uploaded to the cloud along with the blockmap file. After each offload of a volume, the manifest file is updated and replaced on the cloud.
- A Virtual Block Device (VBD) is created using all available blockmap files and is used to map changed blocks and organize them into objects to be sent to cloud resources.

Restore operations

- An offload pool is imported to a vSnap server and the cloud offload restore operation proceeds as if the selected snapshot was available locally. IBM Spectrum Protect Plus is used to restore the snapshots from a vSnap server to the original or an alternate destination.

Expiration operations

As snapshots on a vSnap server expire or are deleted, the blocks that are specific to that snapshot are no longer required for restore operations. When all blocks for a cloud resource are expired, the offloaded data is then expired.

Adding secondary backup storage and creating backup policies

To offload data to secondary storage, the following actions are required:

| Action | How to |
|--|---|
| To offload to a repository server <ul style="list-style-type: none">• Set up IBM Spectrum Protect Plus as an object client in the IBM Spectrum Protect server environment.• Add the storage to IBM Spectrum Protect Plus. | See “Configuring IBM Spectrum Protect server as an offload target” on page 215 and “Adding a repository server as a backup storage provider” on page 219. |

| Action | How to |
|--|---|
| To offload to cloud storage, add the storage to IBM Spectrum Protect Plus. | Follow the instructions for your selected storage type: <ul style="list-style-type: none"> • “Adding Amazon S3 cloud storage as a backup storage provider” on page 211 • “Adding IBM Cloud Object Storage as a backup storage provider” on page 212 • “Adding Microsoft Azure cloud storage as a backup storage provider” on page 213 • “Adding a repository server as a backup storage provider” on page 219 |
| Create a backup policy that includes the storage. | See “Create backup policies” on page 62 . |

Example deployments

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.

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

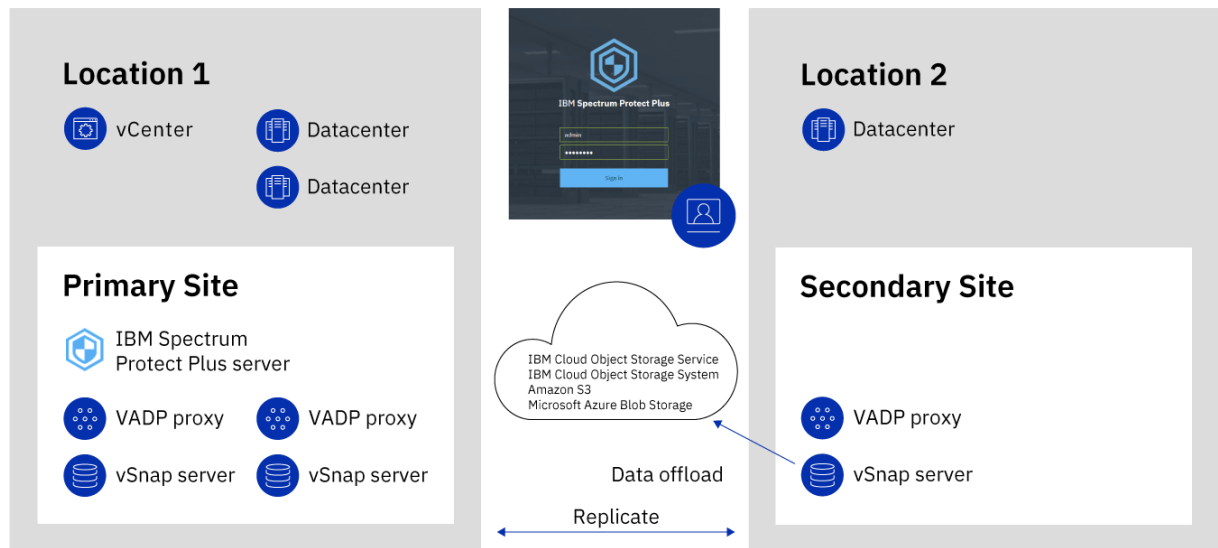


Figure 2. IBM Spectrum Protect Plus deployment across two geographical locations with offload to cloud storage

Figure 2 shows the same deployment as Figure 1.

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

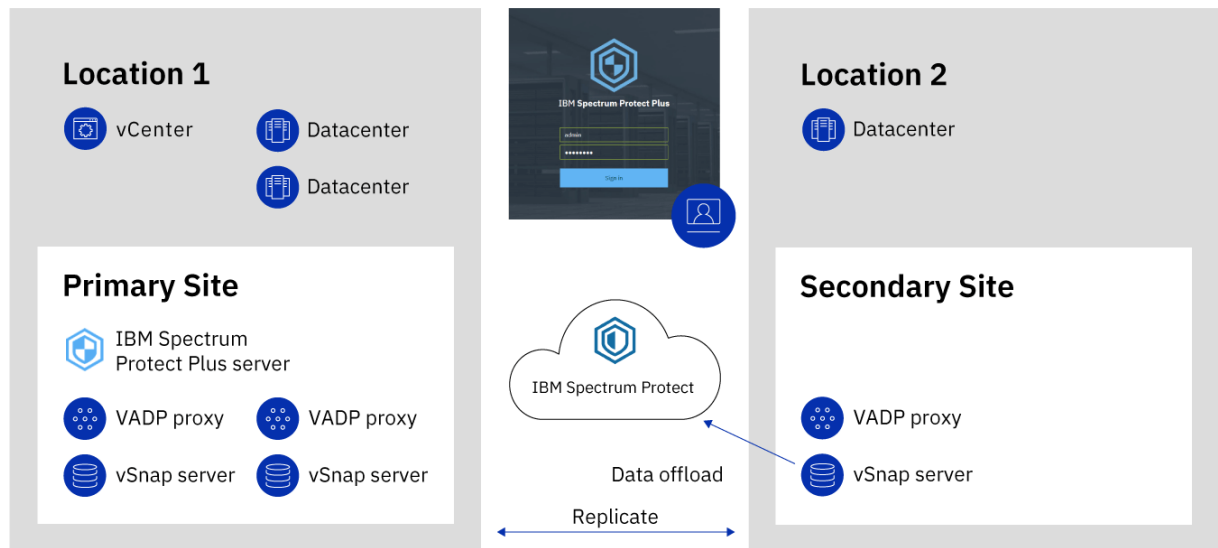


Figure 3. IBM Spectrum Protect Plus deployment across two geographical locations with offload to IBM Spectrum Protect

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 79.

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.

IBM Spectrum Protect Plus on IBM Cloud

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

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

IBM Cloud for VMware Solutions provides the following major benefits:

Global reach

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

Streamlined integration

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

Automated deployment and configuration

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

Simplification

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

Expansion and contraction flexibility

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

Single management console

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

Available features in IBM Spectrum Protect Plus on IBM Cloud

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

However, IBM Spectrum Protect Plus on IBM Cloud supports only VMware environments.

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

The current version of IBM Spectrum Protect Plus and IBM Spectrum Protect Plus on IBM Cloud might not be the same. To find the documentation for the version of IBM Spectrum Protect Plus on IBM Cloud that you are using, go to the [online product documentation](#) and select the product version.

For more information

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

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

IBM Spectrum Protect Plus on the AWS cloud platform

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

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

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

Deploying IBM Spectrum Protect Plus to AWS

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

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.

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

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

Component requirements

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

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

IBM Spectrum Protect Plus support for third-party platforms, applications, services, and hardware 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 machine installation

IBM Spectrum Protect Plus is installed as a virtual appliance. Before you deploy IBM Spectrum Protect Plus 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.

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).

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

- 64-bit 8-core machine
- 48 GB memory
- 536 GB disk storage for virtual machine

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](#).

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. Note that later browser versions might also 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 Spectrum Protect requirements

When you are offloading data using Amazon S3, to the IBM Spectrum Protect Server repository, ensure that IBM Spectrum Protect is at Version 8.1.7 or later.

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 secure connections (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 micro-service 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. |

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

| Port | Protocol | Firewall | Service | Description |
|------|----------|----------|--|---|
| 5672 | AMQP | Blocked | RabbitMQ | Message framework used to manage messages produced and consumed in the IBM Spectrum Protect Plus appliance. |
| 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 run 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. |
| 9090 | TCP | Accept | DOCUMENTATION | Default port for the IBM Spectrum Protect Plus help system |

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

| Port | Protocol | Firewall | Service | Description |
|-------------|-----------------|-----------------|----------------|---|
| 27017 | TCP | Blocked | MongoDB | Persists configuration related documents for IBM Spectrum Protect Plus. |
| 27018 | TCP | Blocked | MongoDB | 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. |

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

| Port | Protocol | Service | Description |
|------|----------|-----------------------------------|--|
| 443 | TCP | VMware vCenter | Client connections to vCenter. |
| 636 | TCP | LDAP | Active directory services (SSL) |
| 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. |

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

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

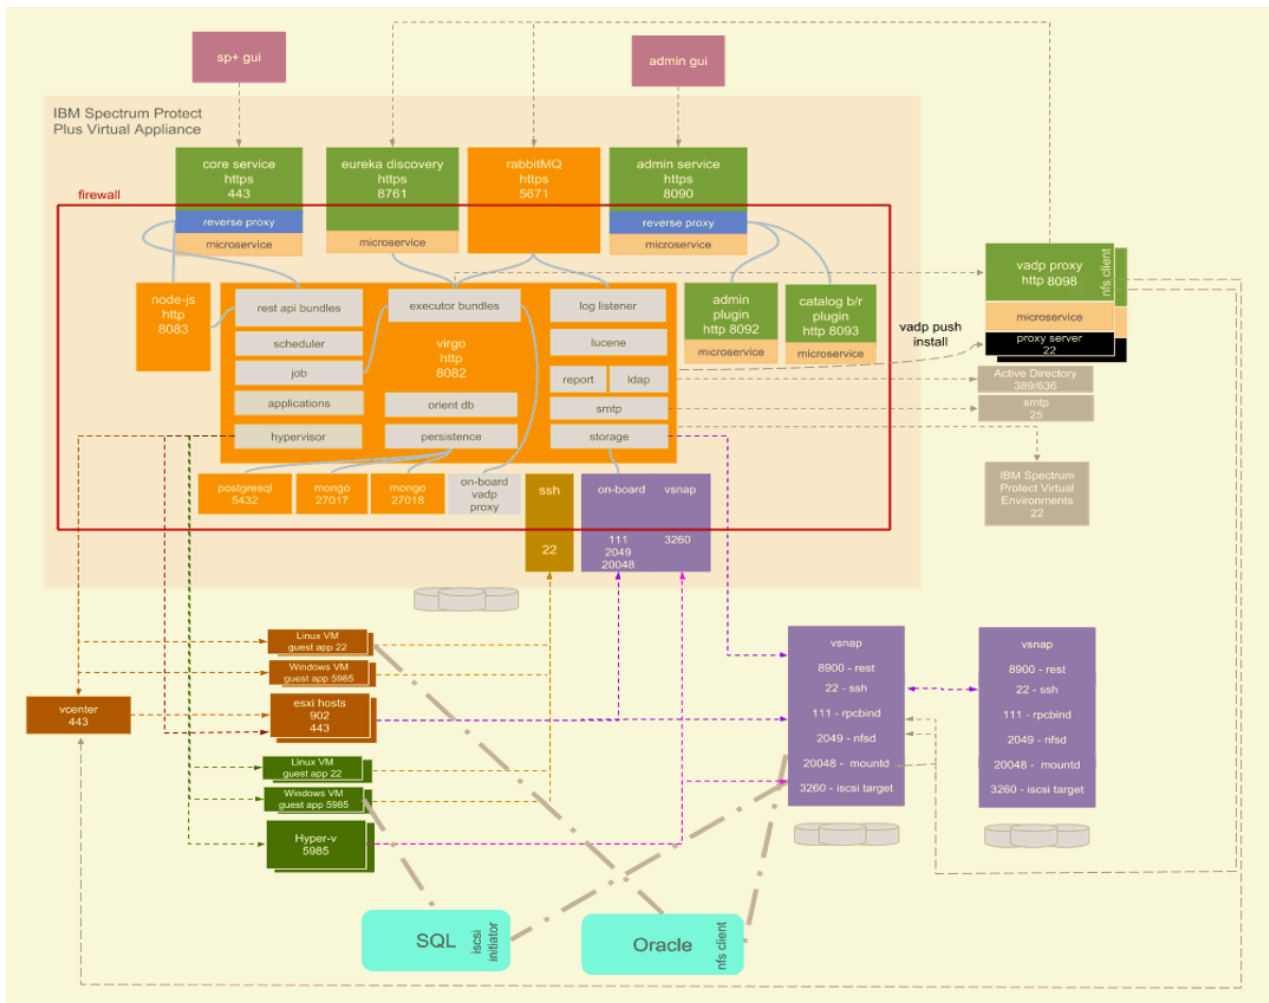


Figure 4. IBM Spectrum Protect Plus virtual appliance

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, more vSnap servers might be required.

Memory should be adjusted based on backup capacity for more efficient deduplication. For more information and sizing guidance, see [IBM Spectrum Protect Plus Blueprints](#).

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

Note: The Linux Network Management service must be installed and running.

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 that you have met the following requirements:

- 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

IBM Spectrum Protect Plus V10.1.3 provides new functionality that requires the kernel levels supported in RHEL 7.5 and CentOS 7.5. Use IBM Spectrum Protect Plus V10.1.2 for physical vSnap V10.1.2 installations if you need to use operating systems earlier than RHEL 7.5 and CentOS 7.5.

The following Linux operating systems are supported for IBM Spectrum Protect Plus V10.1.3 physical vSnap server installations:

- CentOS 7.1804 (7.5) (x86_64)
- CentOS 7.1810 (7.6) (x86_64) (beginning with v10.1.3 patch1)
- RedHat Enterprise Linux 7.5 (x86_64)
- RedHat Enterprise Linux 7.6 (x86_64) (beginning with v10.1.3 patch1)

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

- CentOS Linux 7.3.1611 (x86_64)
- CentOS Linux 7.4.1708 (x86_64)
- Red Hat Enterprise Linux 7.3 (x86_64)
- Red Hat Enterprise Linux 7.4 (x86_64)

vSnap server ports

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

| 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). |
| 137 | TCP | Accept | SMB/CIFS | Used for SMB/CIFS data transfer to and from vSnap (internal). |

Table 4. Incoming vSnap firewall connections (continued)

| Port | Protocol | Firewall | Service | Description |
|-------|----------|----------|----------|---|
| 138 | TCP | Accept | SMB/CIFS | Used for SMB/CIFS data transfer to and from vSnap (internal). |
| 139 | TCP | Accept | SMB/CIFS | Used for SMB/CIFS data transfer to and from vSnap (internal). |
| 445 | TCP | Accept | SMB/CIFS | Used for SMB/CIFS data transfer to and from vSnap (internal). |
| 2049 | TCP | Accept | NFS | Used for NFS data transfer to and from vSnap (internal). |
| 8900 | TCP | Accept | HTTPS | vSnap REST APIs |
| 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 from the IBM Spectrum Protect Plus appliance onto the proxies. This processing has been tested 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](#).

VADP proxies are supported only in 64-bit quad core and higher 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

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

- 64-bit quad core processor
- 8 GB RAM required, 16 GB recommended

- 60 GB free disk space

For more information and sizing guidance, see [IBM Spectrum Protect Plus Blueprints](#).

Increase of used CPUs and concurrency on the VADP proxy server, requires the memory allocated on the proxy server to be increased.

The proxy must be able 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.

Each proxy must have a fully qualified domain name and must be able to resolve and reach the 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 secure connections (HTTPS/SSL).

| <i>Table 5. Incoming VADP proxy firewall connections</i> | | | | |
|--|----------|----------|---------|---|
| Port | Protocol | Firewall | Service | Description |
| 22 | TCP | Accept | SSH | Port 22 is used to push the VADP proxy to the host node. |
| 8098 | TCP | Accept | VADP | Default port for TLS-based REST API communications between the IBM Spectrum Protect Plus server and the VADP proxy. |

| <i>Table 6. Outgoing VADP proxy firewall connections</i> | | | | |
|--|----------|----------|---------------------------|---|
| 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 on 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 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.

Cloud offload requirements

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

Disk cache area

For all functionality relating to offloading or restoring from the cloud, the vSnap server requires a disk cache area on the vSnap server.

- During offload operations, this cache is used as a temporary staging area for objects that are pending upload to the cloud endpoint.
- During restore operations, it is used to cache downloaded objects as well as to store any temporary data that may be written into the restore volume.

Most of the cache space is freed up at the end of each offload or restore operation, but a small amount may continue to be used to cache metadata that is used to speed up subsequent operations. The cache area must be configured in the form of an XFS filesystem mounted at `/opt/vsnap-data` on the vSnap server. If this mount point is not configured, offload or restore jobs fail with this error: Cloud functionality disabled: Data disk `/opt/vsnap-data` is not configured.

Note:

Do not unmount or manipulate files under `/opt/vsnap-data` while any offload or restore jobs are active. Once you have ensured that no jobs are active, it is safe to run any maintenance activities such as unmounting and reconfiguring the cache area. The data stored under `/opt/vsnap-data` is also safe to delete as long as no offload or restore jobs are active. Deleting this data may result in the vSnap server needing to re-download it from the cloud endpoint during the next offload or restore operation, which may introduce a delay during the job.

For new installations of vSnap V10.1.3

When the vSnap is deployed as a virtual appliance, the cache area is already present as a preconfigured 128 GB data disk mounted at `/opt/vsnap-data`. When the vSnap is installed on a custom server, the cache area must be configured manually.

For systems upgraded from vSnap V10.1.2 to V10.1.3

A default preconfigured cache area of 128 GB may already be present and mounted at `/opt/vsnap-data` if the system was previously deployed as a virtual appliance with vSnap V10.1.2. If the system was previously upgraded from vSnap V10.1.1, the cache area is not present. Use the **df** command on the vSnap server to confirm the presence of this mount point. If the preconfigured mount point is not present, it must be configured manually. For more information about sizing and installing cache, see the [Cloud offload configuration document](#).

Certificate requirements

- Self-signed certificates: If the cloud endpoint or repository server uses a self-signed certificate, the certificate must be specified in Privacy Enhanced Mail (PEM) format, when registering the cloud or repository server in the IBM Spectrum Protect Plus user interface.
- Certificates signed by private Certificate Authority: If the cloud endpoint or repository server uses a certificate signed by a private Certificate Authority (CA), the endpoint certificate must be specified (in PEM format) when registering the cloud or repository server in the IBM Spectrum Protect Plus user interface. In addition, the root/intermediate certificate of the private CA, must be added to the system certificate store in each vSnap server with the following procedure:
 - Login to the vSnap server console as the `serveradmin` user, and upload the private CA certificates (in PEM format) to a temporary location.
 - Copy each certificate file to the system certificate store directory

```
/etc/pki/ca trust/source/anchors/ $ sudo cp /tmp/private-ca-cert.pem  
/etc/pki/ca-trust/source/anchors/
```

- Run the following command to update the system certificate bundle to incorporate the added custom certificate:

```
$ sudo update-ca-trust
```

- Certificates signed by public Certificate Authority: If the cloud endpoint uses a public CA-signed certificate, no special action is needed. The vSnap server validates the certificate by using the default system certificate store.

Network requirements

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

| Port | Protocol | Service | Description |
|------|----------|---------|---|
| 443 | TCP | HTTPS | Allows vSnap to communicate with Amazon S3, Azure, or IBM Cloud Object Storage endpoints. |
| 9000 | TCP | HTTPS | Allows vSnap to communicate with IBM Spectrum Protect (Repository Server) endpoints. |

If there are any firewalls or network proxies that run SSL Interception or Deep Packet Inspection for traffic between vSnap servers and cloud endpoints, this may interfere with SSL certificate validation on the vSnap servers. This interference may cause cloud offload job failures. To prevent this, the vSnap servers must be exempted from SSL interception and inspection in the firewall or proxy configuration.

Cloud provider requirements

Native lifecycle management is not supported. IBM Spectrum Protect Plus manages the lifecycle of uploaded objects automatically using an incremental-forever approach where older objects may still be used by newer snapshots. Automatic or manual expiration of objects outside of IBM Spectrum Protect Plus will lead to data corruption. For more information about cloud providers using SSL certificates that are self-signed or signed by a private Certificate Authority, see Certificate requirements on this page.

- Amazon S3 cloud storage offload requirements: When the cloud provider is registered in IBM Spectrum Protect Plus, an existing bucket in one of the supported storage tiers must be specified: S3 Standard, S3 Intelligent-Tiering, S3 Standard-Infrequent Access, and S3 One Zone-Infrequent Access.
- IBM Cloud Object Storage offload requirements: When the cloud provider is registered in IBM Spectrum Protect Plus, an existing bucket must be specified. If the specified bucket has a WORM policy that locks objects for a certain time period, IBM Spectrum Protect Plus automatically detects the configuration and deletes snapshots after the WORM policy removes the lock.
- Microsoft Azure offload requirements: When the cloud provider is registered in IBM Spectrum Protect Plus, an existing container in a hot or cool storage account must be specified.
- Repository Server offload requirements: When the cloud provider is registered in IBM Spectrum Protect Plus, you cannot use an existing bucket. IBM Spectrum Protect Plus creates a uniquely named bucket for its own use.

For quick start information to help you to set up and offload data to specific cloud providers, see [Data offload to cloud object storage with IBM Spectrum Protect Plus](#).

Hypervisor requirements

Review the hypervisor requirements for IBM Spectrum Protect Plus.

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

Hyper-V requirements

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

- Hyper-V Server 2016 or Microsoft Hyper-V on Windows Server 2016
- Microsoft Hyper-V on Windows Server 2019 (Clustering is not supported)

Backup and restore of shared virtual hard disks (shared VHDX) is not supported. For known issues and limitations, see <https://www.ibm.com/support/docview.wss?uid=ibm10843068>.

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.

Note: The DiskPart automount parameter must be enabled on the Hyper-V server. For more information about enabling the automount parameter, see [https://technet.microsoft.com/en-us/library/cc753703\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/cc753703(v=ws.11).aspx).

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 you are registering the cluster in IBM Spectrum Protect Plus, register the Failover Cluster Manager.

VMware requirements

The following VMware vSphere versions are supported:

- vSphere 5.5 and 5.5 update and patch levels
- vSphere 6.0 and 6.0 update and patch levels
- vSphere 6.5 and 6.5 update and patch levels
- vSphere 6.7 and 6.7 update and patch levels

Ensure that 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.

File indexing and restore requirements

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

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

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

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

VMware requirements

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

Windows requirements

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

Windows operating system and file system requirements

- Windows Server 2008 R2
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2019 core

Supported file systems

- NTFS
- ReFS
- FAT

Supported disk storage types

Basic disks with

- MBR partitions
- GPT partitions

Restriction: The backup or restore of files on dynamic disks is not supported.

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.



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

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

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

Ensure that 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.

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

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 `hostAddress` 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.vmdir.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 are specified for the virtual machine must include a user with the following privileges:

- The user identity must have the **Log on as a service** right, which is assigned through the Administrative Tools control panel on the local machine **Local Security Policy > Local Policies > User Rights Assignment > Log on as a service**.
- 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. 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 other file systems, but only the file systems listed here are eligible for file indexing and restore.

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.

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

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

Space requirements

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

Software requirements

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

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 is 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 is 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 re-signature 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 Universally Unique Identifier (UUID) is modified. To resolve this issue, update `xfsprogs` to version 4.2.0 or above.

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.

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 <agent>`

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

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

Microsoft Exchange Server requirements

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

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

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

Configuration

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

Application Versions

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

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

Operating Systems

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

Note: Windows Server 2019 Core Installations are supported, however the Granular restore feature is not supported on core installations.

Additional Notes

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

For information about virtualization support for Exchange Server, see the [“Prerequisites for Microsoft Exchange Server”](#) on page 144.

Software

Ensure the supported version of Windows x64 bit is installed.

The following prerequisites from Microsoft are required and must be installed prior to using IBM Spectrum Protect Plus.

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

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

Note: MAPI/CDO is required for Microsoft Exchange Server 2013 only. It is not required if you are running Microsoft Exchange Server 2016 or Exchange Server 2019.

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

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

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

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

Privileges

The IBM Spectrum Protect Plus agent users have the following privileges:

Microsoft Exchange Server is protected by role-based authentication. In order to get the Microsoft Exchange agent to work in your IBM Spectrum Protect Plus environment, you must set up the appropriate privileges. For more information, see [“Privileges” on page 144](#).

Ports

The following ports are used by IBM Spectrum Protect Plus agent users. Ports that are indicated with "Accept" in the Firewall Rule column use secure connections (HTTPS or SSL).

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

Table 9. Outgoing IBM Spectrum Protect Plus agent firewall connections

| Port | Protocol | Service | Description |
|-------|----------|----------------|---|
| 3260* | TCP | vSnap iSCSI | iSCSI vSnap target port used for mounting LUNS for backup and recovery |
| 137 | UDP | vSnap SMB/CIFS | vSnap SMB or CIFS target port used for mounting filesystem shares for transaction log backup and recovery |
| 138 | UDP | vSnap SMB/CIFS | vSnap SMB or CIFS target port used for mounting filesystem shares for transaction log backup and recovery |
| 139 | TCP | vSnap SMB/CIFS | vSnap SMB or CIFS target port used for mounting filesystem shares for transaction log backup and recovery |
| 445 | TCP | vSnap SMB/CIFS | vSnap SMB or CIFS target port used for mounting filesystem shares for transaction log backup and recovery |

*iSCSI initiator is required on this node.

Hardware

| System | Disk Space | Disk Space for Granular Restore Operations |
|---|--|---|
| x64: Compatible hardware supported by the operating system and the Microsoft Exchange Server. | A minimum of 200 MB of disk space for the product to be installed. | At least 2.1 GB disk space for "Additional Microsoft prerequisites", which will be installed automatically if they are not already. |

Db2 requirements

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

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

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

Configuration

The following IBM 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.
 - Red Hat Enterprise Linux 7 and later maintenance levels and mod levels.
 - SUSE Linux Enterprise Server 11.0 SP4 and later maintenance levels and mod levels
 - SUSE Linux Enterprise Server 12.0 SP1 and later maintenance levels and mod levels

Additional notes

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

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

- IBM Db2 is configured for single partitioned databases only. Backups of multi-partition IBM Db2 databases are not supported.
- IBM Db2 pureScale® is not supported.
- IBM Db2 archive logging is activated and IBM Db2 is in recoverable mode.
- Logical volumes holding IBM Db2 table spaces (data and temporary table spaces), the local database directory, and IBM Db2 log files are managed by Logical Volume Manager (LVM2) on Linux and by JFS2 on AIX respectively. LVM2 on Linux and JFS2 on AIX are used for creating temporary volume snapshots. The logical volume grows in size with data as it changes on the source volume while the snapshot exists. For more information, see [“LVM2 and JFS2” on page 121](#).

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.
Note: The supported Linux x86_64 operating systems include the required bash and sudo packages.
- Python version 2.6.x or 2.7.x must be installed.
- Ensure the supported version of Linux x86_64, or AIX is installed.

Connectivity

Ensure that the following connectivity criteria are in place:

- 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.
- On AIX, ensure that the NFS communication is configured with reserved ports by using the command:
`nfsd -p -o nfs_use_reserved_port=1`

Authentication and privileges

The IBM 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 IBM 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 IBM Db2 inventory using /usr/local/bin/db2ls. IBM Spectrum Protect Plus requires this privilege to discover and collect information about IBM Db2 instances and databases.

Ports

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

Table 10. Incoming IBM Spectrum Protect Plus agent firewall connections

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

Table 11. Outgoing IBM Spectrum Protect Plus agent firewall connections

| Port | Protocol | Service | Description |
|-------|----------|---------------------|---|
| 111 | TCP | vSnap RPC Port Bind | Allows clients to discover ports that Open Network Computing (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. |

MongoDB requirements

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

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

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

MongoDB configuration requirements

The following MongoDB database versions are supported:

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

Operating systems

The following operating systems are supported:

- On Linux x86_x64:
 - Red Hat Enterprise Linux 6.8 and later maintenance levels and mod levels.
 - CentOS 6.8 and later maintenance levels and mod levels.
 - Red Hat Enterprise Linux 7 and later maintenance levels and mod levels.
 - CentOS 7 and later maintenance levels and mod levels.
 - SUSE Linux Enterprise Server 12.0 SP1 and later maintenance levels and mod levels.

Additional notes

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

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

- MongoDB is configured as a standalone instance or replica set. Back up operations of MongoDB sharded cluster instances are not supported. A backup always includes all databases in the instance.
- The MongoDB instance is configured to use the WiredTiger Storage Engine
- The user in the MongoDB application server registration in IBM Spectrum Protect Plus must be able to retrieve server information and status from the MongoDB admin database.
- Logical volumes of MongoDB data and log paths are managed by Linux Logical Volume Manager (LVM2). LVM2 is used for creating temporary volume snapshots. The database files and the journal must be located on a single volume. The logical volume grows in size with data as it changes on the source volume while the snapshot exists. For more information, see [“Linux LVM2 ” on page 162](#).

Software

Review the following software requirements:

- Python version 2.6 or 2.7 must be installed.
- When the MongoDB application server runs RHEL or CentOS 6.x, ensure that the `openssl` package is at version 1.0.1e-57 or above. Run `yum update openssl` to update to this requirement.
- Ensure that the supported version of Linux x86_64 is installed.

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

Authentication and privileges

The MongoDB server must be registered in IBM Spectrum Protect Plus using an operating system user that exists on the MongoDB server. This user is referred to as *IBM Spectrum Protect Plus agent user*.

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

On MongoDB, SSL-based encryption and certificate-based authentication is not supported.

On MongoDB Enterprise Editions, only on-storage encryption is supported.

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

- Privileges to run commands as the root user and as the MongoDB software owner user by using sudo. IBM Spectrum Protect Plus requires this privilege for 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 run the standard MongoDB server module `/usr/local/bin/mongodb`. IBM Spectrum Protect Plus requires this privilege to use the pymongo API to connect to the MongoDB servers by using the assigned DNS/IP name and port. This mechanism is used to gather information about MongoDB instances and databases.
- If the MongoDB server is protected by role-based authentication, to get the MongoDB agent to work in your IBM Spectrum Protect Plus environment, you must set up the appropriate privileges. See [Chapter 13, “Managing user access,”](#) on page 245 and [“Roles for MongoDB”](#) on page 161.

Ports

The following ports are used by MongoDB servers. Ports that are marked as Accept use secure connections (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 Computing (ONC) clients require to communicate with ONC servers. |
| 2049 | TCP | vSnap NFS | Used for NFS file sharing via vSnap. |

Table 13. Outgoing MongoDB server firewall connections (continued)

| Port | Protocol | Service | Description |
|-------|----------|-----------------|---|
| 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.

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

Configuration requirements

Database versions

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

Note:

- For Oracle 12c multitenant databases, IBM Spectrum Protect Plus supports protection and recovery of the container database, including all pluggable databases under it. Granular recovery of specific PDBs can be performed via Instant Disk Restore recovery combined with RMAN.
- Oracle 12c or higher multithreaded configurations are not supported.
- Oracle 18c only standalone configurations are supported.

Operating systems

- 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

- 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.
- RAC databases must be configured such that the RMAN Snapshot Control File location points to shared storage accessible to all cluster instances.

Software

- 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.x or 2.7.x must be installed.
- **Red Hat Enterprise Linux/CentOS 6.x only:** 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

- 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 for the rest of this topic).

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 36.

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 `cron.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: `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
```

Ports

The following ports are used by IBM Spectrum Protect Plus Agents. 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 Computing (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.

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

Configuration

Database versions

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

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

Operating systems

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

Windows Remote Shell (WinRM) must be enabled.

Cluster drive configurations utilizing dynamic disks is not supported. Standalone dynamic disks configurations are supported. An iSCSI route must be enabled between the SQL system and vSnap server. For more information, see [Microsoft iSCSI Initiator Step-by-Step Guide](#).

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

In-Memory OLTP

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

- The maximum restore file path must be fewer than 256 characters, which is an SQL requirement. 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

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 meets the minimum file size threshold requirement. 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

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 journaling, or changing of the cataloged USN ID.
- The file size is less than the tracked threshold size, which by default is 1MB .
- The file size is smaller after a previous backup job.
- A file is added after a previous 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, any option can be selected. . When secondary replica is preferred, and more than one secondary replica is available, the IBM Spectrum Protect Plus job 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 with 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 address must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. Both must have a Windows Remote Management (WinRM) service that is listening on port 5985. The fully qualified domain name must be resolvable and can be routed from the IBM Spectrum Protect Plus server and from the vSnap server.

The user identity must have sufficient rights to install and start the IBM Spectrum Protect Plus Tools Service on the node, including **Log on as a service** rights. 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

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.

Privileges

The IBM Spectrum Protect Plus agent user of an SQL server needs to have:

- public and sysadmin permissions
- Windows local administration permission, which is required by VSS framework, and volume/disk access
- permission to access cluster resources in a SQL Server Always On and SQL FCI environment.

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 the following Microsoft article: [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 the following Microsoft article: [Securing Data and Log Files](#).

Ports

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

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

| 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. |
| 137 | UDP | vSnap SMB/CIFS | vSnap SMB/CIFS target port used for mounting filesystem shares for transaction log backup and recovery. |
| 138 | UDP | vSnap SMB/CIFS | vSnap SMB/CIFS target port used for mounting filesystem shares for transaction log backup and recovery. |
| 139 | TCP | vSnap SMB/CIFS | vSnap SMB/CIFS target port used for mounting filesystem shares for transaction log backup and recovery. |
| 445 | TCP | vSnap SMB/CIFS | vSnap SMB/CIFS target port used for mounting filesystem shares for transaction log 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 743897](#).

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.

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

Related concepts

[“Updating IBM Spectrum Protect Plus components” on page 73](#)

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

Related tasks

[“Installing IBM Spectrum Protect Plus as a VMware virtual appliance” on page 41](#)

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 43](#)

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 47](#)

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 11](#) and [“Hypervisor requirements” on page 23](#).
- Download the virtual appliance template installation file CNZF0EN . ova from Passport Advantage Online. For information about downloading files, see [technote 743897](#).

- Run MD5 Checksum on the downloaded template installation file. Ensure that the generated checksum matches the one provided in the MD5 Checksum file, which is part of the software download.
- During deployment, you will be prompted to enter network properties from the VMware user interface. You can enter a static IP address configuration, or leave all fields blank to use a DHCP configuration.
- To reassign a static IP address after deployment, you can use the NetworkManager text user interface (nmtui) tool. For more information, see [“Assigning a static IP address” on page 44.](#)

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 OVF or OVA file**
2. Specify the location of the CNZF0EN . ova 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 45. |
| Start IBM Spectrum Protect Plus from a supported web browser. | See “Start IBM Spectrum Protect Plus” on page 62. |

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 11](#) and [“Hypervisor requirements ” on page 23.](#)
- Download the installation file CNZF1EN . exe from Passport Advantage Online. For information about downloading files, see [technote 743897.](#)
- 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 CNZF1EN . exe 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 44](#).

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 45 . |
| Start IBM Spectrum Protect Plus from a supported web browser. | See “Start IBM Spectrum Protect Plus” on page 62 . |

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 41](#)

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 43](#)

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 62 . |

Installing iSCSI initiator utilities

You must install Internet Small Computer System Interface (iSCSI) utilities if iSCSI mounted storage devices are directly connected to the IBM Spectrum Protect Plus appliance or to a vSnap server. After the

iSCSI initiator utilities are installed, iSCSI mounted storage devices can be connected to the appliance or server on which the package is installed.

About this task

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

Procedure

1. Log on to the appliance or server that is to be directly connected to the iSCSI mounted storage.
 - For the IBM Spectrum Protect Plus appliance, use the Secure Shell (SSH) protocol and authenticate with the appropriate administrative credentials.
 - For a vSnap server, use SSH or access the server directly and authenticate with the appropriate administrative credentials.
2. Install the iSCSI initiator utilities by running the following command:

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

Chapter 3. Installing 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. An onboard vSnap server resides on a partition of the IBM Spectrum Protect Plus appliance and is registered and initialized in IBM Spectrum Protect Plus. In smaller backup environments, the onboard vSnap server might be sufficient.

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

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

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

1. Install the vSnap server.
2. 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 11.
2. Download the installation package. Different installation files are provided for installation on physical or virtual machines. Ensure that you download the correct files for your environment. For more information about downloading files, see [technote 743897](#).

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 17 for supported operating systems.
The minimum installation configuration is sufficient, but you can also install additional packages including a graphical user interface (GUI). The root partition must have at least 8 GB of free space after installation.
2. Edit the `/etc/selinux/config` file to change the SELinux mode to Permissive.
3. Run `setenforce 0` to apply the setting immediately without requiring a restart.
4. Download the vSnap installation file CNZF4EN .run from Passport Advantage® Online. For information about downloading files, see [technote 743897](#).

5. 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 50 . |

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

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

Before you begin

For easier network administration, use a static IP address for the virtual machine. Assign the address by using the NetworkManager text user interface (nmtui) tool. For instructions, see [“Assigning a static IP address” on page 44](#), Work with your network administrator when configuring network properties.

Procedure

1. Download the server and proxy template installation file `CNZF2EN.ova` from Passport Advantage Online. For information about downloading files, see [technote 743897](#).
2. To deploy the vSnap server, take one of the following actions:
 - If you are using the vSphere Client to deploy the vSnap server, from the **File** menu, click **Deploy OVF Template**.
 - If you are using the vSphere Web Client, click **Create/Register VM**, then click **Deploy a virtual machine from an OVF or OVA file**. Click **Next**.
3. Specify the location of the `CNZF2EN.ova` 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 that is 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. If you are using a Dynamic Host Configuration Protocol (DHCP) configuration, leave all fields blank.

Restriction: A default gateway must be properly configured before deployment of the OVF template. Multiple DNS strings are supported, and must be separated by commas without the use of spaces.

The network prefix should be specified by a network administrator. The network prefix must be entered using CIDR notation; valid values are 1 - 32.
10. Provide details of the VADP configuration, including the IP address of the IBM Spectrum Protect Plus appliance.

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

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

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

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

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

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

What to do next

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

| Action | How to |
|--|---|
| Add the vSnap server to IBM Spectrum Protect Plus and configure the vSnap environment. | See “Managing vSnap servers” on page 50 . |
| Configure the VADP environment. | See “Setting options for VADP proxies” on page 98 . |

Installing a virtual vSnap server in a Hyper-V environment

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

Before you begin

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

Procedure

1. Download the vSnap installation file CNZF3EN.exe from Passport Advantage Online. For information about downloading files, see [technote 743897](#).
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 44](#).

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 50 . |

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 56](#).

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 51.

Password

Enter the password for the user.

7. Click **Save**.

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

What to do next

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

| Action | How to |
|--|--|
| Initialize the vSnap server. | See “ Initializing the vSnap server ” on page 53. |
| Expand the vSnap storage pool. | See “ Expanding a vSnap storage pool ” on page 54. |
| If necessary, configure and administer certain aspects of vSnap, such as network configuration or storage pool management. | “ vSnap server administration reference ” on page 56 |

Related tasks

“[Start IBM Spectrum Protect Plus](#)” on page 62


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

Editing settings for a vSnap server

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

Procedure

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

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

Deleting a vSnap server

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

Before you begin

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



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

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

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

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

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

- Do not delete the vSnap server. If you delete the vSnap server, recovery points that are associated with the server will be purged and backups will be rebased.
- Edit the settings for the vSnap server to specify the new host name or IP address. To edit the settings for a vSnap server, follow the instructions [“Editing settings for a vSnap server” on page 51](#).

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


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

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

1. Delete the vSnap server and run a maintenance job to ensure that recovery points that are associated with the old vSnap pool are purged from IBM Spectrum Protect Plus.
 - Incremental backups of the data that was present in the old pool will no longer be possible.
 - Recovering data that was present in the old pool will no longer be possible.
2. On the vSnap server, create a pool.
3. Add the vSnap server back into IBM Spectrum Protect Plus. To add a vSnap server to IBM Spectrum Protect Plus, see [“Adding a vSnap server as a backup storage provider” on page 50](#).
 - Subsequent runs of backup jobs will automatically create volumes on this or another vSnap server in the same site and will perform new base backups.

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 57](#).

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

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

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

Completing a simple initialization

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

About this task

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

Procedure

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

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

Initialize with Encryption

Enable encryption of backup data on the vSnap server.

Initialize

Initialize the vSnap server without encryption enabled.

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

Completing an advanced initialization

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

Procedure

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

1. Log in to the vSnap server console with the user ID `serveradmin`. The initial password is `sppDP758`.

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 56.

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 57. |

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.

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

If IBM Spectrum Protect Plus reports that a vSnap server is reaching its storage capacity, the vSnap storage pool must be expanded. To expand a vSnap storage pool, you must first add virtual or physical disks on the vSnap server, either by adding virtual disks to the vSnap virtual machine or adding physical disks to the vSnap physical server. See the vSphere documentation for information about creating additional virtual disks.


Before you begin

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

Procedure

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

1. In the navigation pane, click **System Configuration > Backup Storage > Disk**.

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

Establishing a replication partnership for a vSnap server



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

Before you begin

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

Procedure

To establish a replication partnership, complete the following steps:

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

What to do next


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

| Action | How to |
|---|---|
| Select the Backup Storage Replication option in the SLA policy that is associated with the backup job. | See “Creating an SLA policy” on page 79 |

Changing offload throughput rate

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

Procedure

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

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

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




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

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

Note: To clear a timeslot, click the timeslot. The scheduled selections are listed underneath the schedule table.

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

vSnap server administration reference

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

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. See [“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 47](#)

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 50](#)

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. If job logs report that a vSnap server is reaching its storage capacity, additional disks can be added to the vSnap pool. Alternatively, creating new SLA policies will force backups to use an alternate vSnap.

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

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

Managing encryption

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

vSnap encryption utilizes the following algorithm:

Cipher name

Advanced Encryption Standard (AES)

Cipher mode

xts-plain64

Key

256 bits

Linux Unified Key Setup (LUKS) header hashing

sha256

Managing encryption keys

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

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

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

Detecting disks

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

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

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

Showing disks

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

The USED AS column in the output shows whether each disk is in use. Any disk that is unformatted and unpartitioned is marked as unused, otherwise they are marked as used by the partition table or file system that is discovered on them.

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

Showing storage pool information

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

Creating a storage pool

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

To complete an advanced initialization, use the `vsnap pool create` command to create a storage pool manually. Before you run the command, ensure that one or more unused disks are available as described in [“Showing disks” on page 58](#). 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 58.

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

Command line: Run the `vsnap pool expand` command. For information about available options, pass the `--help` flag for any command or subcommand.

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

Network management

Configure and administer network services for a vSnap server.

Showing network interface information

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

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

mgmt

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

nfs

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

iscsi

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

smb

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

repl

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

Modifying services associated with network interfaces

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

The following options are required:

--id <id>

Enter the ID of the interface to update.

--services <services>

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

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

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

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 56.
2. Run the following commands:

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

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

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

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

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

Results

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

Chapter 4. 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.

As shown in the following table, the initial installation and configuration tasks are completed by the IBM Spectrum Protect Plus *infrastructure administrator*. By default, the `admin` user account is created for use by the infrastructure administrator.

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

| Action | Owner | |
|---|--|---|
| 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 66 | 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 68 | 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 secondary backup storage, which are optional features. For information about how to set up these features in the SLA policy, see [“Creating an SLA policy” on page 79](#).

Backup copies of data are called snapshots.

Procedure

To create an SLA policy, complete the following steps:

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

Retention

Specify the retention period for the backup snapshots.

Disable Schedule

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

Frequency

Enter the frequency for backup operations.

Start Time

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

Target Site

Select the target backup site.

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 6. 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 68.

Related concepts

[“Replicate backup-storage data ”](#) on page 8

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.

[“Managing SLA policies for backup operations”](#) on page 79

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 253.

To create custom user roles and resource groups, see [“Creating a resource group”](#) on page 246 and [“Creating a role”](#) on page 250

Procedure

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

1. In the navigation pane, click **Accounts > User**.
2. Click **Add User**. The **Add User** pane is displayed.
3. Click **Select the type of user or group you want to add > Individual new user**.
4. Enter a name and password for the application administrator.
5. In the **Assign Role** section, select **VM Admin**.

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

The screenshot shows the 'Add User - User Information and Role' form. The navigation pane on the left is expanded to 'Accounts > User'. The form has a dark blue header with the title 'Add User - User Information and Role'. Below the header, there is a dropdown menu for 'Select the type of user or group you want to add' with 'Individual new user' selected. The 'Username' field contains 'vmadmin'. The 'Password' field is masked with dots and has a 'Show' toggle. Below the password field, there is a note: 'Password must contain at least 8 characters.' The 'ASSIGN ROLE' section has several checkboxes: 'Application Admin', 'Backup Only', 'Restore Only', 'SYSADMIN', 'Self Service', and 'VM Admin' (which is checked). The 'PERMISSION GROUPS' section has two expandable items: 'Certificate' and 'Cloud'. At the bottom of the form, there are 'Cancel' and 'Continue >' buttons.

Figure 7. 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.

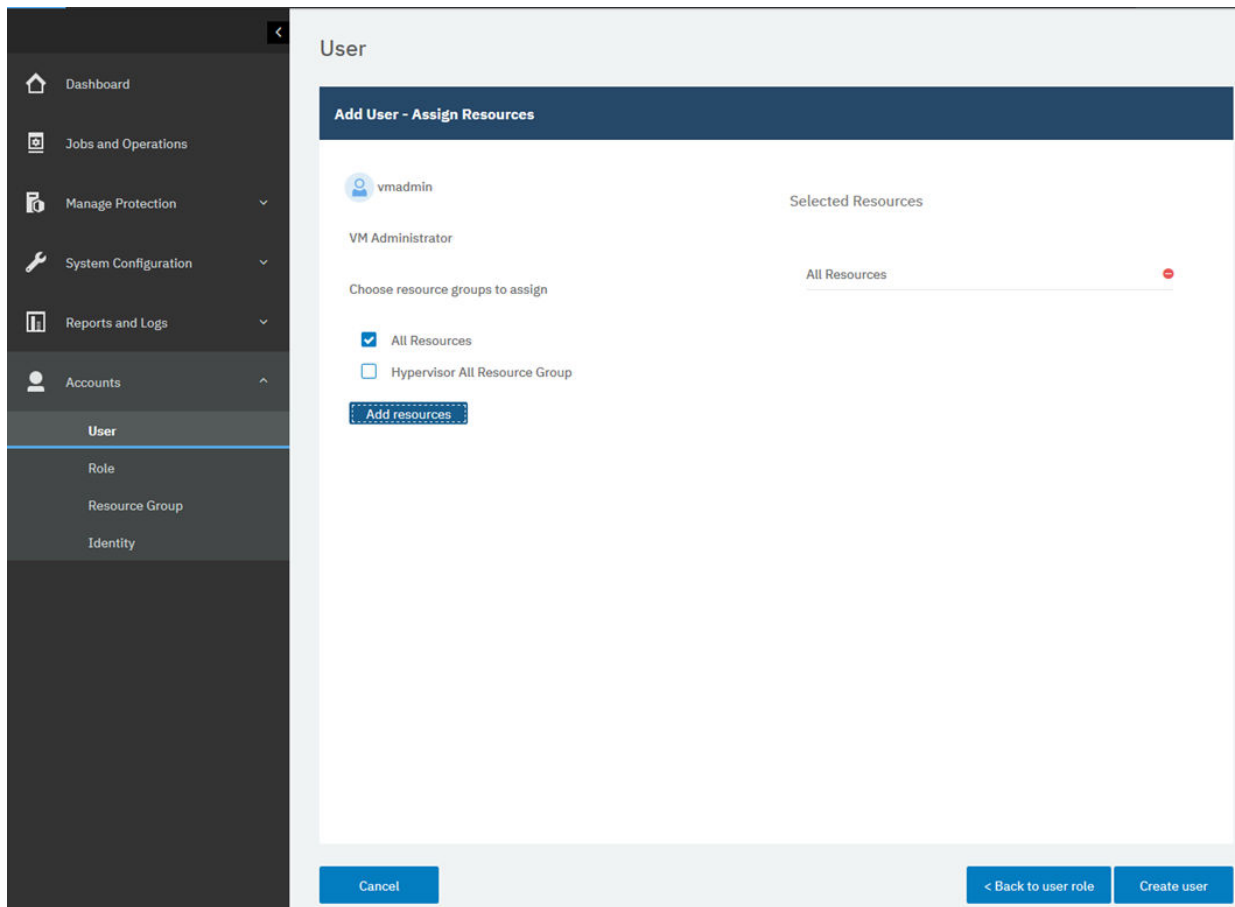


Figure 8. Selecting a resource group for the user account

8. Click **Create user**.

Related concepts

[“Managing user access” on page 245](#)

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 83](#) and [Chapter 8, “Protecting applications,” on page 119](#).

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.

4. In the **Options** section, configure the following option:

Maximum number of VMs to process concurrently per ESX server and per SLA

Set the maximum number of concurrent VM snapshots to process on the ESX server.

The following example shows populated fields.

The screenshot displays the 'Manage vCenter' configuration page in IBM Spectrum Protect Plus. The left sidebar shows the navigation menu with 'Backup' highlighted. The main content area is titled 'Backup' and includes a 'Manage vCenter' button. Below this, the 'vCenter Properties' section contains the following fields: Hostname/IP (192.0.2.0), Use existing user (unchecked), Username (admin_192.0.2.0), Password (masked with dots), Port (443), and a checked 'Use SSL' checkbox. The 'Options' section features a field for 'Maximum number of VM's to process concurrently per ESX server and per SLA' set to 3. At the bottom of the form are 'Cancel' and 'Save' buttons.

Figure 9. Adding a vCenter Server instance

5. Click **Save**.

IBM Spectrum Protect Plus confirms a network connection, adds the resource to the database, and then catalogs the resource. If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to verify and possibly fix the connections.

Add resources to a job definition

Before you can back up a resource, you must create a job definition that associates the resource with one or more backup policies, also referred to as SLA policies.

About this task

For example purposes, this task describes how 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 83 and [Chapter 8, “Protecting applications,”](#) on page 119.

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:

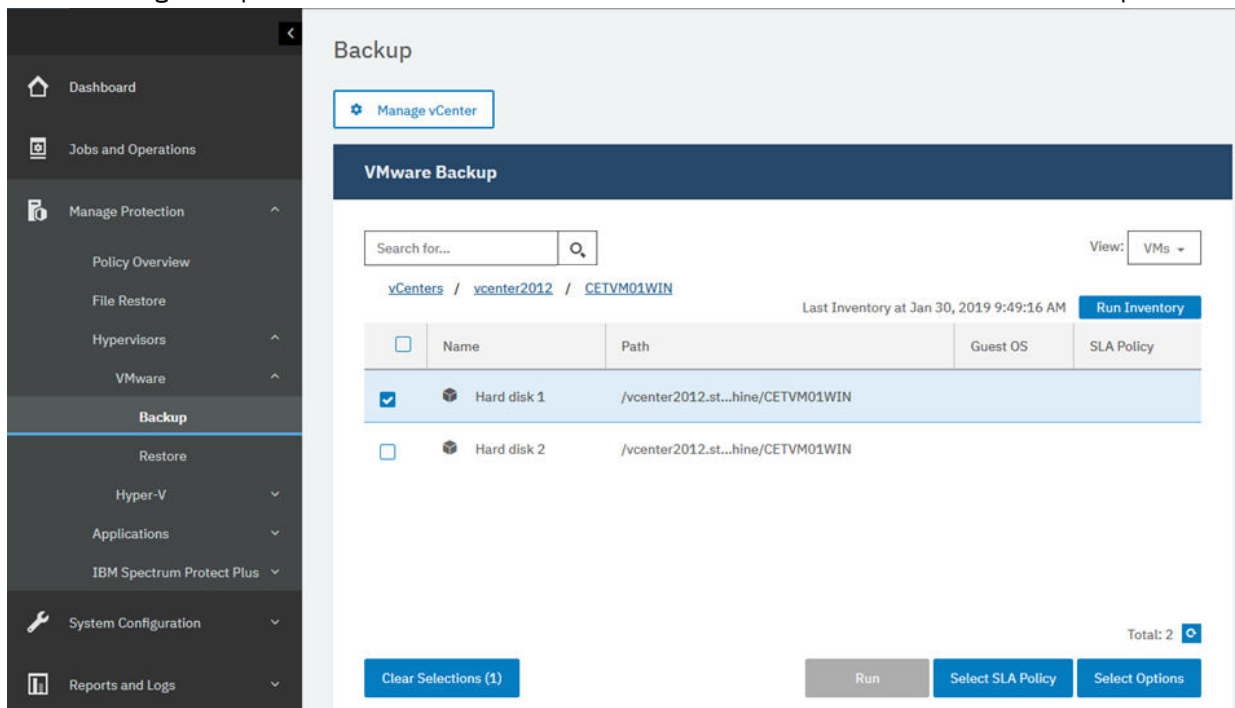


Figure 10. Selecting resources for backup

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

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

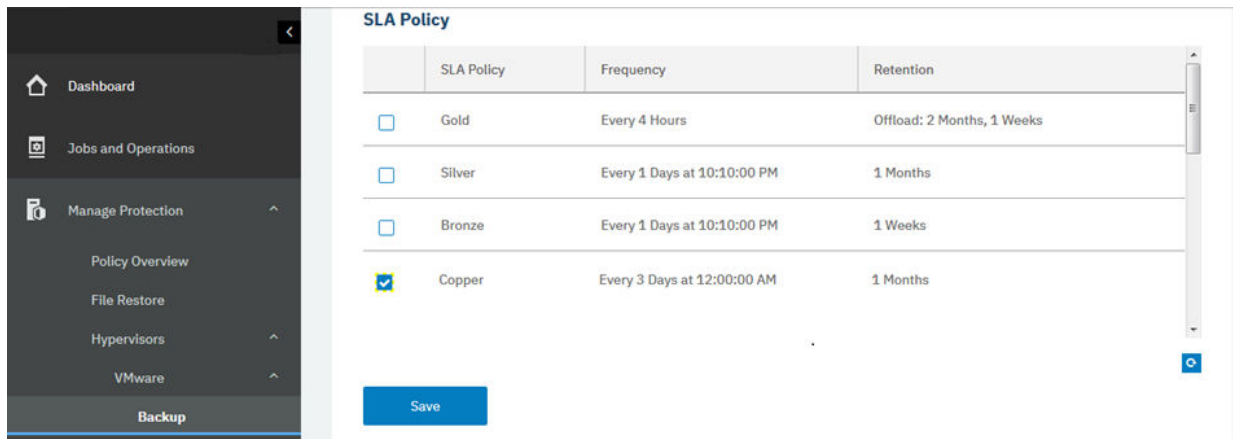


Figure 11. Selecting an SLA policy

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 91.
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 94.

Results

The job runs as defined by the SLA policies that you selected, or you can manually run the job by clicking **Jobs and Operations** and then clicking the **Policy and Job List** tab. For instructions, see [“Start a job”](#) on page 69.

Related concepts

[“Protecting IBM Spectrum Protect Plus”](#) on page 201

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** and then click the **Policy and Job List** tab.
2. Click the **Actions** menu that is associated with the job that you want to start and click **Start** as shown in the following example:

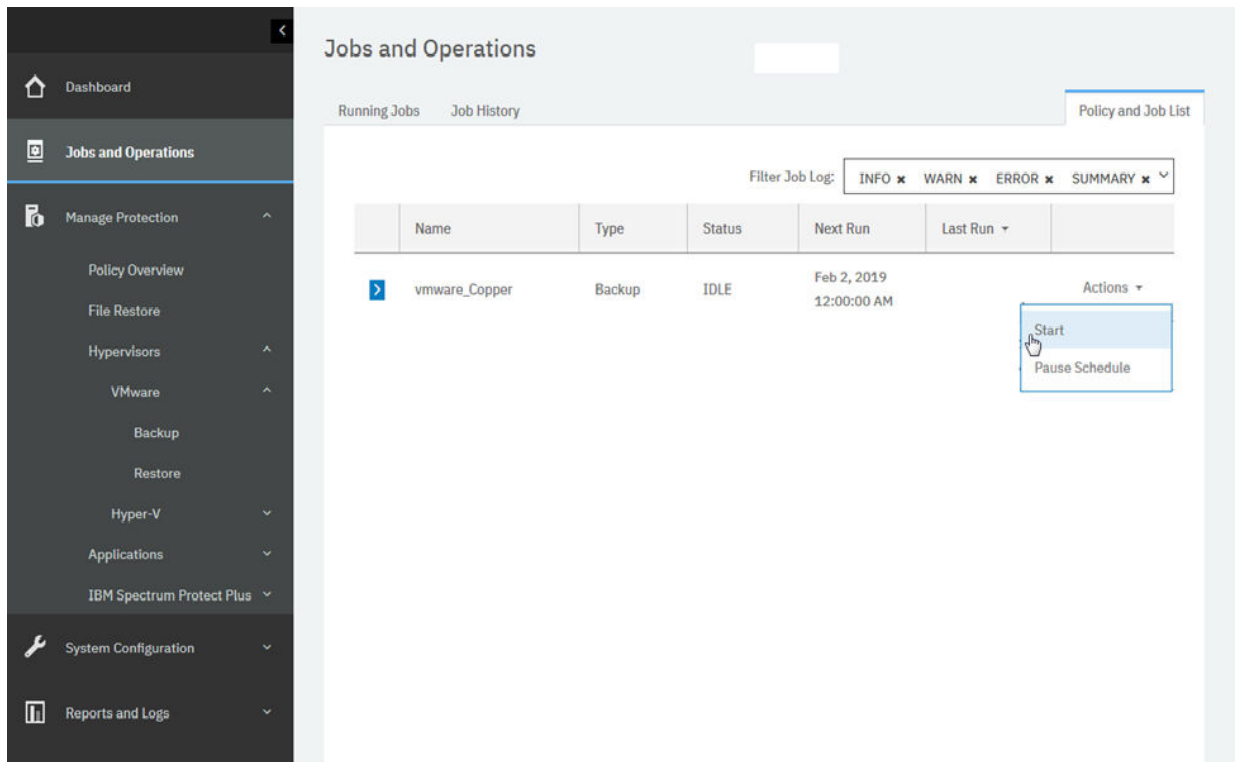


Figure 12. Starting a job

What to do next

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

Related concepts

[“Managing data protection jobs” on page 205](#)

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:

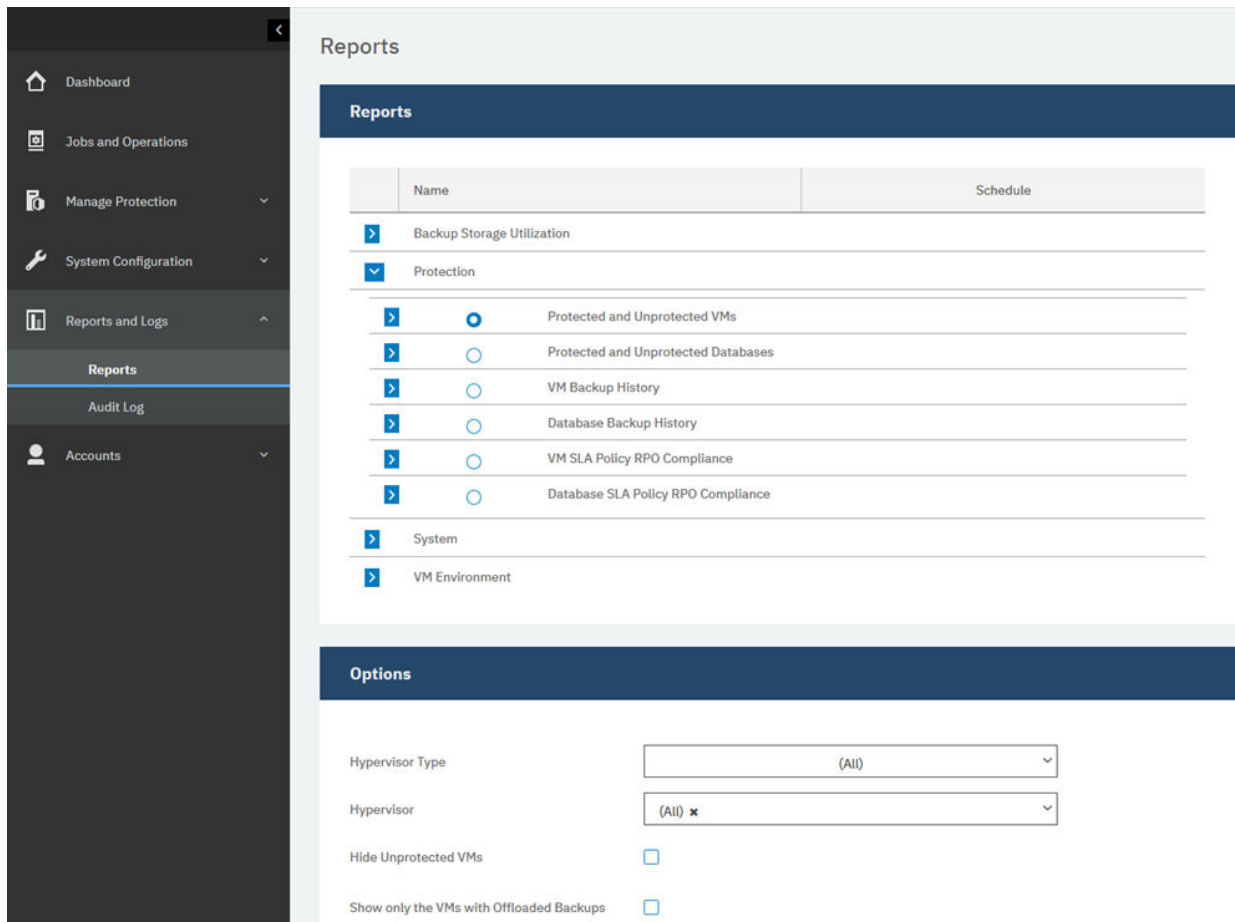


Figure 13. Selecting a report to run

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

- To run the report with custom parameters, set the parameters in the **Options** section, and click **Run**. Parameters are unique to each report.
- To run the report with default parameters, click **Run**.

Related concepts

[“Managing reports and logs” on page 237](#)

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

For information about available update files and how to obtain them from an IBM download site, see [technote 743897](#).

Before you update IBM Spectrum Protect Plus components, review the hardware and software requirements for the components to confirm any changes that might have occurred from previous versions.

Review the following restrictions and tips:

- You must separately update vSnap servers that are not on IBM Spectrum Protect Plus virtual appliances.
- The update process through the administrative console updates IBM Spectrum Protect Plus features and the underlying infrastructure components including the operating system and file system. Do not use another method to update these components.
- Do not update any of the underlying components for IBM Spectrum Protect Plus unless the component is provided in an IBM Spectrum Protect Plus update package. Infrastructure updates are managed by IBM update facilities. The administrative console is the primary means for updating IBM Spectrum Protect Plus features and underlying infrastructure components including the operating system and file system.

Take the following actions:

- Before you update components, it is important that you back up your IBM Spectrum Protect Plus environment as described in [“Backing up the IBM Spectrum Protect Plus application”](#) on page 201.
- After IBM Spectrum Protect Plus is updated, it cannot roll back to a previous version without a virtual machine snapshot. Create a virtual machine snapshot of your environment before you update IBM Spectrum Protect Plus. If you later want to roll back IBM Spectrum Protect Plus to an earlier version, you must have a virtual machine snapshot. After the upgrade is completed successfully, remove the virtual machine snapshot.

Updating the IBM Spectrum Protect Plus virtual appliance

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

Before you begin

Updating your IBM Spectrum Protect Plus environment to Version 10.1.3 is supported only from V10.1.2. If you are using V10.1.1, you must update to V10.1.2 and then update to V10.1.3. For instructions about how to update from V10.1.1, see [Updating the IBM Spectrum Protect Plus virtual appliance to version 10.1.2](#).

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

1. Ensure that you have backed up your IBM Spectrum Protect Plus environment as described in [“Backing up the IBM Spectrum Protect Plus application”](#) on page 201.
2. Download the prerequisite IBM Spectrum Protect Plus update named CNZF6EN . iso to a directory on the computer that is running the browser for the administrative console.

Tip: The file name CNZF6EN . iso might differ based on download source.

3. Ensure that no jobs are running during the update procedure. Pause the schedule for any jobs that have a status of IDLE or COMPLETED.

For a list of download images, including the required operating system update for the virtual appliance, see [technote 743897](#).

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. The default user name is <code>admin</code> and password is <code>password</code> . |
| System (recommended) | 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 **Manage updates**.
4. Click **Browse** to browse for the prerequisite update file named `CNZF6EN.iso`, and then click **Upload Update Image**.

The update process begins once the update image is uploaded to the appliance.

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

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

5. Clear the browser cache.
HTML content from previous versions of IBM Spectrum Protect Plus might be stored in the cache.
6. Start the updated version of IBM Spectrum Protect Plus.
7. Navigate to **Jobs and Operations** and then click the **Policy and Job List** tab.
8. Select **Release Schedule** from the **Actions** list for the paused jobs that are associated with the vSnap server.

Related tasks

[“Updating vSnap servers” on page 75](#)

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

Updating your vSnap servers to version 10.1.3 is supported only from version 10.1.2. If you are using version 10.1.1, you must update to version 10.1.2 and then update to version 10.1.3. For instructions about how to update to version 10.1.2, see [Updating vSnap servers to version 10.1.2](#).

Test restore jobs need to be completed prior to initiating an update to vSnap. Jobs that are not completed or not canceled when an upgrade is initiated will not be visible once the update has completed. If jobs are not visible once the update has completed, re-run test restore jobs.

You might also be required to update the operating system for the vSnap servers prior to updating the servers. For operating system requirements, see [“Component requirements” on page 11](#).

To check the current version and operating system for your vSnap servers, complete the following steps:

1. Log on to the vSnap server as the `serveradmin` user. If you are using IBM Spectrum Protect Plus 10.1.1, log in by using the root account.
2. To check the vSnap server version and operating system, use the vSnap command-line interface to issue the following command:

```
vsnap system info
```

Ensure that no jobs that use the vSnap server are running during the update procedure. Pause the schedule for any jobs that have a status of IDLE or COMPLETED.

Updating the operating system for a physical vSnap server

If you have installed the vSnap server on a machine that is running Red Hat Enterprise Linux, you must update the operating system to version 7.4 or 7.5 before you update the vSnap server. Linux v7.5 is required to use the IBM Spectrum Protect Plus storage offload feature. For instructions about how to update the operating system, see the Red Hat Enterprise Linux documentation.

Related tasks

[“Updating a vSnap server” on page 76](#)

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

Updating the operating system for a virtual vSnap server

If the operating system is CentOS Linux version 7.4 or earlier, you must update the operating system before you update the vSnap server. To update the operating system, follow the instructions in [Updating vSnap servers to version 10.1.2](#). The version 10.1.2 installation includes CentOS Linux version 7.5.

Related tasks

[“Updating a vSnap server” on page 76](#)

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

Updating a vSnap server

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

Before you begin

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

1. Ensure that you have backed up your IBM Spectrum Protect Plus environment as described in [“Backing up the IBM Spectrum Protect Plus application ” on page 201](#).
2. If you are updating from IBM Spectrum Protect Plus 10.1.1, you must update to version 10.1.2 and then update to version 10.1.3. For instructions about how to update to version 10.1.2, see [Updating vSnap servers to version 10.1.2](#).
3. Download the vSnap update file CNZF4EN .run and copy it to a temporary location on the vSnap server. For information about downloading files, see [technote 743897](#).

Procedure

To update a vSnap server, complete the following steps:

1. Log on to the vSnap server as the **serveradmin** user.
2. From the directory where the CNZF4EN .run file is located, make the file executable and run the installer by issuing the following commands:

```
chmod +x CNZF4EN.run
```

```
sudo ./CNZF4EN.run
```

The vSnap packages are installed.

3. Start the updated version of IBM Spectrum Protect Plus.
4. Navigate to **Jobs and Operations** and then click the **Policy and Job List** tab.
5. Select **Release Schedule** from the **Actions** list for the jobs that are associated with the vSnap server.

Updating VADP proxies

Updating the IBM Spectrum Protect Plus virtual appliance automatically updates all the VADP proxies that are associated with the virtual appliance. In rare scenarios such as loss of network connectivity, you must update the VADP proxy manually.

Before you begin



Before you begin, ensure that you have backed up your IBM Spectrum Protect Plus environment as described in [“Backing up the IBM Spectrum Protect Plus application ” on page 201](#).

Procedure

If a VADP proxy update is available for external proxies during a restart of the IBM Spectrum Protect Plus virtual appliance, the update will be automatically applied to any VADP proxy associated with an identity. To associate a VADP proxy with an identity, navigate to **System Configuration > VADP Proxy**. Click the options icon **⋮** and select **Set Options**. Through the User setting, select a previously entered username and password for the VADP proxy server.

To update a VADP proxy manually, complete the following steps:

1. Navigate to the **System Configuration > VADP Proxy** page in IBM Spectrum Protect Plus.

2. The **VADP Proxy** page displays each proxy server. If a newer version of the VADP proxy software is available, an update icon  displays in the **Status** field.
3. Ensure that there are no active jobs that use the proxy, and then click the update icon .

The proxy server enters a suspended state and installs the latest update. When the update completes, the VADP proxy server automatically resumes and enters an enabled state.

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

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

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

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

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

3. Set the appropriate permissions on the file.

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

What to do next

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

| Action | How to |
|----------------------------|---|
| Run the VMware backup job. | <p>See “Backing up VMware data” on page 91.</p> <p>The proxies are indicated in the job log by a log message similar to the following text:</p> <p>Run remote vmdkbackup of MicroService: <a href="http://<proxy>">http://<proxy> <i>nodename</i>, IP:<i>proxy_IP_address</i></p> |

Related tasks

“Creating VADP proxies” on page 96

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

Applying early availability updates

Early availability updates provide fixes for authorized program analysis reports (APARs) and minor issues between IBM Spectrum Protect Plus releases. These updates are available in bundles from the Fix Central Online website.

About this task

Early availability updates might not contain fixes for all IBM Spectrum Protect Plus components.

For instructions about how to obtain and install interim fixes, see the download information that is published when the fixes are available.

Chapter 6. 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.

If you edit an existing SLA policy by changing the cloud offload source, offload destination type, or target offload server options, associated jobs will start a full base backup, not an incremental backup, during the next job run.

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.

If a snapshot replication task is started before an initial backup to a vSnap server is completed, errors in the job log indicate that no recovery points exist for the database. After the initial backup to the vSnap server is completed, run the replication task again to replicate the snapshots as configured in the SLA policy.

Procedure

To create an SLA policy, complete the following steps:

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

Retention

Specify the retention period for the backup snapshots.

Disable Schedule

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

Frequency

Enter a frequency for backup operations.

Start Time

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

Target Site

Select 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 55.

Backup Storage Replication

Select this option to enable replication.

Disable Schedule

Select this check box to create the replication relationship without defining a frequency or start time.

Frequency

Enter a frequency for replication operations.

Start Time

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

Target Site

Select 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, set the following options to offload data to cloud storage or to a repository server. Data is backed up to the vSnap server for short term protection, and then offloaded to the selected cloud storage or repository server for longer-term protection. During the first offload of a backup volume, the snapshot is backed up in full. After the first offload of the base snapshot is completed, subsequent offloads are incremental and capture cumulative changes since the last offload. Cloud or repository server restore operations can be performed from any available vSnap server.

Offload to Cloud Storage

Select this option to offload data to cloud storage or to a repository server.

Disable Schedule

Select this check box to create the offload relationship without defining a frequency or start time.

Frequency

Enter a frequency for offload operations.

Start Time

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

Same retention as source selection

Select this option to use the same retention policy for the cloud offload backup as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

Restriction: Offload retention options are disabled if a server that uses write once read many (WORM) retention is selected in the **Target Offload Server** field.

Offload source

Click the source for the offload operation:

Main Policy Destination

The source for the offload operation is the target site that is defined in the **Main Policy** section.

Replication Policy Destination

The source for the offload operation is the target site that is defined in the **Replication Policy** section.

Offload Destination Type

Click **Cloud Servers** or **Repository Servers**.

Target Offload Server

Click the server to which you want to offload data.

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 250 |
| Create a backup job definition that uses the SLA policy. | See the backup topics in Chapter 7, “Protecting hypervisors,” on page 83 and Chapter 8, “Protecting applications,” on page 119 . |

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 23.

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

Note: An NFS volume may be mounted to any number of datacenters that belong to the same vCenter. If an NFS volume is mounted on more than one datacenter, vCenter treats the same volume as two different datastores. IBM Spectrum Protect Plus treats this as a single datastore and combines all of the VMs and VMDKs residing on the datastore from all of the datacenters on which the datastore is mounted. Any SLA selection against this datastore will cause all of the VMs from the different datacenters to be backed up or restored in IBM Spectrum Protect Plus.

Adding a vCenter Server instance

When a vCenter Server instance is added to IBM Spectrum Protect Plus, an inventory of the instance is captured, enabling you to complete backup and restore jobs, as well as run reports.

Procedure

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

1. In the navigation pane, click **Manage Protection > Hypervisors > VMware > 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. 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 250. |

Related concepts

[“Managing identities”](#) on page 255

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 91

Use a backup job to back up VMware resources such as virtual machines, datastores, folders, vApps, and datacenters with snapshots.

[“Restoring VMware data”](#) on page 99

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

Virtual machine privileges

vCenter Server privileges are required for the virtual machines that are associated with a VMware provider. These privileges are included in the vCenter Administrator role.

If the user that is associated with the provider is not assigned to the Administrator role for an inventory object, the user must be assigned to a role that has the following required privileges. Ensure that the privileges are propagated to child objects. For instructions, refer to the VMware documentation about adding a permission to an inventory object.

| vCenter Server Object | Required Privileges |
|-----------------------|---|
| Alarm | <ul style="list-style-type: none"> • Acknowledge alarm • Set alarm status |

| vCenter Server Object | Required Privileges |
|------------------------------|---|
| Cryptographic Operations | <ul style="list-style-type: none"> • Add disks • Direct access • Encrypt • Encrypt new • Manage encryption policies |
| Datacenter | <ul style="list-style-type: none"> • Create datacenter • Reconfigure datacenter |
| Datastore | <ul style="list-style-type: none"> • Allocate space • Browse datastore • Configure datastore • Low level file operations • Remove file • Update virtual machine files |
| Datastore Cluster | <ul style="list-style-type: none"> • Configure a datastore cluster |
| Distributed switch | <ul style="list-style-type: none"> • Create • Delete • Host operation • Modify • Move • Network I/O Control operation • Policy operation • Port configuration option • Port setting operation • VSPAN operation |
| ESX Agent Manager | <ul style="list-style-type: none"> • Config • Modify • View |
| Extension | <ul style="list-style-type: none"> • Register extension |
| Folder | <ul style="list-style-type: none"> • Create folder • Delete folder • Move folder • Rename folder |

| vCenter Server Object | Required Privileges |
|-------------------------------------|---|
| Global | <ul style="list-style-type: none"> • Cancel task • Diagnostics (used for troubleshooting, not required for operations) • Disable methods • Enable methods • Licenses • Log event • Manage custom attributes • Set custom attribute • Settings |
| Host > Configuration | <ul style="list-style-type: none"> • Advanced settings • Storage partition configuration |
| Inventory Service > vSphere Tagging | <ul style="list-style-type: none"> • Assign or Unassign vSphere Tag • Create vSphere Tag • Create vSphere Tag Category • Modify UsedBy Field for Category • Modify UsedBy Field for Tag |
| Network | <ul style="list-style-type: none"> • Assign network • Configure • Move network • Remove |
| Resource | <ul style="list-style-type: none"> • Apply recommendation • Assign a vApp to resource pool • Assign virtual machine to resource pool • Create resource pool • Migrate powered off VM • Migrate powered on VM • Modify resource pool • Move resource pool • Query vMotion • Remove resource pool • Rename resource pool |
| Sessions | <ul style="list-style-type: none"> • View and stop sessions |
| Storage views | <ul style="list-style-type: none"> • Configure service • View |
| Tasks | <ul style="list-style-type: none"> • Create task • Update task |

| vCenter Server Object | Required Privileges |
|------------------------------------|--|
| Virtual Machine > Configuration | <ul style="list-style-type: none"> • Add existing disk • Add new disk • Add or remove device • Advanced • Change CPU count • Change resource • Configure managedBy • Disk change tracking • Disk lease • Display connection settings • Extend virtual disk • Host USB device • Memory • Modify device settings • Query Fault Tolerance compatibility • Query unowned files • Raw device • Reload from path • Remove disk (detach and remove virtual disk) • Rename • Reset guest information • Set annotation • Settings • Swapfile placement • Unlock virtual machine • Upgrade virtual machine compatibility |
| Virtual Machine > Guest Operations | <ul style="list-style-type: none"> • Guest Operation Modifications • Guest Operation Program Execution • Guest Operation Queries |

| vCenter Server Object | Required Privileges |
|-------------------------------|---|
| Virtual Machine > Interaction | <ul style="list-style-type: none"> • Answer question • Backup operation on virtual machine • Configure CD media • Configure floppy media • Console interaction • Create screenshot • Defragment all disks • Device connection • Disable Fault Tolerance • Enable Fault Tolerance • Guest operating system management by VIX API • Inject USB HID scan codes • Perform wipe or shrink operations • Power Off • Power On • Record session on VM • Replay session on VM • Reset • Resume Fault Tolerance • Suspend • Suspend Fault Tolerance • Test failover • Test restart Secondary VM • Turn Off Fault Tolerance • Turn On Fault Tolerance • VMware Tools install |
| Virtual Machine > Inventory | <ul style="list-style-type: none"> • Create from existing • Create new • Move • Register • Remove • Unregister |

| vCenter Server Object | Required Privileges |
|---|--|
| Virtual Machine > Provisioning | <ul style="list-style-type: none"> • Allow disk access • Allow read-only disk access • Allow virtual machine download • Allow virtual machine files upload • Clone template • Clone virtual machine • Create template from virtual machine • Customize • Deploy template • Mark as template • Mark as virtual machine • Modify customization specification • Promote disks • Read customization specifications |
| Virtual Machine > Service configuration | <ul style="list-style-type: none"> • Allow notifications • Allow polling of global event notifications • Manage service configurations • Modify service configurations • Query service configurations • Read service configurations |
| Virtual Machine > Snapshot management | <ul style="list-style-type: none"> • Create snapshot • Remove snapshot • Rename snapshot • Revert to snapshot |
| Virtual Machine > vSphere Replication | <ul style="list-style-type: none"> • Configure replication • Manage replication • Monitor replication |

| vCenter Server Object | Required Privileges |
|-----------------------|--|
| vApp | <ul style="list-style-type: none"> • Add VM to vApp • Assign resource pool to vApp • Assign vApp to another vApp • Clone • Create • Delete • Export • Import • Move • Power Off • Power On • Rename • Suspend • Unregister • View OVF Environment • vApp application configuration • vApp instance configuration • vApp managedBy configuration • 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 resources such as 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 83.
- Configure SLA policies. For more instructions, see [“Create backup policies”](#) on page 62.
- 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 245 and [“Managing user accounts”](#) on page 253.
- 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.
- When backing up VMware virtual machines, IBM Spectrum Protect Plus downloads .vmx, .vmxf, and .nvram files if necessary, and then it transfers those files to the vSnap as needed. For this to work successfully, the IBM Spectrum Protect Plus appliance must be able to resolve and access all protected ESXi hosts; and when communicating with an ESXi host, the correct IP address must be returned.
- A virtual machine is decommissioned when it expires based on the retention parameters in the SLA policy.
- 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** and then clicking the **Policy and Job List** tab.

When the job definition is saved, available virtual machine disks (VMDKs) in a virtual machine are discovered and are shown when **VMs and Templates** is selected in the **View** filter. By default, these VMDKs are assigned to the same SLA policy as the virtual machine. If you want a more granular backup operation, you can exclude individual VMDKs from the SLA policy. For instructions, see [“Excluding VMDKs from the SLA policy for a job”](#) on page 94.

5. To edit options before you create the job definition, click **Select Options**.
In the **Backup Options** section, set the following job definition options:

Skip Read-only datastores

Skip datastores that are mounted as read-only.

Skip temporary datastores mounted for Instant Access

Exclude temporary Instant Access datastores from the backup job definition.

VADP Proxy

Select a VADP proxy to balance the load.

Priority

Set the backup priority of the selected resource. Resources with a higher priority setting are backed up first in the job. Click the resource that you want to prioritize in the **VMware Backup** section, and then set the backup priority in the **Priority** field. Set 1 for the highest priority resource or 10 for the lowest. When a priority value is not set, automatically assigns a priority of 5 is set by default.

In the **Snapshot Options** section, set the following job definition options:

Make VM snapshot application/file system consistent

Enable this option to turn on application or filesystem consistency for the virtual machine snapshot. All VSS-compliant applications such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL, and the system state are quiesced. VMDKs and virtual machines can be instantly mounted to restore data that is related to quiesced applications.

VM Snapshot retry attempts

Set the number of times that IBM Spectrum Protect Plus attempts to capture an application or file-consistent snapshot of a virtual machine before the job is canceled. If the **Fall back to unquiesced snapshot if quiesced snapshot fails** option is enabled, an unquiesced snapshot will be taken after the retry attempts.

Fall back to unquiesced snapshot if quiesced snapshot fails

Enable to fall back to a non-application or non-file-system consistent snapshot if the application consistent snapshot fails. Selecting this option ensures that an unquiesced snapshot is taken if environmental issues prohibit the capture of an application or file-system consistent snapshot.

In the **Agent Options** section, set the following job definition options:

Truncate SQL logs

To truncate application logs for SQL Server during the backup job, enable the **Truncate SQL logs** option. The credentials must be established for the associated virtual machine by using the Guest OS user name and Guest OS Password option within the backup job definition. When the virtual machine is attached to a domain, the user identity follows the default *domain\name* format. If the user is a local administrator, the format *local_administrator* is used.

The user identity must have local administrator privileges. On the SQL Server server, the system login credential must have the following permissions:

- SQL Server sysadmin permissions must be enabled.
- The **Log on as a service** right must be set. For more information about this right, see [Add the Log on as a service Right to an Account](#).

IBM Spectrum Protect Plus generates log files for the log truncation function and copies them to the following location on the IBM Spectrum Protect appliance:

```
/data/log/guestdeployer/latest_date/latest_entry/vm_name
```

Where *latest_date* is the date that the backup job and log truncation occurred, *latest_entry* is the universally unique identifier (UUID) for the job, and *vm_name* is the 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 cloud resources or repository servers.

Catalog file metadata

Turn on file indexing for the associated snapshot. When file indexing is completed, individual files can be restored by using the **File Restore** pane in IBM Spectrum Protect Plus. Credentials must be established for the associated virtual machine by using an SSH key, or a **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: SSH Keys are not valid authorization mechanism for Windows platforms.

Note: File indexing and file restore are not supported from restore points that were offloaded to cloud resources or repository servers.

Exclude Files

Enter directories to skip when file indexing is performed. Files within these directories are not added to the IBM Spectrum Protect Plus catalog and are not available for file recovery. Directories can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*). Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *. Separate multiple filters with a semicolon.

Use existing user

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 **Policy Options** field that is associated with the job in the **SLA Policy Status** section. Set the additional policy options:

Pre-scripts and Post-scripts

Pre-scripts and post-scripts are scripts that can be run before or after a job runs. Windows-based machines support Batch and PowerShell scripts while Linux-based machines support shell scripts.

In the **Pre-script** or **Post-script** section, select an uploaded script and a script server where the script will run. Scripts and script servers are configured by using the **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

Exclude Resources

Exclude specific resources from the backup job by using single or multiple exclusion patterns. Resources can be excluded by using an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *.

Separate multiple filters with a semicolon.

Force Full Backup of Resources

Force base backup operations for specific virtual machines or databases in the backup job definition. Separate multiple resources with a semicolon.

9. To save any additional options that you configured, click **Save**.

What to do next

After you define a backup job, you can complete the following actions:

| Action | How to |
|--|--|
| If you are using a Linux environment, consider creating VADP proxies to enable load sharing. | See “Creating VADP proxies” on page 96 . |
| Create a VMware restore job definition. | See “Restoring VMware data” on page 99 . |

Related concepts

[“Configuring scripts for backup and restore operations” on page 208](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and Batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Starting jobs” on page 206](#)

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

Excluding VMDKs from the SLA policy for a job

After you save a backup job definition, you can exclude individual VMDKs in a virtual machine from the SLA policy that is assigned to job.

Before you begin

Excluding one or more VMDKs from a backup operation can impact the success of recovery. Consider the following scenarios before excluding a disk from a VM backup operation.

- For Instant Disk Restore, if a VMDK is selected for a restore operations, an existing VM is chosen as the destination. IBM Spectrum Protect Plus mounts the restored disk to the chosen destination VM.
- For Instant VM Restore, if the VMDK that was excluded during a backup contains data that is necessary to boot the virtual machine, then the restored VM may fail to boot.
- For VMs with Windows-based guests, the restored VM may fail to boot if the disk on which the main operating system is installed, typically the C: drive, was excluded during the backup operation.
- For VMs with Linux-based guests, the restored VM may fail:
 - If a disk containing the boot or root partition was excluded during backup.
 - If a disk containing a data (non-root) partition was excluded during backup, and the data volume did not have the 'nofail' option specified in /etc/fstab, then the restored VM may fail.

Procedure

To exclude VMDKs from the SLA policy:

1. In the navigation pane, click **Manage Protection > Hypervisors > VMware > 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}: Release of backup" >> ${log}
#execute release command
cmd="echo \"SELECT pg_stop_backup();\" | sudo /opt/vmware/vpostgres/9.4/bin/psql -U postgres >> ${log} 2>&1"
eval ${cmd}
#set and log end date
today=`date +%Y/%m/%d\ %H:%M:%S`
echo "${today}: Finished thaw script" >> ${log}
```

Managing VADP backup proxies

In IBM Spectrum Protect Plus, you can create proxies to run VMware backup jobs by using vStorage API for Data Protection (VADP) in Linux environments. The proxies reduce demand on system resources by enabling load sharing and load balancing. 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 251](#).

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.

Transport modes describe the method by which a VADP proxy moves data. The transport mode is set as a property of the proxy. Most backup and recovery jobs are later configured to use one or more proxies.

VADP proxies in IBM Spectrum Protect Plus support the following VMware transport modes: SAN, HotAdd, NBDSSL, and NBD.

Although every enterprise differs, and priorities in terms of size, speed, reliability, and complexity vary from environment to environment, the following general guidelines apply to the Transport Mode selection:

- SAN transport mode should be used in a direct storage environment because this mode is fast and generally reliable.
- HotAdd transport mode should be used if the VADP proxy is virtualized. This mode supports all vSphere storage types.
- NBD or NBDSSL transport mode (LAN) is the fallback mode because it works in physical, virtual, and mixed environments. However, with this mode, data transfer speed might be compromised if network connections are slow. NBDSSL mode is similar to NBD mode except that data transferred between the VADP proxy and the ESXi server is encrypted when using NBDSSL.

Creating VADP proxies

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

Before you begin

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.

The connection between the IBM Spectrum Protect Plus virtual appliance and a registered VADP proxy is a bidirectional connection that requires the IBM Spectrum Protect Plus virtual appliance to have connectivity to the VADP proxy, and the VADP proxy to have connectivity to the IBM Spectrum Protect Plus virtual appliance. To ensure a proper connection from the IBM Spectrum Protect Plus virtual appliance to the VADP proxy, verify that the IBM Spectrum Protect Plus virtual appliance can ping the VADP proxy by completing the following steps:

1. Connect to the command line for the IBM Spectrum Protect Plus virtual appliance by using the Secure Shell (SSH) network protocol.
2. Execute `ping <vadm_ip>`, where `<vadm_ip>` is the resolvable IP address of the VADP proxy.

If the ping fails, ensure that the IP address of the VADP proxy is resolvable and is addressable by the IBM Spectrum Protect Plus appliance and that a route exists from the IBM Spectrum Protect Plus appliance to the VADP proxy. If the ping succeeds, ensure there is a proper connection from the VADP proxy to the IBM Spectrum Protect Plus virtual appliance by performing the following procedure:

1. Connect to the command line for the VADP proxy by using Secure Shell (SSH) network protocol.
2. Execute `ping <spectrum_protect_plus_ip>`, where `<spectrum_protect_plus_ip>` is the resolvable IP address of the IBM Spectrum Protect Plus virtual appliance.

If the ping fails, ensure that the IP address of the IBM Spectrum Protect Plus virtual appliance is resolvable and is addressable by the VADP proxy. Ensure that a route exists from the VADP proxy to the IBM Spectrum Protect Plus virtual appliance.

What to do next

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

| Action | How to |
|----------------------------|---|
| Run the VMware backup job. | See “Backing up VMware data” on page 91 . The proxies are indicated in the job log by a log message similar to the following text: Run remote vmdbkbackup of MicroService: <code>http://<proxy></code> <code>nodename, IP:proxy_IP_address</code> |

Related tasks

[“Setting options for VADP proxies” on page 98](#)

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

Setting options for VADP 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 VMs that a proxy can process. If **Use All Resources** is selected, the number of CPUs on the proxy determines the task limit based on the following formula:

1 CPU = 1 VMDK

A CPU is the smallest hardware unit capable of executing a thread. The number of CPUs on a proxy is determined by using the `lscpu` command.

What to do next

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

| Action | How to |
|--|---|
| Run the VMware backup job. | See “Backing up VMware data” on page 91 . The proxies are indicated in the job log by a log message similar to the following text: Run <code>remote vmdkbackup of MicroService:</code> <code>http://<proxy</code> <code>nodename, IP:proxy_IP_address</code> |
| 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 proxies” on page 96](#)

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

Uninstalling VADP proxies

You can remove a VADP proxies from your IBM Spectrum Protect Plus environment.

Procedure

To uninstall VADP proxies from your IBM Spectrum Protect Plus, complete the following steps:

1. From a command prompt, navigate to the directory `/opt/IBM/SPP/uninstall` on the proxy host system.
2. Run the following command:
`./uninstall_vmdkbackup`

Restoring VMware data

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

Before you begin

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 91](#).
- 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 245](#) and [“Managing user accounts” on page 253](#).
- 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.
- 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.
- Windows file indexing and file restore on volumes residing on dynamic disks is not supported.

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 104](#).

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, click the add to restore list icon  at the resource level. Select **Restore by site** or **Restore by cloud/repository server**. If restoring from a Site, click the **Select a site** drop-down menu to choose a site associated with the backup storage server you want to restore from. If restoring from a cloud or repository server, the restore source will be automatically selected.

To select a restore point that is not the latest, expand a resource in the **Restore** pane, and then click the add to restore list icon  that is associated with the restore point. Adding a combination of latest restore points and non-latest restore points to the Restore List is not supported. Click the delete icon  to remove restore points from the **Restore List**.

Additional filtering options are available when viewing non-latest restore points. To view available restore points from sites, cloud resources, or repository servers, expand a resource in the **Restore** pane, then select the source type through the **Filters** menu. Once a source type is selected, all available restore points associated with the source type display. To view restore points on a specific

resource, select it from the drop-down menu adjacent to the source type drop-down menu. For example, if the filter is set to **Sites**, click **Show restore points in all sites** to select a specific site.

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 this option to bypass vCenter Server and to restore data directly to an ESXi host. In other restore scenarios, actions are completed through vCenter Server. If vCenter Server is unavailable, this option restores the virtual machine or virtual machines that contain the components that vCenter Server is dependent on.

When you select an ESXi host, you must specify the host user. You can select an existing user for the host or create a new one.

To create a user, enter a user name, the user ID, and the user password.

If the ESXi host is attached to a domain, the user ID follows the default *domain\name* format. If the user is a local administrator, use the *local_administrator* format.

To restore to an ESXi host, the host must have a standard switch or a distributed switch with ephemeral binding. Review the information in [“Restoring data when vCenter Server or other management VMs are not accessible”](#) on page 105 to ensure that you have the correct environment configured to use this option.

Alternate vSnap

When restoring from a restore point that was offloaded to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server used to complete the restore is the same vSnap server used to complete the backup and offload operations. To reduce load, an alternate vSnap server can be selected to serve as the gateway to complete the restore. To select an alternate vSnap server when restoring a specific, non-latest restore point from a cloud resource or repository server, select **Use alternate vSnap server for the restore job**, then select a server from the **Select alternate vSnap** menu.

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. 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 **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a nonzero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a nonzero return code, the post-script task status is reported as COMPLETED.

When this option is 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:

Make IA clone resource permanent

Enable this option to move the virtual disk to permanent storage and clean up the temporary resources. This is accomplished by kicking off a vMotion operation for the resources in the background. The destination of the vMotion operation is the VM Configuration Datastore. The Instant Access disk is still available for read/write operations while this operation is being performed.

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.

Restore VM tags

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

Append suffix to virtual machine name

Enter a suffix to add to the name of restored virtual machines.

Prepend prefix to virtual machine name

Enter a prefix to add to the name of restored virtual machines.

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 83](#)

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 99](#).

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 83](#)

When a vCenter Server instance is added to IBM Spectrum Protect Plus, an inventory of the instance is captured, enabling you to complete backup and restore jobs, as well as run reports.

Restoring data when vCenter Server or other management VMs are not accessible

IBM Spectrum Protect Plus provides an option to automatically restore data by using an ESXi host if vCenter Server or one of the components that it uses are not accessible. This option restores the virtual machines that contain the components that vCenter Server uses.

Before you begin

To complete this procedure, you must be familiar with the ESXi and vCenter Server user interfaces.

About this task

vCenter Server uses the following components:

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

To use the **ESX host if vCenter is down** option, the ESXi host must have a standard switch or a distributed switch. The distributed switch must have ephemeral binding. If one or both of these switches

are available, you can run a restore operation in IBM Spectrum Protect Plus with the option enabled as described in [“Restoring VMware data” on page 99](#) and no further manual configuration is required.

If neither of these switches is available, you must complete the following steps before you can use the **ESX host if vCenter is down** option.

Procedure

1. Connect to the destination ESXi host user interface and create a standard virtual switch.
The new switch has no port groups or uplinks.
2. Use the Secure Shell (SSH) protocol to connect to the ESXi host.
3. List the distributed switches that are configured on the ESXi host by issuing the following command:

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

4. Identify the physical network interface card (NIC) and the port group of the distributed switch that you want to use for the restore operation.
5. Remove the physical NIC and port group from the distributed switch by issuing the following command:

```
#esxcfg-vswitch -Q physical_vnic -V port_group switch_name
```

6. Add the physical NIC and port group to the new standard switch by issuing the following command:

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

7. In the ESXi host user interface, add a temporary port group and select the standard switch that you created in step [“1” on page 106](#).
The standard switch has one port group and one uplink.
8. Run a restore operation in IBM Spectrum Protect Plus with the **ESX host if vCenter is down** option enabled.
For instructions about running a restore operation, see [“Restoring VMware data” on page 99](#).
9. In the ESXi host user interface for the ESXi host, power on the VMs that are restored.
10. Log in to the vCenter Server user interface and start the migration of the management VMs from the temporary port group that you created in step [“7” on page 106](#) to an available distributed port group.
11. After all of the VMs are migrated to the original port group, reincorporate the physical NIC and the port group into the original distributed switch by taking the following actions. For example purposes, the following commands reference a virtualized Network Interface Card (vNIC) named `vmnic0` that is part of port group 64.
 - a. Remove the network cards (known as vmnics) from a standard switch by issuing the following command:

```
#esxcli network vswitch standard uplink remove --uplink-name=vmnic --vswitch-name=vSwitch
```

For example:

```
#esxcli network vswitch standard uplink remove --uplink-name=vmnic0 --vswitch-name=vered_recovery
```

- b. Add network cards to the distributed switch by issuing the following command:

```
#esxcfg-vswitch -P vmnic -V unused_distributed_switch_port_ID distributed_switch
```

For example:

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

12. Delete the temporary port group and the standard switch from the ESXi host user interface.

13. After the VMs are migrated and accessible, use the ESXi host user interface to unregister, but not delete, the old VMs if the original host is reachable.

By using this method, you avoid creating duplicated information such as names, Media Access Control (MAC) addresses, operating system level IDs, and VM Universal Unique Identifiers (UUIDs). You must complete this step even if you are using a new datastore.

In some vSphere or ESXi versions, the unregister operation can be completed by using the **Remove from inventory** option. This option unregisters a VM from the vCenter Server catalog, but leaves VMDK files on the datastore where the files consume storage space. After you have fully recovered the VM and the environment is successfully running, you can regain the space by manually removing these files from the datastore.

Backing up and restoring Hyper-V data

To protect Hyper-V data, first add Hyper-V servers in IBM Spectrum Protect Plus, and then create jobs for backup and restore operations for the content of the servers.

Ensure that your Hyper-V environment meets the system requirements in [“Hypervisor requirements ”](#) on page 23.

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 SSL connection, you must add the self-signed SSL certificate for the Hyper-V server or a certificate authority (CA) certificate. To upload a certificate, see [“Uploading an SSL certificate from the administrative console” on page 229](#).

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 108](#).

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 250 . |

Related tasks

[“Backing up Hyper-V data” on page 109](#)

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

[“Restoring Hyper-V data” on page 112](#)

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 back up 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 107](#)
- Configure SLA policies. For instructions, see [“Create backup policies” on page 62](#).
- 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 245](#) and [“Managing user accounts” on page 253](#).
- 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 policy 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** and then clicking the **Policy and Job List** tab.

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

In the **Backup Options** section, set the following job definition options:

Skip Read-only datastores

Enable to skip datastores mounted as read-only.

Skip temporary datastores mounted for Instant Access

Enable to exclude temporary Instant Access datastores from the backup job definition.

Priority

Set the backup priority of the selected resource. Resources with a higher priority setting are backed up first in the job. Click the resource that you want to prioritize in the **VMware Backup** section, and then set the backup priority in the **Priority** field. Set 1 for the highest priority resource or 10 for the lowest. When a priority value is not set, automatically assigns a priority of 5 is set by default.

In the **Snapshot Options** section, set the following job definition options:

Make VM snapshot application/file system consistent

Enable this option to turn on application or filesystem consistency for the virtual machine snapshot.

VM Snapshot retry attempts

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

In the **Agent Options** section, set the following job definition options:

Truncate SQL logs

To truncate application logs for SQL during the Backup job, enable the **Truncate SQL logs** option. Note that credentials must be established for the associated virtual machine through the Guest OS Username and Guest OS Password option within the backup job definition. The user identity follows the default *domain\name* format if the virtual machine is attached to a domain. The format *local_administrator* is used if the user is a local administrator.

The user identity must have local administrator privileges. Additionally, on the SQL server, the system login credential must have SQL sysadmin permissions enabled, as well as the **Log on as a service** right. For more information about this right, see [Add the Log on as a service Right to an Account](#).

IBM Spectrum Protect Plus generates logs pertaining to the log truncation function and copies them to the following location on the IBM Spectrum Protect Plus appliance:

```
/data/log/guestdeployer/latest_date/latest_entry/vm_name
```

Where *latest_date* is the date that the backup job and log truncation occurred, *latest_entry* is the universally unique identifier (UUID) for the job, and *vm_name* is the hostname or IP address of the VM where the log truncation occurred.

Restriction: File indexing and file restore are not supported from restore points that were offloaded to an IBM Spectrum Protect server.

Catalog file metadata

To turn on file indexing for the associated snapshot, enable the Catalog file metadata option. After file indexing is complete, individual files can be restored by using the **File Restore** pane in IBM Spectrum Protect Plus. Note that credentials must be established for the associated virtual machine by using an SSH key, or a Guest OS Username and Guest OS Password option in the backup job definition. Ensure that the virtual machine can be accessed from the IBM Spectrum Protect Plus appliance either by using DNS or hostname. Note that SSH keys are not a valid authorization mechanism for Windows platforms.

Exclude Files

Enter directories to skip when file indexing is performed. Files within these directories are not added to the IBM Spectrum Protect Plus catalog and are not available for file recovery. Directories can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*). Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *. Separate multiple filters with a semicolon.

Use existing user

Enable to select a previously entered username and password for the provider.

Guest OS Username/Password

For some tasks (such as cataloging file metadata, file restore, and IP reconfiguration), credentials must be established for the associated virtual machine. Enter the username and password, and ensure that the virtual machine can be accessed from the IBM Spectrum Protect Plus appliance either through DNS or hostname.

The default security policy uses the Windows NTLM protocol, and the user identity follows the default *domain\name* format if the Hyper-V virtual machine is attached to a domain. The format *local_administrator* is used if the user is a local administrator.

6. To troubleshoot a connection to a hypervisor virtual machine, use the **Test** function.

The **Test** function verifies communication with the virtual machine and tests DNS settings between the IBM Spectrum Protect Plus appliance and the virtual machine. To test a connection, select a single virtual machine, then click **Select Options**. Select **Use existing user** and select a previously entered user name and password for the resource. 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 **Policy Options** field that is associated with the job in the **SLA Policy Status** section. Set the additional policy options:

Pre-scripts and Post-scripts

Pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job level. Windows-based machines support Batch and PowerShell scripts while Linux-based machines support shell scripts.

In the **Pre-script** or **Post-script** section, select an uploaded script and a script server where the script will run. Scripts and script servers are configured on the **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

Exclude Resources

Exclude specific resources from the backup job through single or multiple exclusion patterns. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *.

Separate multiple filters with a semicolon.

Force Full Backup of Resources

Force base backup operations for specific virtual machines or databases in the backup job definition. Separate multiple resources with a semicolon.

9. To save any additional options that you configured, click **Save**.

What to do next

After you define a backup job, complete the following action:

| Action | How to |
|--|--|
| Create a Hyper-V restore job definition. | See “Restoring Hyper-V data” on page 112 . |

Related concepts

[“Configuring scripts for backup and restore operations” on page 208](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and Batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Starting jobs” on page 206](#)

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 109](#).
- 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, [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 245 and [“Managing user accounts”](#) on page 253.
- Windows file indexing and file restore on volumes residing on dynamic disks is not supported.

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, click the add to restore list icon  at the resource level. Select **Restore by site** or **Restore by cloud/repository server**. If restoring from a Site, click the **Select a site** drop-down menu to choose a site associated with the backup storage server you want to restore from. If restoring from a cloud or repository server, the restore source will be automatically selected.

To select a restore point that is not the latest, expand a resource in the Restore pane, and then click the add to restore list icon  that is associated with the restore point. Adding a combination of latest

restore points and non-latest restore points to the Restore List is not supported. Click the delete icon  to remove restore points from the **Restore List**.

Additional filtering options are available when viewing non-latest restore points. To view available restore points from sites, cloud resources, or repository servers, expand a resource in the **Restore** pane, then select the source type through the **Filters** menu. Once a source type is selected, all available restore points associated with the source type display. To view restore points on a specific resource, select it from the drop-down menu adjacent to the source type drop-down menu. For example, if the filter is set to **Sites**, click **Show restore points in all sites** to select a specific site.

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.

Alternate vSnap

When restoring from a restore point that was offloaded to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server used to complete the restore is the same vSnap server used to complete the backup and offload operations. To reduce load, an alternate vSnap server can be selected to serve as the gateway to complete the restore. To select an alternate vSnap server when restoring a specific, non-latest restore point from a cloud resource or repository server, select **Use alternate vSnap server for the restore job**, then select a server from the **Select alternate vSnap** menu.

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 **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

Advanced Options

Set the following advanced job definition options:

Make IA clone resource permanent

Enable this option to move the virtual disk to permanent storage and clean up the temporary resources. This is accomplished by kicking off a live migration of the resources in the background. The destination of the live migration operation is the VM Configuration volume. The Instant Access disk is still available for read/write operations while this operation is being performed.

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.

Append suffix to virtual machine name

Enter a suffix to add to the name of restored virtual machines.

Prepend prefix to virtual machine name

Enter a prefix to add to the name of restored virtual machines.

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 109](#)

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

[“Adding a Hyper-V server” on page 107](#)

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 23](#).
- 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 cloud resources or repository servers.
- When restoring files in a Resilient File System (ReFS) environment, restores from newer versions of Windows Server to earlier versions are not supported. For example, restoring a file from Windows Server 2016 to Windows Server 2012.
- File cataloging, backup, point-in-time restores, and other operations that invoke the Windows agent will fail if a non-default local administrator is entered as the **Guest OS Username** when defining a backup job. A non-default local administrator is any user that has been created in the guest OS and has been granted the administrator role.

This occurs if the registry key `LocalAccountTokenFilterPolicy` in `[HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]` is set to 0 or not set. If the parameter is set to 0 or not set, a local non-default administrator cannot interact with WinRM, which is the protocol IBM Spectrum Protect Plus uses to install the Windows agent for file cataloging, send commands to this agent, and get results from it.

Set the `LocalAccountTokenFilterPolicy` registry key to 1 on the Windows guest that is being backed up with Catalog File Metadata enabled. If the key does not exist, navigate to `[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]` and add a DWord Registry key named `LocalAccountTokenFilterPolicy` with a value of 1.

To help avoid issues that can result from time zone differences, use an NTP server to synchronize time zones across resources. For example, you can synchronize time zones for storage arrays, hypervisors, and application servers that are in your environment.

If the time zones are out of sync, you might experience errors during application registration, metadata cataloging, inventory, backup, or restore, or file restore jobs. For more information about identifying and resolving timer drift, see [Time in virtual machine drifts due to hardware timer drift](#)

Hyper-V considerations

Only volumes on SCSI disks are eligible for file cataloging and file restore.

Linux considerations

If data 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 **Manage Protection > File Restore**.
2. Enter a search string to search for a file by name, and then click the search icon . For more information about using the search function, see [Appendix A, "Search guidelines,"](#) on page 261.
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 91](#)

Use a backup job to back up VMware resources such as virtual machines, datastores, folders, vApps, and datacenters with snapshots.

[“Restoring VMware data” on page 99](#)

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.

Note: IBM Spectrum Protect Plus may create folders on application servers when applications are registered in SPP. Folders created by IBM Spectrum Protect Plus will need to remain for the product to function properly and should not be removed. If there is a need to remove SPP created folder, unregister the application from SPP. SPP will perform a cleanup of folders associated with the registration.

You may not assign more than one application per machine as an application server to a resource group. For example, if SQL and Exchange occupy the same machine and both are registered in SPP, only one of those can be added as an application server to a given resource group.

Db2

After you successfully add your IBM Db2 instances to IBM Spectrum Protect Plus, you can start to protect your Db2 data. Create Service Level Agreements (SLAs) to back up your data, and define retention policies to maintain the data by using Service Level Agreements.

Ensure that your Db2 environment meets the system requirements. For more information, see [“Db2 requirements”](#) on page 29.

Prerequisites for Db2

All prerequisites for the IBM Spectrum Protect Plus Db2 application server must be met before you start protecting Db2 resources with IBM Spectrum Protect Plus.

Requirements for the IBM Spectrum Protect Plus Db2 application server are available here, [Db2 requirements](#).

Space prerequisites

Ensure that you have enough space on the Db2 database management system, in the volume groups for the backup operation, and on the target volumes for copying files during the restore operation. For more information about space requirements, see [Space requirements for Db2 protection](#). When you are restoring data to an alternative location, allocate extra dedicated volumes for the copy and restore processes. The data paths for table spaces and logs on the target host are the same as the paths on the original host, but the paths are on separate volumes. This setup is needed to allow copying of data from the mounted vSnap to the target host. Ensure that dedicated local database directories are allowed for each database in your volume setup.

More configuration requirements

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

- Db2 archive logging is activated, and Db2 is in recoverable mode.
- Ensure that the effective file size **ulimit -f** for the IBM Spectrum Protect Plus agent user and the Db2 instance user, is set to **unlimited**. Alternatively, set the value to a sufficiently high value to allow copying of the largest database files in your backup and restore jobs. If you change the **ulimit** setting, restart the Db2 instance to finalize the configuration.
- If you are running IBM Spectrum Protect Plus in an AIX or Linux environment, ensure that the installed sudo version is at the recommended level. For more information, see technote [2013790](#). Then, set sudo privileges as described in [“Setting sudo privileges for Db2”](#) on page 121.

- In a Linux environment, ensure that the Linux utility package `util-linux-ng` or `util-linux` package is current.
- Unicode characters in MongoDB file path names cannot be handled by IBM Spectrum Protect Plus. All names must be in ASCII.
- The database table spaces, online logs, and the local database directory can be on one or separate dedicated logical volumes that are managed by either LVM2 or JFS2. For layout two examples, see the following pictures. In the first picture, two types of volume groups shown. In the second picture, all volumes for data and logs are on one volume group.

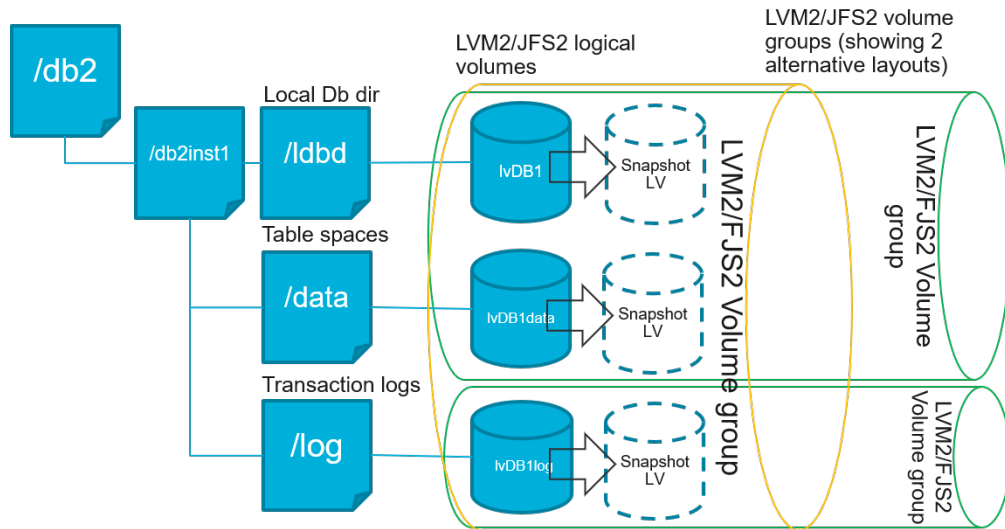


Figure 14. Logical volume layout examples

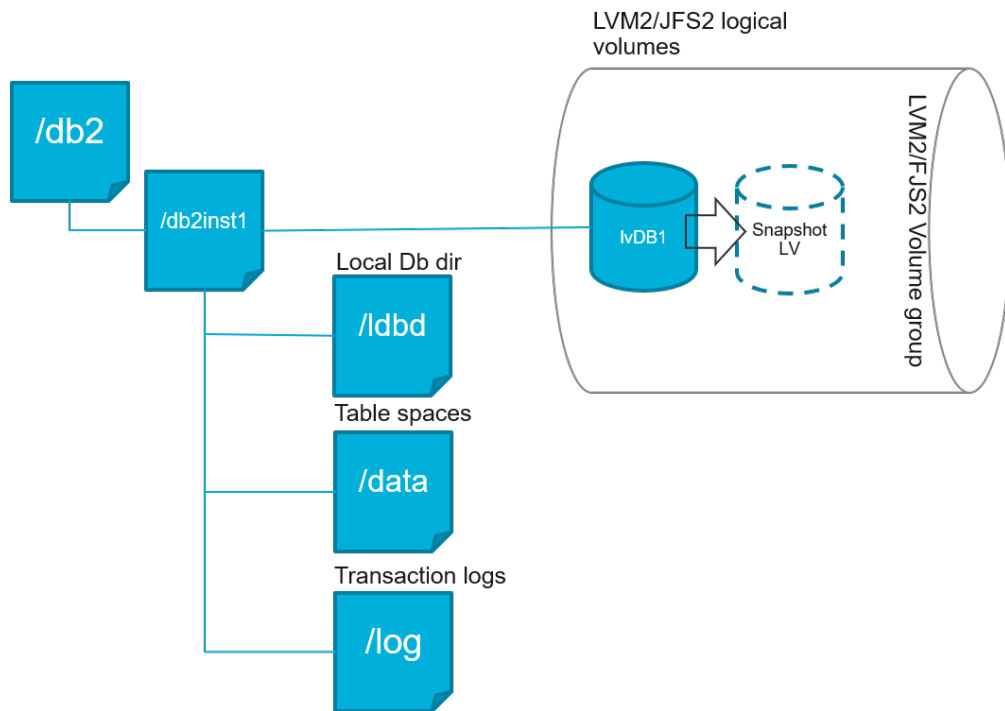


Figure 15. Single logical volume layout example

- Ensure that your Db2 logical volume setup does not include nested mount points.

Space requirements for Db2 protection

Before you start backing up Db2 databases, ensure you have enough free disk space on the target and source hosts, and in the vSnap repository. Extra free disk space is required on the volume groups for creating temporary Logical Volume Manager (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 the volume groups with Db2 data have sufficient free space.

LVM snapshots

LVM snapshots are point-in-time copies of LVM logical volumes. They are space-efficient snapshots with the changed data updates from the source logical volume. LVM snapshots are created in the same volume group as the source logical volume. The IBM Spectrum Protect Plus Db2 agent uses LVM snapshots to create a temporary, consistent point-in-time copy of the Db2 database.

The IBM Spectrum Protect Plus Db2 agent creates an LVM snapshot which is then mounted, and is copied to the vSnap repository. The duration of the file copy operation depends on the size of the Db2 database. During file copying, the Db2 application remains fully online. After the file copy operation finishes, the LVM snapshots are removed by the IBM Spectrum Protect Plus Db2 agent in a cleanup operation.

For AIX, no more than 15 snapshots can exist for each JFS2 file system. Internal and external JFS2 snapshots cannot exist at the same time for the same file system. Ensure that no internal snapshots exist on the JFS2 volumes as these snapshots can cause issues when the IBM Spectrum Protect Plus Db2 agent is creating external snapshots.

For every LVM or JFS2 snapshot logical volume containing data, allow at least 10 percent of its size as free disk space in the volume group. If the volume group has enough free disk space, the IBM Spectrum Protect Plus Db2 agent reserves up to 25 percent of the source logical volume size for the snapshot logical volume.

LVM2 and JFS2

When you run a Db2 backup operation, Db2 requests a snapshot. This snapshot is created on a Logical Volume Management (LVM) system or a Journalized File System (JFS) for each logical volume with data or logs for the selected database. In Linux systems, the logical volumes are managed by LVM2 with `lvcreate` commands. On AIX, the logical volumes are managed by JFS2 and created with the `JFS2 snapshot` command as external snapshots.

A software-based LVM2 or JFS2 snapshot is taken as a new logical volume on the same volume group. The snapshot volumes are temporarily mounted on the same machine that runs the Db2 instance so that they can be transferred to the vSnap repository.

On 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 version of the `sudo` program. For the Db2 application server, you must set up `sudo` in a specific way that might be different from other application servers.

Before you begin

To determine the correct version of `sudo` to be installed, see technote [2013790](#).

About this task

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

Procedure

1. Create an application server user by issuing the following command:

```
useradd -m <agent>
```

where `agent` specifies the name of the IBM Spectrum Protect Plus agent user.

2. Set a password for the new user by issuing the following command:

```
passwd <agent>
```

3. To enable superuser privileges for the agent user, set the `!requiretty` setting. At the end of the `sudo` configuration file, add the following lines:

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

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

Adding a Db2 application server

To start protecting your Db2 data, you must add the 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.



Figure 16. Adding a Db2 agent.

3. In the **Application Properties** section, enter the host address.
4. Choose to specify a user or use an SSH key.
5. For a user, either select an existing user, or enter a user ID and password. If you are using an SSH key, choose the key from the menu.

The user must have `sudo` privileges set up.

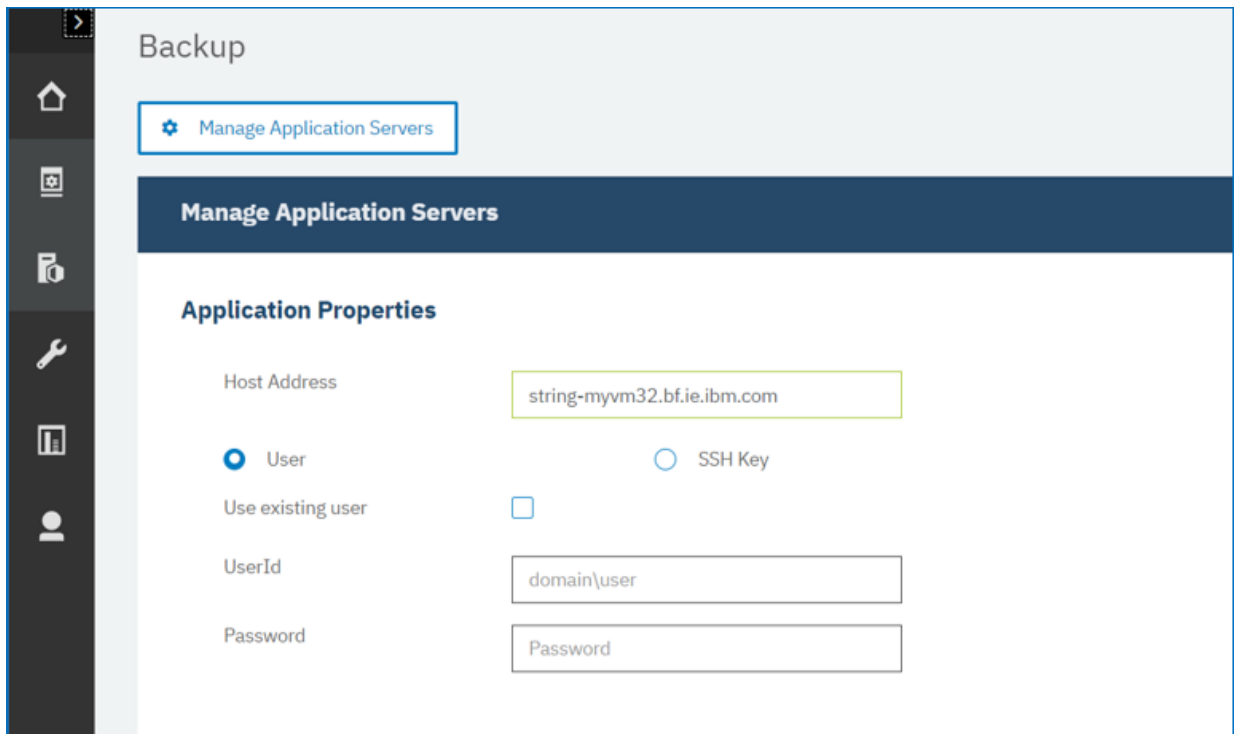


Figure 17. Managing agent users.

6. 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, review the job log. Go to **Jobs and Operations**, click the **Policy and Job List** tab, and look for the latest Application Server Inventory log entry. Databases must be detected to ensure that they can be protected. For instructions about running an inventory, see [Detecting Db2 resources](#).

Detecting Db2 resources

After you add 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 the databases available for protection with IBM Spectrum Protect Plus.

Before you begin

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

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**.

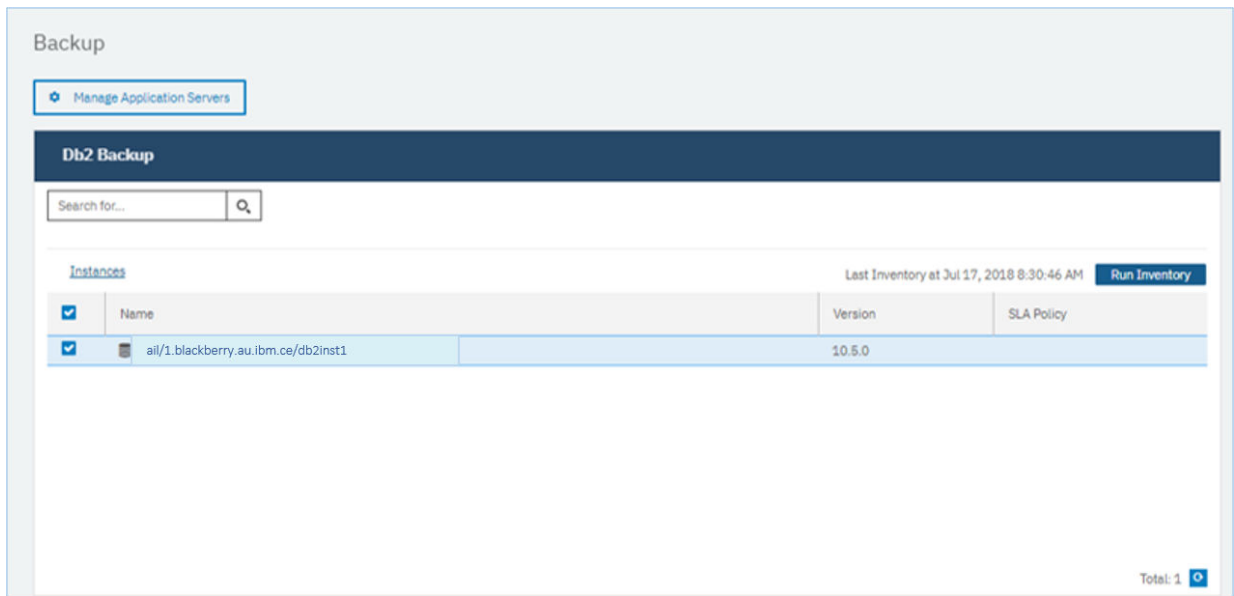


Figure 18. Detecting Db2 resources.

When the inventory is running, the button changes to show **Inventory In Progress**. You can run an inventory on any available application servers, but you can run only one inventory process at a time. To monitor the inventory job, navigate to **Jobs and Operations**, click the **Policy and Job List** tab, and look for the latest Default Application Server Inventory log entry.

3. Click on an instance to open a view that shows the databases that are detected for that instance. If any databases are missing from the **Instances** list, check your Db2 application server and rerun the inventory. In some cases, certain databases are marked as ineligible for backup; hover over the database to reveal the reason why.

Tip: To return to the list of instances, click the **Instances** hypertext in the **Backup Db2** pane.

What to do next

To start protecting Db2 databases that are cataloged in the selected instance, apply a service level agreement (SLA) policy to the instance. For instructions about setting an SLA policy, see [Defining an SLA policy](#).

Testing the Db2 connection

After you add a Db2 application server, you can test the connection. The test verifies communication with the server and the DNS settings between IBM Spectrum Protect Plus and the Db2 server. It also checks for the correct sudo permissions for the user.

Procedure

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

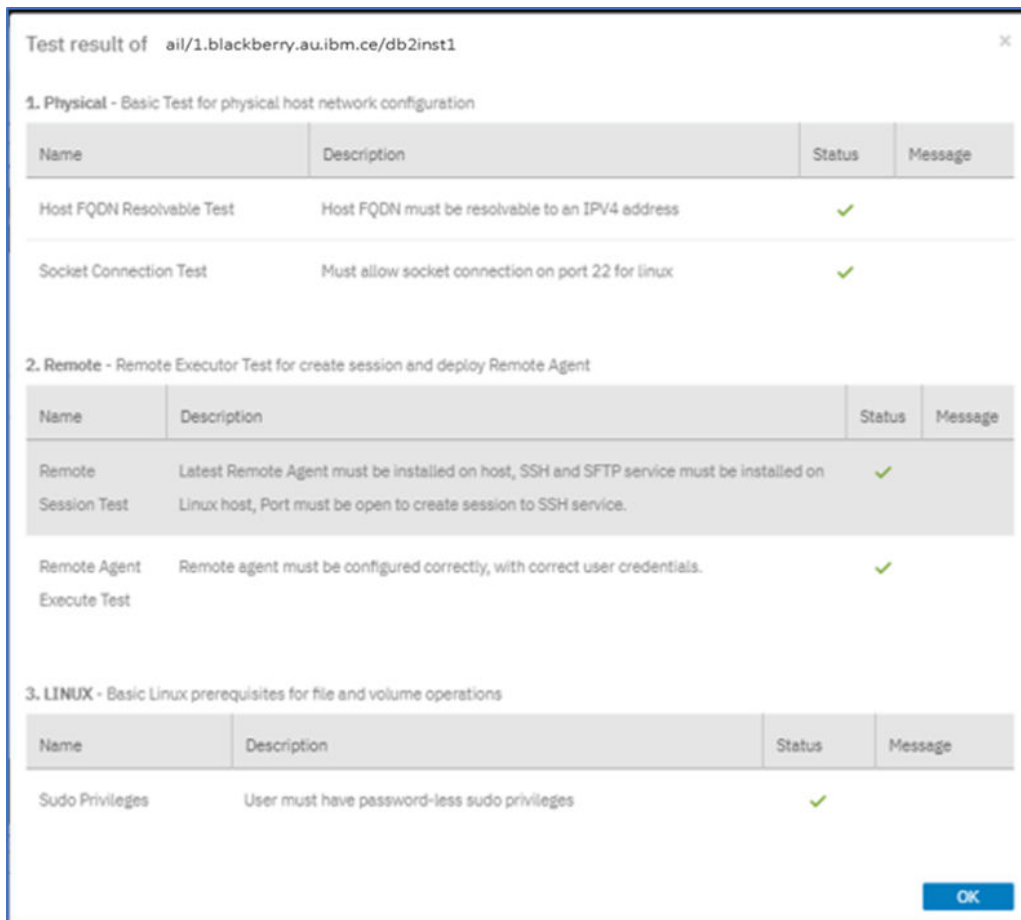


Figure 19. Testing the connection.

The test report shows a list of the tests. It consists of a test for the physical host network configuration, and tests for the remote server installation on the host, which checks SSH and SFTP on the host. The third test checks for operating system prerequisites and correct sudo privileges.

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.

Procedure

1. From the navigation menu, expand **Manage Protection > Applications > Db2 > Backup**.
2. Select an instance or database for backing up by choosing one of the following actions:
 - Select an entire instance in the **Instances** pane by clicking the check-box beside the instance name. Any databases added to this instance are automatically assigned to the SLA policy that you choose.
 - Select a specific database in an instance by clicking the instance name, and choosing a database from the list of databases in that instance.

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](#).

The screenshot shows a configuration pane titled "Options". At the top left is the title "Options" in blue. Below it is a checkbox labeled "Enable Log Backup", which is currently unchecked. Underneath the checkbox is a text input field labeled "Maximum Parallel Streams per Database" containing the number "1". At the bottom left of the pane is a blue button labeled "Save".

Figure 20. Backup pane with the Enable Log Backup option

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

After your Db2 databases are listed for each of your Db2 instances, select and apply a service level agreement (SLA) policy to start protecting your data.

Procedure

1. From the navigation menu, expand **Manage Protection > Applications > Db2 > Backup**.
2. Select a 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.

You can backup an entire instance with all the data associated, or you can choose to backup just one database.

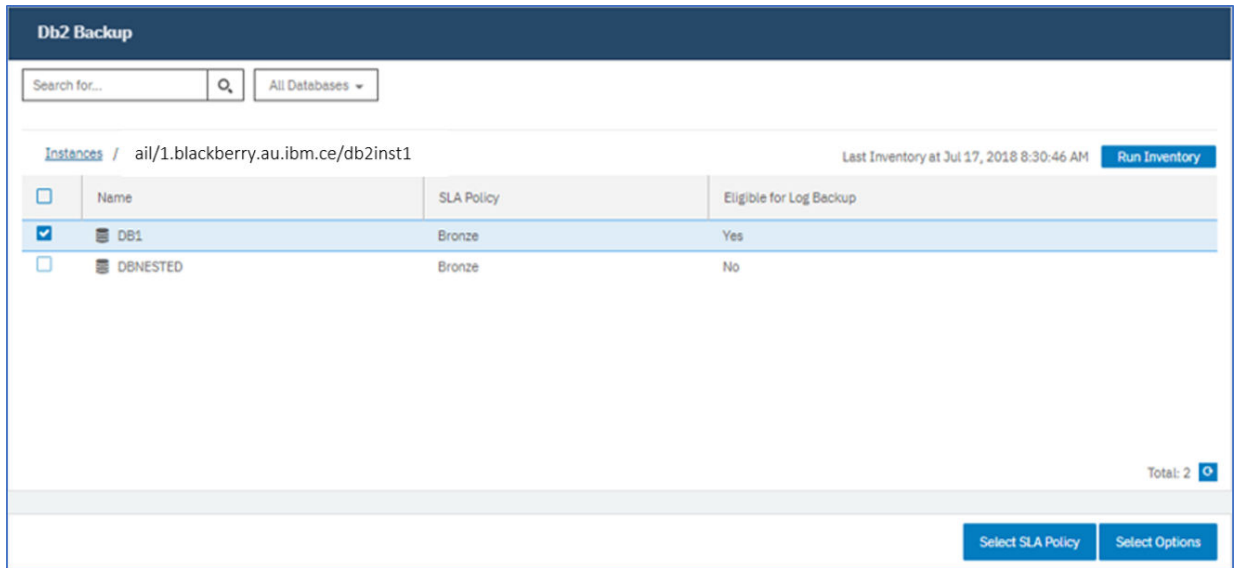


Figure 21. Db2 Backup pane showing databases in an instance

3. Click **Select SLA Policy** and choose an SLA policy, **Gold**, **Silver**, or **Bronze**. Save your choice.

Predefined choices of Gold, Silver, or Bronze, each have different frequencies and retention rates. You can create a custom SLA policy or edit an existing policy, by navigating to **Policy Overview > SLA Policies**.

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.

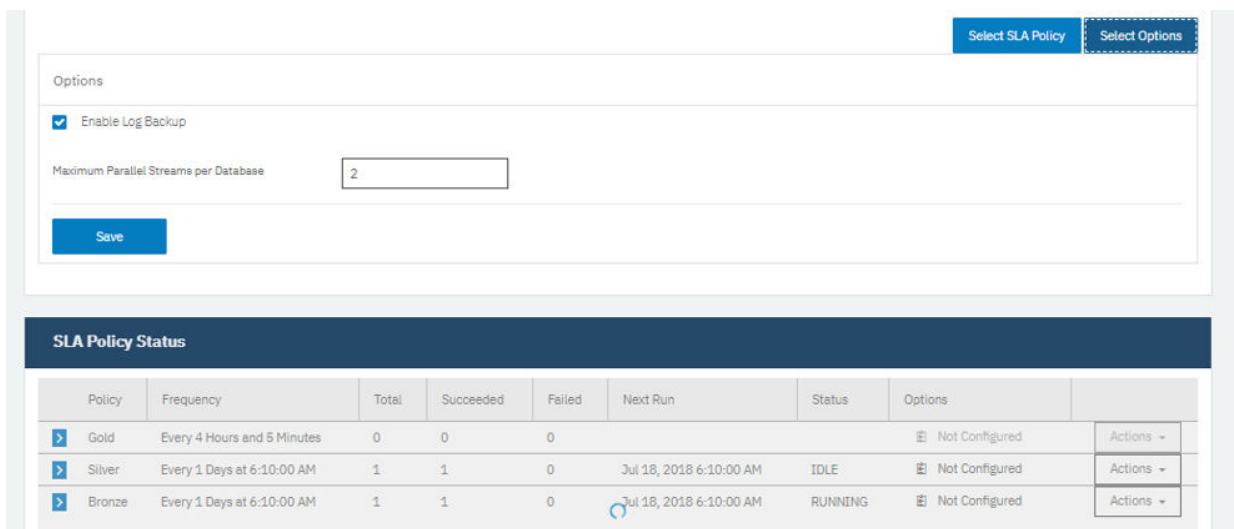


Figure 22. Backup options and SLA policies

- Configure the SLA policy by clicking the icon in the **Options** column of the **SLA Policy Status** table.
To read about more SLA configuration options, see [“Setting SLA configuration options for a backup job”](#) on page 128.
- 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.

| Policy | Frequency | Total | Succeeded | Failed | Next Run | Status | Options | Actions |
|--------|----------------------------|-------|-----------|--------|-------------------------|---------|----------------|-----------|
| Gold | Every 4 Hours | 0 | 0 | 0 | | | Not Configured | Actions - |
| Silver | Every 1 Days at 6:10:00 AM | 0 | 0 | 0 | | | Not Configured | Actions - |
| Bronze | Every 1 Days at 6:10:00 AM | 1 | 1 | 0 | Aug 31, 2018 6:10:00 AM | RUNNING | Not Configured | Actions - |

Auto Refresh Total: 3

Figure 23. SLA policies

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

To cancel a job after it has started, click **Actions** > **Cancel**.

Setting SLA configuration options for a backup job

After you set up a service level agreement (SLA) for your backup job, you can choose to configure more options for that job. You can run scripts, exclude resources from the backup operation, and force a full base backup copy of a database if required.

Procedure


- In the **Policy Options** column of the **SLA Policy Status** table for the job you are configuring, click the clipboard icon  to specify extra configuration options.
If the job is already configured, click on the icon to edit the configuration.

Figure 24. Specifying SLA configuration options

2. Click **Pre-Script** and define your pre-script configuration by choosing one of the following options:
 - Click **Use Script Server** and select an uploaded script from the menu.
 - Do not click **Use Script Server**. Select an application server from the list to run the script at that location.
3. Click **Post-Script** and define your post-script configuration by choosing one of the following options:
 - Click **Use Script Server** and select an uploaded script from the menu.
 - Do not click **Use Script Server**. Select an application server from the list to run the script at that location.

Scripts and script servers are configured on the **System Configuration > Script** page. For more information about working with scripts, see [Configuring scripts](#).

4. To continue running the job when the script that is associated with the job fails, select **Continue job/task on script error**.

If this option is selected, the backup or restore operation is reattempted and the script task status is reported as COMPLETED when the script completes processing with a nonzero return code. If this option is not selected, the backup or restore is not reattempted and the script task status is reported as FAILED.
5. To exclude resources from a backup job, specify the resources to exclude from the job. Enter an exact resource name in the **Exclude Resources** field. If you are unsure of a name, use wildcard asterisks that are specified before the pattern (**text*) or after the pattern (*text**). Multiple wildcards can be entered with standard alphanumeric characters and the following special characters: - _ and *. Separate entries with a semicolon.
6. To create a full new backup of a resource, enter the name of that resource in the **Force full backup of resources** field. Separate multiple resources with a semicolon.

The full backup creates a full new backup of that resource and replaces the existing backup of that resource for one occurrence only. After the full backup completes, the resource is backed up incrementally as before.

Log backups

Archived logs for databases contain committed transaction data. This transaction data can be used to run a rollforward data recovery when you are running a restore operation. Using archive log backups enhances the recovery point objective for your data.

Ensure that you select the **Enable Log Backups** option to allow rollforward recovery when you set up a backup job or service level agreement (SLA) policy. When selected for the first time, you must run a backup job for the SLA policy to activate log archiving to IBM Spectrum Protect Plus on the database. This backup creates a separate volume on the vSnap repository, which is mounted persistently on the Db2 application server, 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 127](#).

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 from the vSnap repository, define a job that restores data from either the latest backup or 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 125.
- 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 245.

Before you start a restore operation to an alternative instance, ensure that the file system structure on the source machine is matched on the target machine. This file system structure includes table spaces, online logs, and the local database directory. Ensure that dedicated volumes with sufficient space are allocated to the file system structure. Db2 must be at the same version level on the source and target hosts for all restore operations, and an instance of the same name must exist on each host. For more information about space requirements, see [Space requirements for Db2 protection](#). For more information about prerequisites and setup, see [Prerequisites for Db2](#).

Procedure

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

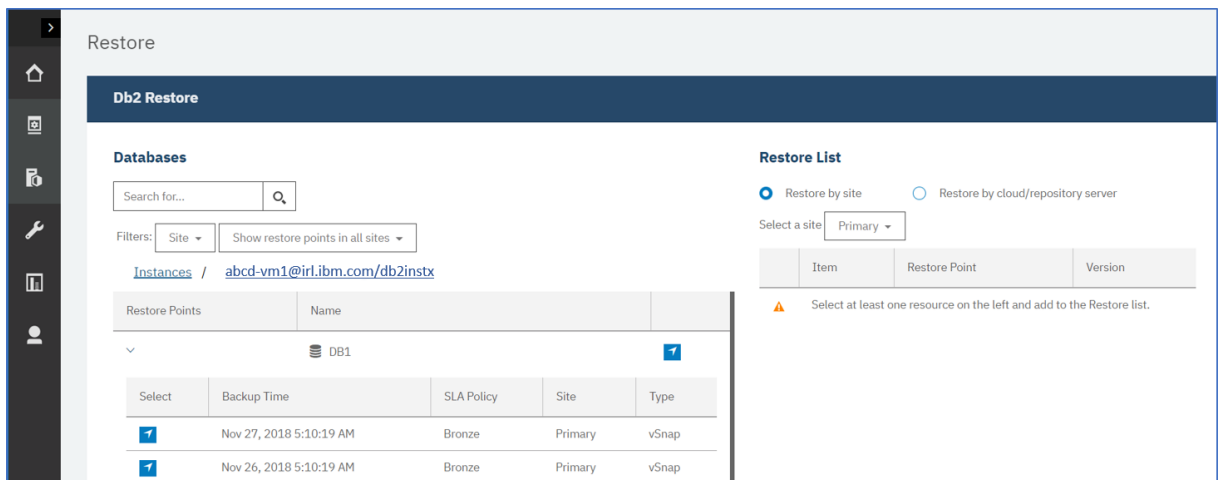



Figure 25. Restore panel for Db2

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

4. Choose the latest backup or an earlier backup from the **Restore Points** list, and select **Restore by site** or **Restore by cloud/repository server** :
 - To restore the latest backup, click the add icon  next to the database name on the right of the **Restore Points** table.

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





Figure 26. Selecting a database for point-in-time restore

- 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 .

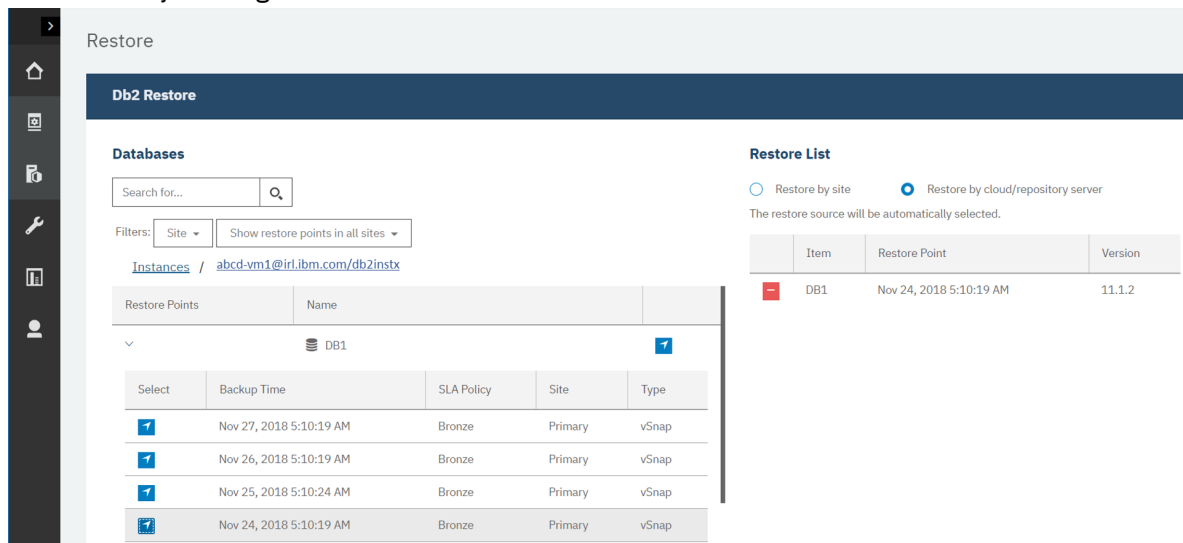



Figure 27. Adding a specific backup the restore list

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

For information about using cloud repositories for backup and restore jobs, see [“Offload to secondary backup storage”](#) on page 6, and [“IBM Spectrum Protect Plus on IBM Cloud”](#) on page 9.



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. To define options for the restore job, click **Options**.

- 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 alternative 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 data to an alternative 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, Jan 1, 2019 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.

Application options are not available for instant access restore jobs.

- **Overwrite existing databases.** Choose this option to replace existing databases that have the same names during the restore recovery process. If this option is not selected, the restore job fails when databases with the same name are found during the restore operation. If you select this option, ensure that the Db2 log directory and the Db2 mirror log directory have no data.



Attention: Ensure that no other databases share the 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.
- **Mount Point Prefix.** For instant access restore operations, specify the prefix for the path where the mount point is to be directed.

9. Specify **Scripts Settings** by choosing one or more of the following actions:

- 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 **System Configuration > Script** page to configure scripts and script servers.
- Select **Post-script** to select an uploaded script and an application or script server where the post-script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Navigate to the **System Configuration > Script** page to configure scripts and script servers.
- Select **Continue job/task on script error** to continue running the job when the script associated with the job fails. When this option is enabled and the pre-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 with a FAILED status.

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 28. Restore job definition in Manage Restore Jobs

To cancel the job, navigate to **System Configuration** and then click the **Policy and Job List** tab. Find the restore job that you want to cancel. Click **Actions**, and select **Cancel**.

When the test restore operation is listed in the **Active Resources** pane, select **Actions > Cancel** to cancel that process. If the status is not updated, click **Refresh** to update the list.

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 clicking 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

You can restore a database backup to its original instance on the original host. You can restore to the latest backup or an earlier Db2 database backup version. When you restore a database to its original instance, you cannot rename it. This restore option runs a full production restoration of data, and existing data is overwritten at the target site if the **Overwrite existing databases** option is selected.

Before you begin

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 125.
- 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 245.

Before you start a restore operation to an alternative instance, ensure that the file system structure on the source machine is matched on the target machine. This file system structure includes table spaces, online logs, and the local database directory. Ensure that dedicated volumes with sufficient space are allocated to the file system structure. Db2 must be at the same version level on the source and target hosts for all restore operations, and an instance of the same name must exist on each host. For more information about space requirements, see [Space requirements for Db2 protection](#). For more information about prerequisites and setup, see [Prerequisites for Db2](#).

Procedure

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

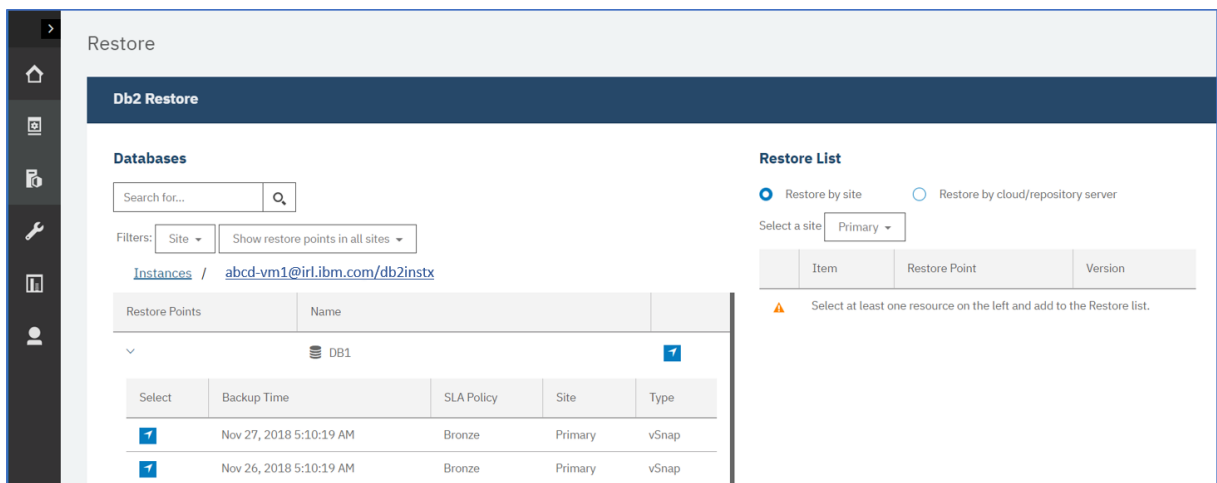



Figure 29. Restore panel for Db2

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

4. Choose the latest backup or an earlier backup from the **Restore Points** list, and select **Restore by site** or **Restore by cloud/repository server** :
 - To restore the latest backup, click the add icon  next to the database name on the right of the **Restore Points** table.
- To restore data with recovery options to a specific point-in-time, you must select the overall database.

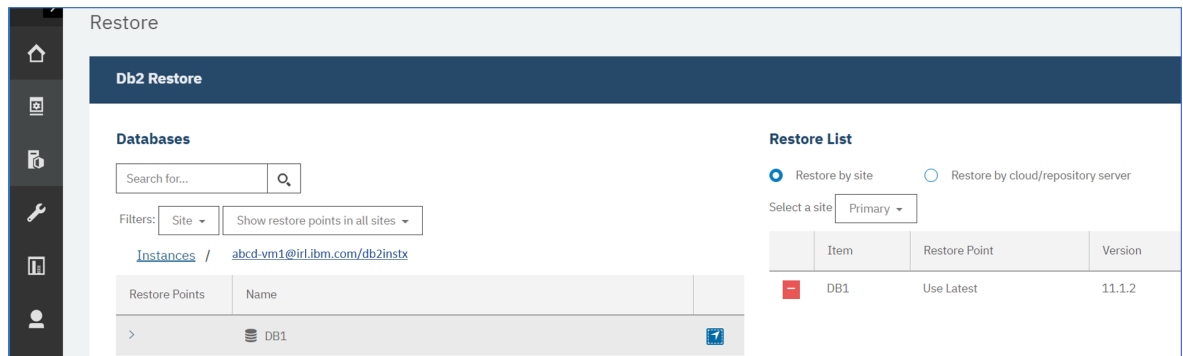



Figure 30. Selecting a database for point-in-time restore

- 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 .

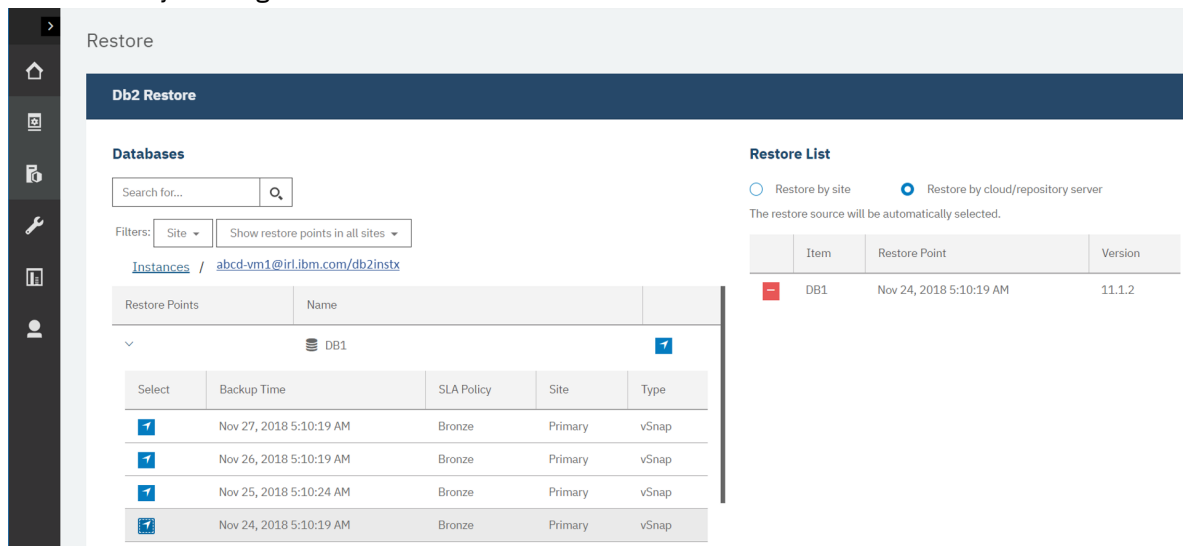



Figure 31. Adding a specific backup the restore list

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

For information about using cloud repositories for backup and restore jobs, see [“Offload to secondary backup storage”](#) on page 6, and [“IBM Spectrum Protect Plus on IBM Cloud”](#) on page 9.



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

- To define options for the restore job, click **Options**.

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

- Click **Restore to original instance** to restore data to the original server for the production restore operation.

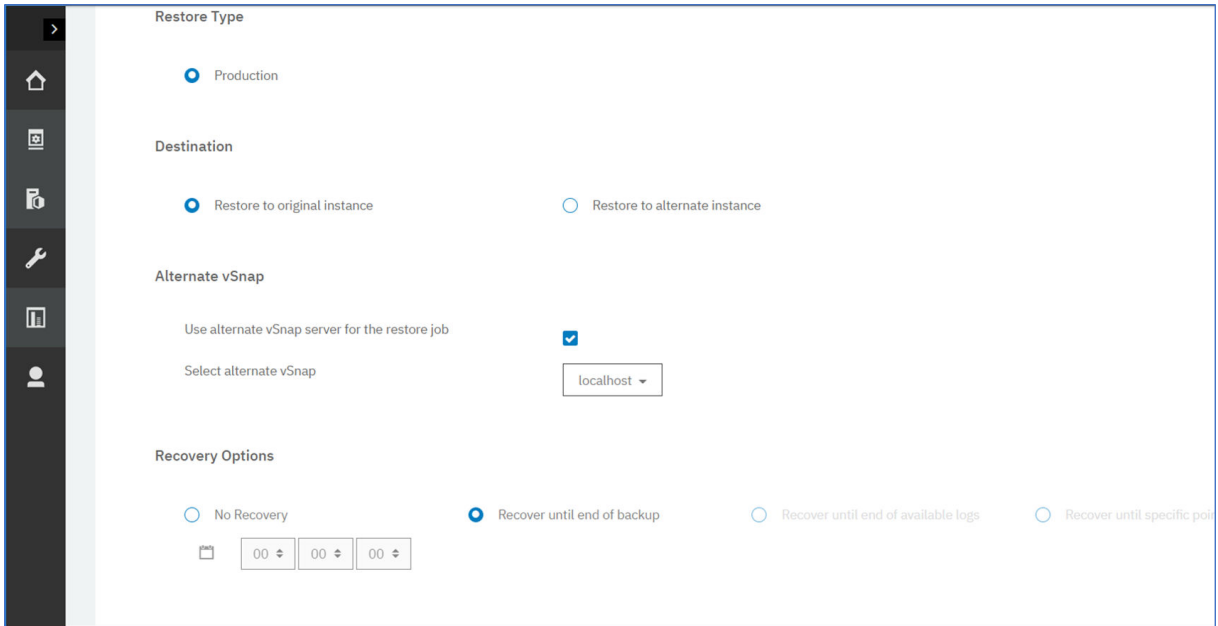


Figure 32. Restore options

- When you are restoring data from the cloud or a repository server, choose an alternative vSnap for the restore operation by selecting the **Use alternate vSnap server for the restore job**, and selecting a vSnap from the menu.

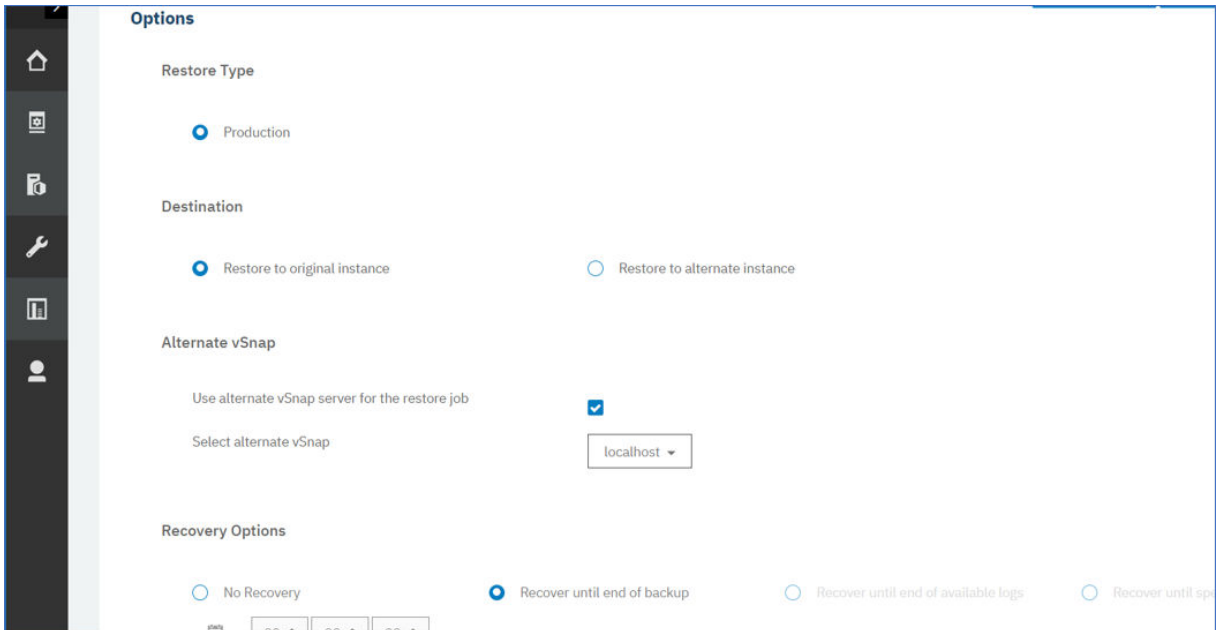


Figure 33. Restoring from the cloud or from a different vSnap.

- 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, Jan 1, 2019 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. Edit **Application Options**.

- **Overwrite existing databases.** Select this option to overwrite 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.

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.
- **Mount Point Prefix.** For instant access restore operations, specify the prefix for the path where the mount point is to be directed.

11. Specify **Scripts Settings** by choosing one or more of the following actions:

- 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 **System Configuration > Script** page to configure scripts and script servers.
- Select **Post-script** to select an uploaded script and an application or script server where the post-script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Navigate to the **System Configuration > Script** page to configure scripts and script servers.
- Select **Continue job/task on script error** to continue running the job when the script associated with the job fails. When this option is enabled and the pre-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 with a FAILED status.

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.

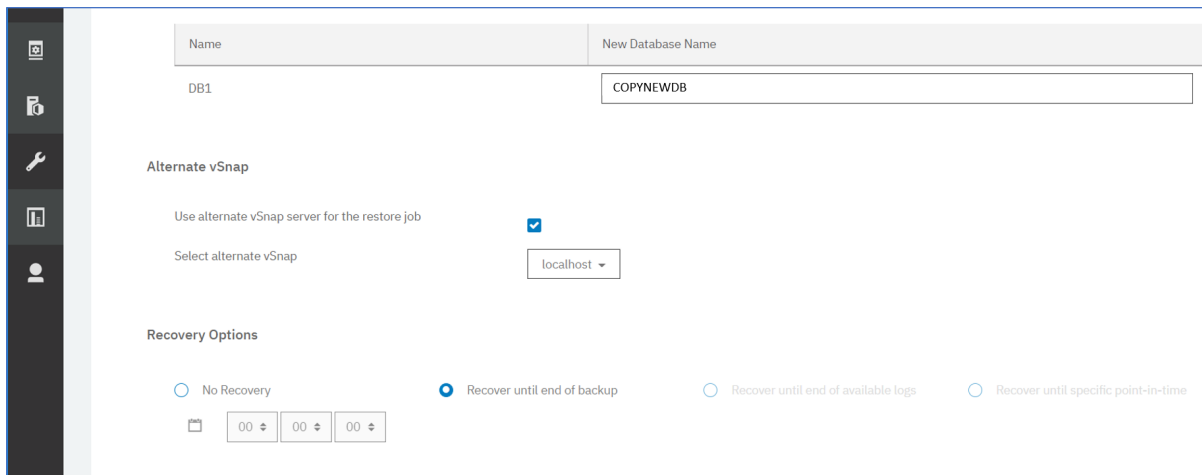


Figure 34. Restore job definition in Manage Restore Jobs

To cancel the job, navigate to **System Configuration** and then click the **Policy and Job List** tab. Find the restore job that you want to cancel. Click **Actions**, and select **Cancel**.

Restoring Db2 databases to an alternative instance

You can restore a Db2 database to a Db2 instance on an alternative host. You can also choose to restore a database to an instance with a different name and rename the database. This process creates an exact copy of the database on a different host in a different instance. If you are restoring a resource to an alternative location, you can restore the same resource multiple times without specifying different target hosts.

Before you begin

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 125.
- 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 245.

Before you start a restore operation to an alternative instance, ensure that the file system structure on the source machine is matched on the target machine. This file system structure includes table spaces, online logs, and the local database directory. Ensure that dedicated volumes with sufficient space are allocated to the file system structure. Db2 must be at the same version level on the source and target hosts for all restore operations, and an instance of the same name must exist on each host. For more information about space requirements, see [Space requirements for Db2 protection](#). For more information about prerequisites and setup, see [Prerequisites for Db2](#).

Restriction: If data exists on the local database directory to which you are restoring the database backup to, and the **Overwrite existing databases** option is not selected, the restore operation fails. No other data can share the local database directory where the backup will be restored. When the **Overwrite existing databases** option is selected, any existing data is removed and the local database directory on the alternate host.

About this task

Ensure that the disk paths for the redirected restore operation include the instance name and the database name. The information is needed for all types of paths: database paths, container paths, storage paths, and log and mirror log paths.

Procedure

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

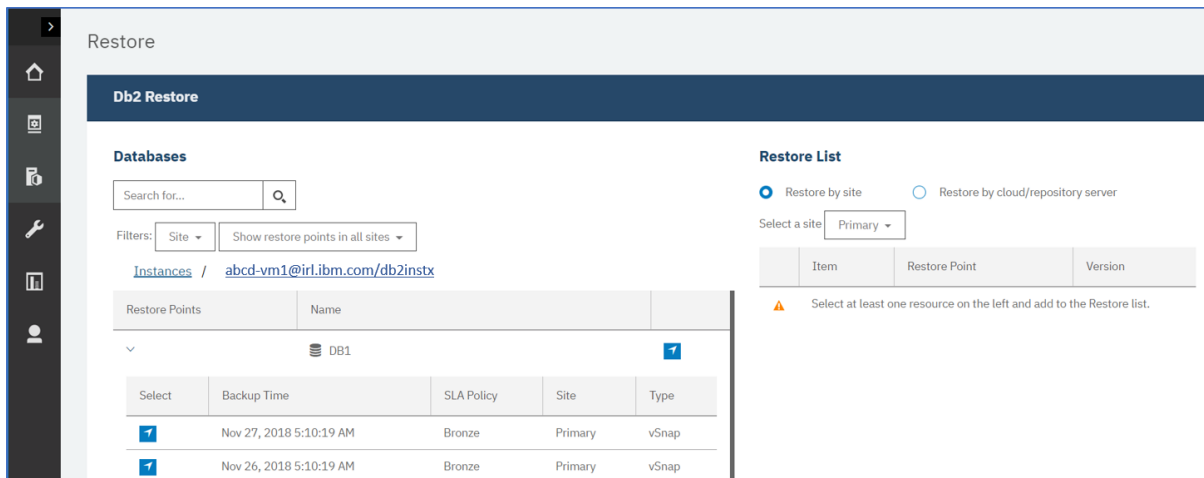


Figure 35. Restore panel for Db2

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

4. Choose the latest backup or an earlier backup from the **Restore Points** list, and select **Restore by site** or **Restore by cloud/repository server** :

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

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

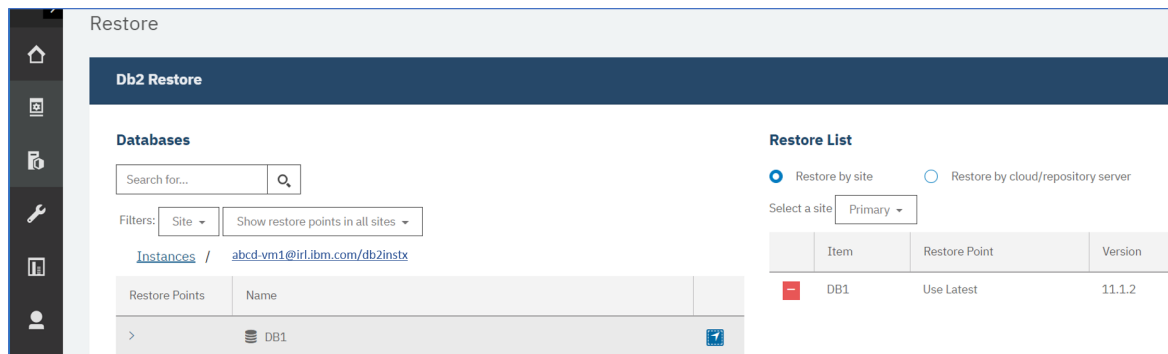



Figure 36. Selecting a database for point-in-time restore

- 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 37. Adding a specific backup to the restore list

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

For information about using cloud repositories for backup and restore jobs, see [“Offload to secondary backup storage”](#) on page 6, and [“IBM Spectrum Protect Plus on IBM Cloud”](#) on page 9.



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. Select 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 alternative 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 data to an alternative instance.

6. Click **Options** and choose **Destination > Restore to alternate instance**.

7. 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.

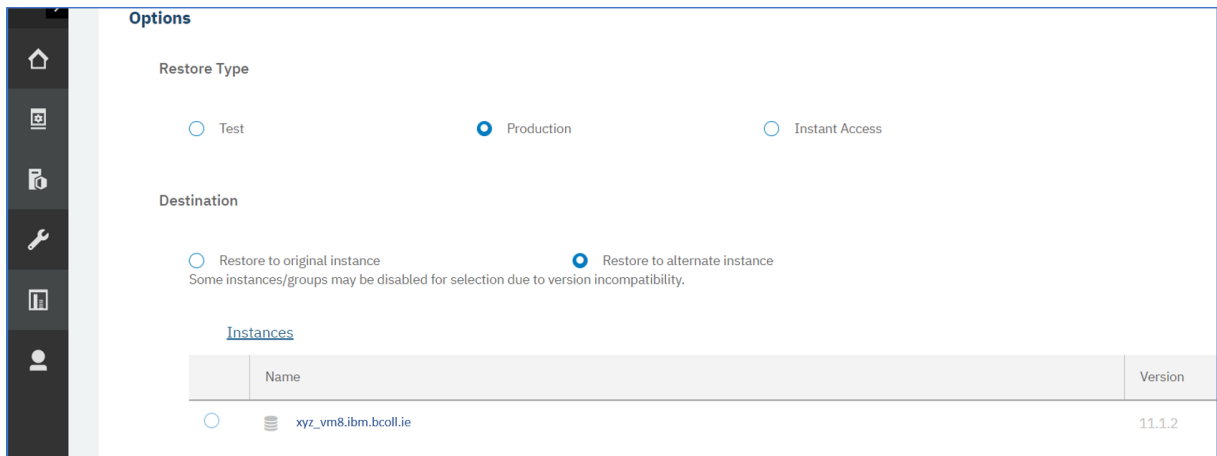


Figure 38. Restoring data to an alternative instance

8. When you are restoring data from the cloud or a repository server, choose an alternative vSnap for the restore operation by selecting the **Use alternate vSnap server for the restore job**, and selecting a vSnap from the menu.

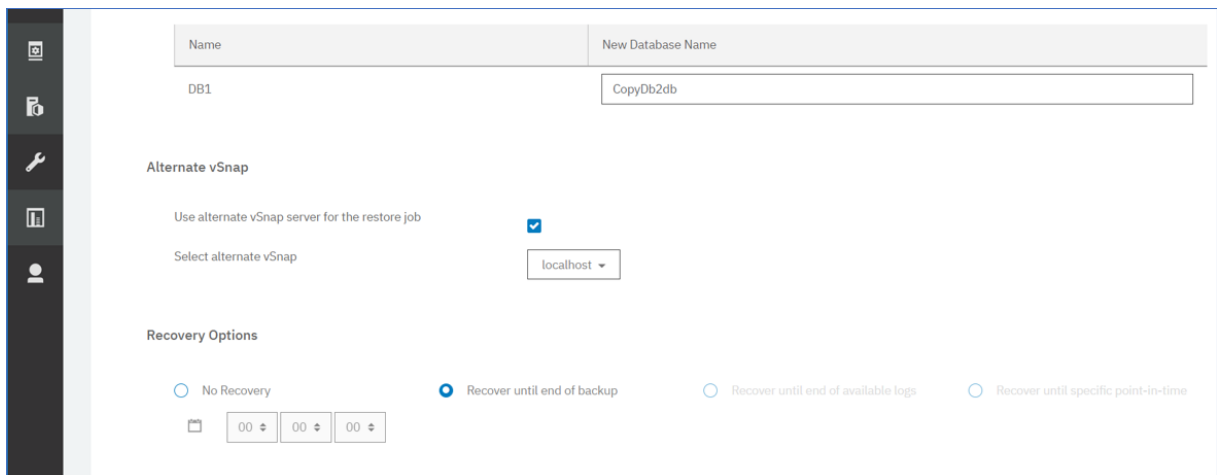


Figure 39. Restoring from an alternate vSnap

9. Optional: Enter a new database name in the **New Database Name** field.
10. 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, Jan 1, 2019 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.

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

Application options are not available for instant access restore jobs.

- **Overwrite existing databases.** Choose this option to replace existing databases that have the same names during the restore recovery process. If this option is not selected, the restore job fails when databases with the same name are found during the restore operation. If you select this option, ensure that the Db2 log directory and the Db2 mirror log directory have no data.



Attention: Ensure that no other databases share the 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.

12. 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.
- **Mount Point Prefix.** For instant access restore operations, specify the prefix for the path where the mount point is to be directed.

13. Specify **Scripts Settings** by choosing one or more of the following actions:

- 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 **System Configuration > Script** page to configure scripts and script servers.
- Select **Post-script** to select an uploaded script and an application or script server where the post-script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Navigate to the **System Configuration > Script** page to configure scripts and script servers.
- Select **Continue job/task on script error** to continue running the job when the script associated with the job fails. When this option is enabled and the pre-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 with a FAILED status.

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

15. 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.

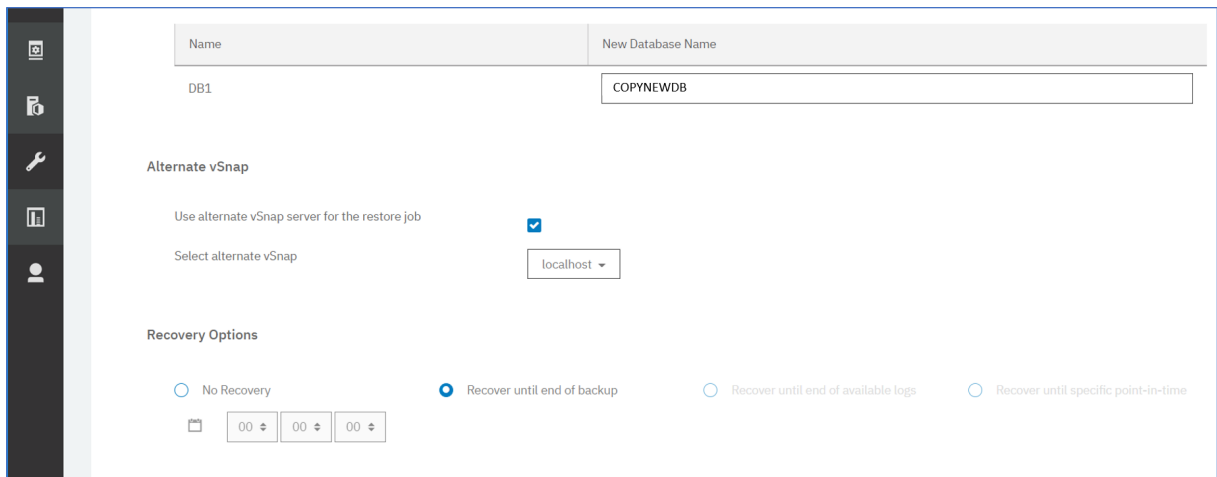


Figure 40. Restore job definition in Manage Restore Jobs

To cancel the job, navigate to **System Configuration** and then click the **Policy and Job List** tab. Find the restore job that you want to cancel. Click **Actions**, and select **Cancel**.

When the test restore operation is listed in the **Active Resources** pane, select **Actions > Cancel** to cancel that process. If the status is not updated, click **Refresh** to update the list.

Microsoft Exchange Server

After you successfully register a Microsoft Exchange Server, you can start to protect Microsoft Exchange data with IBM Spectrum Protect Plus. Define a service level agreement (SLA) policy to create backup jobs with specific schedules, retention policies, and scripts.

Prerequisites for Microsoft Exchange Server

Ensure that all prerequisites for your Microsoft Exchange application are met before you start protecting Microsoft Exchange databases with IBM Spectrum Protect Plus.

For more information, see [“Microsoft Exchange Server requirements” on page 27](#).

Virtualization support

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

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

Privileges

To help ensure that a Microsoft Exchange agent can work in your IBM Spectrum Protect Plus environment, you must set up appropriate privileges.

Role-based access control

For IBM Spectrum Protect Plus security, users who are logged on to the Exchange Server must have role-based access control (RBAC) permissions to access mailboxes and to complete mailbox restore tasks.

If your user name is authorized by the security policy in your organization, you can add user names in the Exchange Organization Management role group or subgroups. A user whose name is in the Exchange Organization Management role group or subgroups can complete mailbox restore operations. A user whose name is not in the Exchange Organization Management role group or subgroups might experience slower performance when completing restore operations.

You must define a minimum set of management roles and role scope for the Exchange user.

You must assign the following management roles to each Exchange user: Active Directory Permissions, Databases, Disaster Recovery, Mailbox Import Export, View-Only Configuration, and View-Only Recipients.

To restore an Exchange public folder mailbox, the Exchange user must also have the Public Folders management role. To restore mail to a Unicode PST file, the Exchange user must have the Mailbox Import Export management role.

To assign management roles to a user, use an Exchange Powershell cmdlet as shown in the following example:

```
New-RoleGroup -Name "My Admins" -Roles "Active Directory Permissions",  
"Databases", "Disaster Recovery", "Mailbox Import Export", "Public Folders",  
"View-Only Configuration", "View-Only Recipients" -Members operator1
```

The preceding example creates a group, My Admins, with minimum roles to run the IBM Spectrum Protect Plus Exchange agent, and assigns user operator1 to this group. The operator1 user can run the IBM Spectrum Protect Plus Exchange agent but with limited Exchange privileges; for example, the user cannot create or remove a user mailbox.

Management role scope:

Ensure that the following Exchange objects are in the management role scope for the Exchange user:

- The Exchange Server that contains the required data
- The recovery database that IBM Spectrum Protect Plus creates
- The database that contains the active mailbox
- The database that contains the active mailbox of the user who completes the restore operation

Verify that the Exchange user name is a member of a local Administrator group and has an active Exchange mailbox in the domain. By default, Windows adds the Exchange Organization Administrators group to other security groups, including the local Administrators group. For Exchange users who are not members of the Exchange Organization Management group, you must manually add the user account to the local Administrators group by using the Local Users and Groups tool on the computer of the domain member.

On the computer of the domain member, click **Administrative tools > Computer Management > Local Users and Groups tool**. On a domain controller computer that does not have a local Administrators group or Local Users and Groups tool, manually add the user account to the Administrators group in the domain by clicking **Administrative tools > Active Directory Users and Computers tool**.

Encrypting File System

IBM Spectrum Protect Plus for Exchange requires that Encrypting File System (EFS) is enabled in the local or group domain policy, and a valid Domain Data Recovery Agent (DRA) certificate is available. If a custom group policy is defined and linked to the organizational unit, ensure that the Exchange server is part of the organizational unit.

Adding a Microsoft Exchange application server

When you register Microsoft Exchange Server, an inventory of Exchange databases is added to IBM Spectrum Protect Plus. When the inventory is available, you can start to backup and restore your Exchange databases and run reports.

About this task

To register a Microsoft Exchange application server, you need the IP address or host name.

Procedure

To add a Microsoft Exchange application server, complete the following steps:

1. In the navigation pane, expand **Manage Protection > Applications > Exchange**, and then click **Backup**.
2. In the **Backup** pane, click **Manage Application Servers**, and then click **Add Application Server** to add the host system.
3. In the **Application Properties** form, enter the IP or host address.
4. Enter a user ID in the format of active directory domain and user account (domain\user), and the associated password. This user must have the correct Exchange roles and privileges. For more information about Exchange privileges, see [“Privileges ” on page 144](#).
5. Click **Save**, and repeat the steps to add other Microsoft Exchange instances to IBM Spectrum Protect Plus.

Important: In a database availability group (DAG) environment, register all Microsoft Exchange application servers in the DAG.

What to do next

When you add your Exchange application server to IBM Spectrum Protect Plus, an inventory is automatically run on each instance. Databases must be detected to ensure that they can be backed up, and you can run a manual inventory at any time to detect updates. For instructions about running a manual inventory, see [“Detecting Microsoft Exchange databases by running an inventory” on page 146](#). For instructions about setting up Exchange database backup jobs, see [“Defining a Service Level Agreement backup job” on page 147](#).

Detecting Microsoft Exchange databases by running an inventory

When you add your Microsoft Exchange Server instances to IBM Spectrum Protect Plus, an inventory is run automatically. However, you can run an inventory on an Exchange application server manually at any time to detect updates and list all of the Exchange databases for each instance.

Before you begin

Ensure that you added your Exchange instances to IBM Spectrum Protect Plus. For instructions about adding an Exchange instance, see [“Adding a Microsoft Exchange application server” on page 145](#).

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > Exchange**, and then click **Backup**.
2. Click **Run Inventory**.
When the inventory is running, the button label changes to **Inventory In Progress**. You can run an inventory on any available application server, but you can run only one inventory process at a time.
3. To monitor the inventory job, go to **Jobs and Operations**, click the **Policy and Job List** tab, and look for the latest Application Server Inventory log entry.
4. When the inventory job is complete, on the **Backup** pane, click an Exchange instance to open a view that shows the databases that are detected for that instance. If any databases are missing from the **Instances** list, check your Microsoft Exchange application server and rerun the inventory.

Tip: To return to the list of instances, click the **Instances** hypertext in the Exchange Backup pane.

Testing the Microsoft Exchange connection

After you register a Microsoft Exchange application server and add it to the application server list, test the connection. The test verifies communication between IBM Spectrum Protect Plus and the host application server.

Procedure

1. In the navigation pane, expand **Manage Protection**, and then click **Applications**.
2. Next, click **Exchange**, and then click **Backup**.

3. On the **Backup** screen, click **Manage Application Servers**.

The Microsoft Exchange application servers that are available are shown.

4. Click **Actions** for the Microsoft Exchange application server that you want to test, and then click **Test**.

The test report shows you a list of the tests that ran and their status. Each test procedure includes a test of the physical host network configuration, a remote session test, and a test of Windows prerequisites such as user administrator privileges.

5. Click **OK** to close the test. Run the test again after you fix any issues.

Defining a Service Level Agreement backup job

When your Microsoft Exchange databases are listed for each of your Exchange instances, select and apply a service level agreement (SLA) policy to start protecting your data.

About this task

IBM Spectrum Protect Plus supports single or multiple Microsoft Exchange databases per Exchange backup job. Multiple database backup jobs run sequentially.

Procedure

1. In the navigation pane, click **Manage Protection**, and then click **Applications**.
2. Expand **Exchange**, and then click **Backup**.
3. Select an Exchange instance to back up all the data in that instance, or click an instance name, and then select individual databases that you want to back up.
4. Click **Select SLA Policy** and choose an SLA Policy.
Predefined choices are Gold, Silver, and Bronze, each with different frequencies and retention rates. Gold is the most frequent with the shortest retention rate. You can also create a custom SLA policy or edit an existing policy. For more information see [“Creating an SLA policy” on page 79](#).
5. 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.
6. Configure the SLA policy by clicking the icon in the **Policy Options** column of the **SLA Policy Status** table.
For more information about SLA configuration options, see [“Setting SLA configuration options for a backup job” on page 147](#).
7. To run the policy outside of the scheduled job, select the instance or database and then click **Actions > Start**.
The status changes to **Running** for your chosen SLA. To pause the schedule, click **Actions > Pause Schedule**, and to cancel a job after it has started, click **Actions > Cancel**.

Setting SLA configuration options for a backup job

After you set up a service level agreement (SLA) for your backup job, you can choose to configure more options for that job. Extra SLA options include running scripts, excluding resources from the backup operation, and forcing a full base backup copy if required.

Procedure

1. In the **Policy Options** column of the **SLA Policy Status** table for the job that you are configuring, click the clipboard icon to specify additional configuration options.
2. To define a pre-script configuration, select **Pre-Script** and take one of the following actions:
 - To use a script server, select **Use Script Server** and choose an uploaded script from the **Script** or **Script Server** list.
 - To run a script on an application server, clear the **Use Script Server** check box, and choose an application server from the **Application Server** list.

3. To define a post-script configuration, select **Post-Script** and take one of the following actions:

- To use a script server, select **Use Script Server** and choose an uploaded script from the **Script** or **Script Server** list.
- To run a script on an application server, clear the **Use Script Server** check box, and choose an application server from the **Application Server** list.

Scripts and script servers are configured on the **System Configuration > Script** page. For more information about working with scripts, see [Configuring scripts](#).

4. Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails.

If this option is selected, the backup or restore operation is attempted and the script task status is reported as COMPLETED when the script completes processing with a nonzero return code. If this option is not selected, the backup or restore is not attempted and the script task status is reported as FAILED.

5. Specify resources to exclude them from the backup job. Enter an exact resource name in the **Exclude Resources** field. If you are unsure of a name, use wildcard asterisks that are specified before the pattern (**text*) or after the pattern (*text**). Multiple wildcards can be entered with standard alphanumeric characters and the following special characters: - _ and *. Separate entries with a semicolon.

6. If you want to create a full backup of a particular resource, enter the name of that resource in the **Force full backup of resources** field. Separate multiple resources with a semicolon.

A full backup replaces the existing backup of that resource for one occurrence only. After that, the resource is backed up incrementally as before.

7. Click **Save**.

Backing up Microsoft Exchange database logs

You can back up the database transaction logs for Microsoft Exchange databases. Exchange log backups are scheduled by using Windows Task Scheduler. When log backups are available, you can run a rollforward data recovery during a restore operation to ensure that the data is recovered to the latest possible point in time.

About this task

When log backups are enabled, a Task Scheduler task is created on the Exchange server. The task runs a backup operation of your Exchange log files according to the SLA policy.

Procedure

1. In the navigation pane, expand **Manage Protection**, and then click **Applications**.
2. Next, click **Exchange**, and then click **Backup**.
3. Click the Microsoft Exchange instance that you want to protect, and then select the databases whose logs you want to back up.
Tip: The **Eligible for Log Backup** column shows the databases for which you can run log backups. If a database is registered as not eligible for log backup, a hover help explanation is provided.
4. Click **Select Options** and then select **Enable Log Backup**.
5. Enter the frequency of the log backups in days, hours, or minutes.
6. Choose the start date and select the time for the log backups to begin, and then click **Save**.

Results

The database transaction logs are backed up to the vSnap server according to the selected frequency.

Restriction: The database logs are backed up on the preferred node only. Only one Microsoft Exchange instance at a time can write log backups to the vSnap server.

Any log backup issues that occur are displayed in the Alert notifications in IBM Spectrum Protect Plus.

Backing up Exchange databases in a Database Availability Group

You can back up the mailbox databases in a Microsoft Exchange Database Availability Group (DAG) and specify whether to use the active copy or a passive copy of the database for the backup. The Exchange servers in a DAG environment synchronize the data between active and passive copies for high availability.

About this task

By using the information from an inventory job, IBM Spectrum Protect Plus provides a DAG view that displays all of the databases in an Exchange DAG environment. Each database has an active copy on one server in the DAG, and one or more passive copies on the other servers. By default, scheduled backups are taken from the server that the database is active on, but you can select a different server to back up a passive copy of the database.

Procedure

1. In the navigation pane, click **Manage Protection > Applications > Exchange > Backup**.
2. In the **Backup** pane, click the **View** menu and select **Database Availability Group**.
3. Click the Microsoft Exchange DAG that you want to view, and then select the databases to back up.
4. Click **Select Options**. In the **Backup preferred node** list, select the instance to run the backups on.
With the **Backup preferred node** option, you can select a passive copy of the database for the backup.
5. Click **Select SLA Policy** and then select an SLA policy from the list.
6. To create the job definition by using default options, click **Save**.
The DAG databases are scheduled for backup jobs in accordance with the selected SLA policies and the preferred node choices.
7. To run the selected policy outside of the schedule, in the **SLA Policy Status** pane, click **Actions > Start**.

Incremental forever backup strategy

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

The incremental forever backup solution provides the following advantages:

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

The IBM Spectrum Protect Plus incremental forever process includes the following steps:

1. The first backup job creates a VSS snapshot of the Exchange application. As a result, the database files are in an application consistent state. The complete database files are copied to the vSnap location.
2. All subsequent backups create a VSS snapshot of the Exchange application. The database files are in an application consistent state. However, only the change blocks of the database files are copied to the vSnap location.
3. The backups are reconstructed at each point in time that a backup is performed, making it possible to recover the database from any single backup point.

Restoring Microsoft Exchange databases

If data in a Microsoft Exchange database is lost or corrupted, you can restore the data from a backup copy. You can 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.

Before you begin

Ensure that the following requirements are met:

- At least one Microsoft Exchange backup job is defined and ran successfully. For instructions about defining a backup job, see [“Defining a Service Level Agreement backup job”](#) on page 147.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is defining the restore job. For more information about assigning roles, see [Chapter 13, “Managing user access,”](#) on page 245.

Important: For granular restore operations, you must log on to the Exchange application server and use the Microsoft Management Console (MMC) GUI to complete mailbox batch restore and mailbox restore browser tasks.

Procedure

To restore data in a Microsoft Exchange database, take one of the following actions:

- Restore a database to its original instance.
- Restore a database to an alternative instance.
- Restore mailbox data by using the granular restore function.
- Restore a database in a database availability group (DAG).

Restoring a Microsoft Exchange database to the original location

Restore a Microsoft Exchange database to its original location and original instance. Choose between restoring the latest backup or an earlier Exchange database backup version.

Before you begin

Ensure that the following requirements are met:

- At least one Microsoft Exchange backup job is defined and ran successfully.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is defining the restore job. For more information about assigning roles, see [Chapter 13, “Managing user access,”](#) on page 245.

About this task

When you restore a database to its original location, you cannot rename it. This restore option runs a full production restore operation, and existing data is overwritten at the target site.

Procedure

1. In the navigation pane, click **Manage Protection**, and then click **Applications**.
2. Expand **Exchange**, and then click **Restore**.
3. In the **Restore** window, click an Exchange instance from the list to expand that instance and show the databases.
4. Click the database that you want to restore to show the available restore points for that database, and select the backup version to restore.
Restore points are listed with the backup time stamp, SLA policy type, and site information.
5. Click **Options** to review the restore operation options. In the **Destination** pane, select **Restore to original instance** for the production restore operation.

The **Restore Type** is already selected as **Production**; there are no other choices for restoring databases to the original location.

6. Expand **Recovery Options** to select further options.

The default value is **Recovery**.

- **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 manually.
- **Recovery until end of backup**. This option restores the database and applies logs to recover it up to the time that the backup operation was completed.
- **Recover until end of available logs**. This option restores the database and applies all available logs, including logs that are newer than the backup that exists on the application server, to recover the database up to the latest possible time.
- **Recover until specific point-in-time**. This option restores the database and applies logs from the log backup volume to recover the database up to an intermediate, user-specified point in time. This includes all the backup data up to a specific point in time, for example, 5 September 2018 12:15:00 AM. This option is available only if you selected **Enable Log Backup** in the backup job.

7. You can specify further restore options in the **Advanced Options** pane.

- **Run cleanup immediately on job failure**. This option is selected by default to automatically clean up allocated resources as part of a restore operation if the recovery fails.
- **Continue with restores of other selected databases even if one fails**. If one database in the instance fails to be restored successfully, the process continues for all other databases that are being restored.
- **Mount Point Prefix**. When specifying an instant access restore job, type in the prefix for the path where the mount point is to be directed.

8. To save your restore choices, click **Save**.

9. To restore data immediately, click **Restore**.

10. To specify a schedule for the restore operation, click **Manage Jobs**.

Restoring a Microsoft Exchange database to an alternative instance

You can select a Microsoft Exchange database backup and restore it to a local instance on an alternative host. You can also choose to restore the backup to an instance with a different name, or you can rename the database.

Before you begin

Ensure that the following requirements are met:

- Enough disk space and allocated dedicated volumes are available for the copying of files.
- The file system structure on the source server is the same as the file system structure on the target server. This file system structure includes table spaces, online logs, and the local database directory.

Procedure

1. In the navigation pane, click **Manage Protection**, and then click **Applications**.

2. Expand **Exchange**, and then click **Restore**.

3. In the **Restore** window, click an Exchange instance from the list to expand that instance and show the databases.

4. Click the database that you want to restore to show the available restore points for that database, and select the backup version to restore.

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

5. Click **Options** to review the restore operation options.

6. Choose a restore type:

- **Test**. Choose this option to restore the data from the vSnap repository directly. This restore type might be used for testing purposes.

- **Production.** Choose this option to restore the full database with a full-copy data restore operation. This restore operation is for permanent use of the restored database.
 - **Instant Access.** Choose this option to create a copy of the database that you can access to view data instantly. In instant access mode, no further action is taken after IBM Spectrum Protect Plus mounts the share. Use the data for custom recovery of data from the files in the vSnap volume.
7. In the Destination pane, choose **Restore to alternate instance**, and select the target instance that you want to restore the database to.
 8. Optional: In the **New Database Name** field, enter a new database name.
 9. Expand the database name to see the path information. In the **Destination Path** field, add the location of the Exchange database file including the .edb name, and the logs location.

For example, for a database that is named Database_A.edb, enter C:\ExchangeDatabase\Database_A\Database_A.edb, and for the location of the logs (**Source Path E01**), enter D:\ExchangeDatabase\Log\Database_A\
 10. Choose the recovery type for the restore in the **Recovery Options** pane:
 - **No Recovery.** This option skips any rollforward recovery after the restore operation. The database remains in a Rollforward pending state until you decide whether you want to run the rollforward recovery manually.
 - **Recover until end of backup.** This choice recovers only the data that is contained in the backup that is selected.
 - **Recover until end of available logs.** This option restores the database and applies all available logs (including logs newer than the backup that might exist on the application server) to recover the database up to the latest possible time.
 - **Recover until specific point-in-time.** This option restores the database and applies logs from the log backup volume to recover the database up to an intermediate, user-specified point in time. This includes all the backup data up to a specific point in time, for example, 8 June 2018 12:18:00 AM. This option is available only if you selected **Enable Log Backup** in the backup job.
 11. Review the job definition and click **Save**.
 12. To restore the database immediately, click **Restore**.
 13. To specify a schedule for the restore operation, click **Manage Jobs**.

Restoring individual mailbox items by using a granular restore operation

You can restore Microsoft Exchange individual mailbox items by using a granular restore operation and the IBM Spectrum Protect Plus Microsoft Management Console (MMC) GUI.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations. If RBAC permissions were not assigned, you might encounter configuration errors in the IBM Spectrum Protect Plus MMC GUI for each missing role.

Tip:

If you encounter role-based configuration errors in the IBM Spectrum Protect Plus MMC GUI, you can set the required permissions manually to resolve the errors (see “Privileges ” on page 144), or you can run the IBM Spectrum Protect Plus configuration wizard to automatically configure permissions (see step “13” on page 153).

About this task

To start a granular restore operation, complete preparatory steps in the IBM Spectrum Protect Plus GUI, and then log in to the Exchange application server. Then, use the IBM Spectrum Protect Plus MMC GUI to restore user mailbox data from the recovery database that is created by the granular restore operation. A granular restore operation can be used to complete the following tasks:

- You can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode .pst file.

- You can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
- You can restore an archive mailbox or a part of the mailbox, for example, a specific folder.
- You can restore archive mailbox messages to a mailbox that is on the Exchange Server, to an archive mailbox, or to an Exchange Server .pst file.

Procedure

1. In the navigation pane of the IBM Spectrum Protect Plus GUI, click **Manage Protection**, and then click **Applications**.
2. Expand **Exchange**, and then click **Restore**.
3. In the **Restore** window, click an Exchange instance from the list to expand that instance and show the databases.
4. Choose the database that contains the mailbox that you want to restore and expand the list of restore points for that database.
5. Select the backup version that you want to restore from the list of restore points.
6. Click **Options** to review the restore operation options.
7. In the **Restore Type** section, choose **Granular Restore**.
In the **Destination** section, the recovery database name is displayed. The name consists of the existing database name with the suffix `_RDB`.
8. Click **Save**.
9. To start the granular restore operation, click **Restore**.
The recovery database is displayed in the **Active Resources** section.
Tip: Click the **i** icon to display an information message that describes the next steps for completing the granular restore task.
10. Connect to the Exchange application server instance by using Remote Desktop Connection (RDC) or Virtual Network Computing (VNC) if connecting remotely, or by logging on to the Exchange Server machine locally.
The granular restore operation automatically installs and starts the IBM Spectrum Protect Plus MMC GUI on the application server. If the MMC GUI fails to start, start it manually by using the path that is provided in the **Active Resources** information message.
11. In the IBM Spectrum Protect Plus MMC GUI, click the **Protect and Recover Data** node, and select **Exchange Server**.
12. On the **Recover** tab for the Exchange Server instance, click **View > Mailbox Restore Browser** to view the mailbox from the recovery database.
13. Optional: Run the IBM Spectrum Protect Plus configuration wizard:
 - a) In the navigation pane, click **Dashboard > Manage > Configuration > Wizards > IBM Spectrum Protect Plus Configuration**.
 - b) In the **Actions** pane, click **Start**.
The configuration wizard runs the requirements check.
 - c) When the requirements checks have run, click the **Warnings** link next to **User Roles Check**.
 - d) On the message dialog box, to add any missing roles, click **Yes**.
 - e) On the configuration wizard, click **Next**, and then click **Finish**.
14. In the **Mailbox Restore Browser > Source** tree, click the mailbox that contains the items you want to restore, which enables you to browse the individual folders and messages.
Choose from the following actions to select the folder or message to restore.

| <i>Table 18. Previewing and filtering mailbox items</i> | |
|---|--|
| Task | Action |
| Preview mailbox items | a. Select a mailbox item, such as Inbox , to display its contents in the preview pane. b. Click an individual item in the preview pane, such as an email message, to view the message text and details. c. If an item contains an attachment, click the attachment icon to preview its contents. |
| Filter mailbox items | Use the filter options to narrow the list of folders and messages to restore: <ol style="list-style-type: none"> Click Show Filter Options and Add Row. Click the down arrow in the Column Name field and select an item to filter. You can filter by folder name, subject text, and other options. Restriction: You can filter public mailbox folders only by the Folder Name column. When you select All Content, the mailbox items are filtered by attachment name, sender, subject, and message body. In the Operator field, select an operator: Contains. In the Value field, specify a filter value. To specify additional filtering criteria, click Add Row. Click Apply Filter to filter the messages and folders. |

15. When you have selected the mailbox item to restore, in the **Actions** pane, click the restore task that you want to run. Choose from the following options:

- **Restore Folder to Original Mailbox**
- **Restore Messages to Original Mailbox**
- **Save Mail Message Content**

Tip: If you click **Save Mail Message Content**, a Windows Save File window is displayed. Specify the location and message name and click **Save**.

When you choose the restore option, the **Restore Progress** window opens and shows the progress of the restore operation, and the mailbox item is restored.

16. To restore a mailbox item to another mailbox or .pst file, complete the following steps.

Note: You can also restore a complete mailbox to another mailbox or .pst file.

Choose from the actions in the following table:

| Table 19. Restoring a mailbox item to another mailbox or .pst file | |
|--|---|
| Task | Action |
| Restore a mailbox item (or a mailbox) to a different mailbox | <p>a. On the Actions pane, click Open Exchange Mailbox.</p> <p>b. Enter the alias of the mailbox to identify it as the restore destination.</p> <p>c. Drag the source mailbox item (or mailbox) to the destination mailbox on the results pane.</p> <p>Restriction: You cannot drag mail items or subfolders in the Recoverable Items folder to a destination mailbox.</p> |
| Restore a mailbox item (or mailbox) to an Outlook personal folders (.pst) file | <p>a. On the Actions pane, click Open non-Unicode PST File.</p> <p>b. When the Open File window opens, select an existing .pst file or create a .pst file.</p> <p>c. Drag the source mailbox item (or mailbox) to the destination .pst file on the results pane.</p> <p>Restriction: You can use the Mailbox Restore Browser view only with non-Unicode .pst files.</p> |
| Restore a Public Folder | <p>Select this action to restore a public folder to an existing online public folder mailbox.</p> <p>You can filter the mailbox and restore a specific public folder to an existing online public folder. In the Folder to be restored field, enter the name of the public folder that you want to restore.</p> <ul style="list-style-type: none"> • To restore a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name</i>. • To restore all subfolders in a parent folder, use <i>parent_folder_name/*</i>. • If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\). <p>You can also restore all or part of a public folder to a different public folder mailbox than the original mailbox. In the Target public folder mailbox field, specify the destination public folder mailbox that you want to restore to.</p> |

- In the **Actions** pane, click **Close Exchange Mailbox** or **Close PST File** to close the destination mailbox or .pst file.
- When the restore operation for the individual items is finished, return to IBM Spectrum Protect Plus, and on the **Active Resources** pane, click **Actions** > **Cleanup Granular Restore** to end the granular restore process.

Restoring mailboxes by using a granular restore operation

You can restore Microsoft Exchange mailboxes by using a granular restore operation and the IBM Spectrum Protect Plus Microsoft Management Console (MMC) GUI.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations. If RBAC permissions were not assigned, you might encounter configuration errors in the IBM Spectrum Protect Plus MMC GUI for each missing role.

Tip:

If you encounter role-based configuration errors in the IBM Spectrum Protect Plus MMC GUI, you can set the required permissions manually to resolve the errors (see [“Privileges ” on page 144](#)), or you can run the IBM Spectrum Protect Plus configuration wizard to automatically configure permissions (see step [“13” on page 157](#)).

About this task

To start a granular restore operation, complete preparatory steps in the IBM Spectrum Protect Plus GUI, and then log in to the Exchange application server. Then use the IBM Spectrum Protect Plus MMC GUI to restore user mailbox data from the recovery database that is created by the granular restore operation. A granular restore operation can be used to complete the following tasks:

- You can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode .pst file.
- You can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
- You can restore an archive mailbox or a part of the mailbox, for example, a specific folder.
- You can restore archive mailbox messages to a mailbox that is on the Exchange Server, to an archive mailbox, or to an Exchange Server .pst file.

Procedure

1. In the navigation pane of the IBM Spectrum Protect Plus GUI, click **Manage Protection**, and then click **Applications**.
2. Expand **Exchange**, and then click **Restore**.
3. In the **Restore** window, click an Exchange instance from the list to expand that instance and show the databases.
4. Choose the database that contains the mailbox that you want to restore and expand the list of restore points for that database.
5. Select the backup version that you want to restore from the list of restore points.
6. Click **Options** to review the restore operation options.
7. In the **Restore Type** section, choose **Granular Restore**.

In the **Destination** section, the recovery database name is displayed. The name consists of the existing database name with the suffix `_RDB`.

8. Click **Save**.

9. To start the granular restore operation, click **Restore**.

The recovery database is displayed in the **Active Resources** section.

Tip: Click the **i** icon to display an information message that describes the next steps for completing the granular restore task.

10. Connect to the Exchange application server instance by using Remote Desktop Connection (RDC) or Virtual Network Computing (VNC) if connecting remotely, or by logging on to the Exchange Server machine locally.

The granular restore operation automatically installs and starts the IBM Spectrum Protect Plus MMC GUI on the application server. If the MMC GUI fails to start, start it manually by using the path that is provided in the **Active Resources** information message.

11. In the IBM Spectrum Protect Plus MMC GUI, click the **Protect and Recover Data** node, and select **Exchange Server**.
12. On the **Recover** tab for the Exchange Server instance, select **View > Mailbox Restore**.
A list of user mailboxes from all databases that are included in the backup is displayed.
13. Optional: Run the IBM Spectrum Protect Plus configuration wizard:
 - a) In the navigation pane, click **Dashboard > Manage > Configuration > Wizards > IBM Spectrum Protect Plus Configuration**.
 - b) In the **Actions** pane, click **Start**.
The configuration wizard runs the requirements check.
 - c) When the requirements checks have run, click the **Warnings** link next to **User Roles Check**.
 - d) On the message dialog box, to add any missing roles, click **Yes**.
 - e) On the configuration wizard, click **Next**, and then click **Finish**.
14. Select one or more mailboxes from the recovery database to restore. Mailboxes are listed by Mailbox Name, Alias, Server, Database, and Mailbox Type.
You can restore only user mailboxes that are located in the recovery database.
Tip: Mailboxes from other databases are shown in this view for informational purposes only. If the mailbox that you want to restore is not in the recovery database, use this view to determine which Exchange database the user mailbox was assigned to. You can then run the granular restore task again for that database.
15. To complete the restore operation, in the **Actions** pane, click one of the following restore options.

| <i>Table 20. Restore options</i> | |
|---|--|
| Option | Action |
| Restore Mail to Original Location | Restore mail items to their location at the time of the backup operation. |
| Restore Mail to Alternate Location | Restore the mail items to a different mailbox. <ul style="list-style-type: none"> • On the Alternate Mailbox Options window, enter the Mailbox alias name. Tip: If deleted mail items or tasks are flagged in the Recoverable Items folder of a mailbox, the items are restored with the flag attribute to the Flagged Items and Tasks view in the target mailbox. |
| Restore Mail to non-Unicode PST file Restriction: <ul style="list-style-type: none"> • This option is available only for Exchange Server 2013. • Each folder can contain a maximum of 16,383 mail items. | Restore mail items to a non-Unicode personal folders (.pst) file. When you restore mail items to a .pst file with one selected mailbox, you are prompted for a file name. When you restore mail items to a .pst file with more than one selected mailbox, you are prompted for a directory location. Each mailbox is restored to a separate .pst file that reflects the name of the mailbox at the specified directory. If the .pst file exists, the file is used. Otherwise, the file is created. |

| Table 20. Restore options (continued) | |
|---|---|
| Option | Action |
| Restore Mail to Unicode PST file | <p>Restore mail items to a Unicode .pst file.</p> <p>When you restore mail items to a .pst file with one selected mailbox, you are prompted for a file name. When you restore mail items to a .pst file with more than one selected mailbox, you are prompted for a directory location.</p> <p>Tip:</p> <p>You can enter a standard path name (for example, c:\PST\mailbox.pst) or a UNC path (for example, \\server\c\$\PST\mailbox.pst). When you enter a standard path, the path is converted to a UNC path. If the UNC is a non-default UNC path, enter the UNC path directly.</p> <p>Each mailbox is restored to a separate .pst file that reflects the name of the mailbox at the specified directory. If the .pst file exists, the file is used. Otherwise, the file is created.</p> |
| Restore Public Folder Mailbox | <p>Restore a public folder mailbox to an online public folder mailbox.</p> <p>In the Folder to be restored field, enter the name of the public folder that you want to restore:</p> <ul style="list-style-type: none"> • To restore a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name.</i> • To restore all subfolders in a parent folder, use <i>parent_folder_name/*.</i> • If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\). <p>You can also restore all or part of a public folder mailbox to a different public folder mailbox than the original mailbox. In the Target public folder mailbox field, specify the destination public folder mailbox.</p> |

| Table 20. Restore options (continued) | |
|---|---|
| Option | Action |
| Restore Mail to Archive Mailbox | <p>This action applies to a primary mailbox or an archive mailbox. Select this action to restore all or part of either type of mailbox to the original archive mailbox or to an alternative archive mailbox.</p> <p>You can filter the archive mailbox and restore a specific mailbox folder. In the Folder to be restored field, enter the name of the folder in the archive mailbox that you want to restore.</p> <ul style="list-style-type: none"> To restore a subfolder in a parent folder, specify the full folder path in this format: <i>parent_folder_name/sub_folder_name.</i> To restore all subfolders in a parent folder, use <i>parent_folder_name/*.</i> If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\). <p>In the Target archive mailbox field, specify the archive mailbox destination.</p> |
| Exclude recoverable mail items while restoring the mailbox | <p>Apply this action if you are restoring an online, public folder, or archive mailbox to an original mailbox, alternative mailbox, or to a Unicode .pst file.</p> <p>Specify a value of Yes to exclude the mail items in the Recoverable Items folder in mailbox restore operations. No is the default value.</p> |

- When the mailbox restore operation is finished, return to IBM Spectrum Protect Plus, and on the **Active Resources** pane, click **Actions > Cleanup Granular Restore** to end the granular restore process.

Restoring Database Availability Group backups

With IBM Spectrum Protect Plus, you can restore an Exchange Server Database Availability Group (DAG) backup to the original instance or to an alternative instance.

About this task

In a DAG environment, you must restore a database to an active database copy. If you had selected a passive database copy as the preferred target of backup operations, IBM Spectrum Protect Plus attempts to restore the database to this passive copy by default. The restore operation fails. In this situation, you can choose to restore the database to an alternative instance, and then select the active database copy.

Procedure

- In the navigation pane, click **Manage Protection**, and then click **Applications**.
- Expand **Exchange**, and then click **Restore**.
- In the **Restore** pane, click the **View** menu and select **Database Availability Groups**.
- In the **Availability Groups** list, click an Exchange instance to see the list of restore points for that instance and select the backup versions that you want to restore.
- Click **Options** to review the restore operation options.
- In the **Restore Type** section, choose **Production**, **Test**, or **Instant Access**.



Attention: When you choose the restore type and destination, you must select an active node as the destination; otherwise, the restore operation fails.

7. In the **Destination** section, choose **Restore to original instance** or **Restore to alternate instance**. Then, specify the destination of the restore operation.
8. Choose the recovery options, and any other details for the restore operation, and click **Save**.
9. To restore the database immediately, click **Restore**.
10. To specify a schedule for the restore operation, click **Manage Jobs**.

MongoDB

After you successfully add MongoDB instances to IBM Spectrum Protect Plus, you can start to protect the data in your MongoDB databases. Create service level agreement (SLA) policies to back up and maintain MongoDB data.

Ensure that your MongoDB environment meets the system requirements. For more information, see [“MongoDB requirements” on page 31](#).

Prerequisites for MongoDB

All system requirements and prerequisites for the IBM Spectrum Protect Plus MongoDB application server must be met before you start protecting MongoDB data with IBM Spectrum Protect Plus.

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

To meet the prerequisites for MongoDB, complete the following checks and actions.

1. Ensure you have met the space prerequisites, as described in [Space requirements for MongoDB protection](#).
2. Set the file size limit for the MongoDB instance user with the command **ulimit -f** to unlimited. Alternatively, set the value to sufficiently high to allow the copying of the largest database files in your backup and restore jobs. If you change the **ulimit** setting, restart the MongoDB instance to finalize the configuration.
3. If you are running MongoDB in an AIX or Linux environment, ensure that the installed sudo version is at a supported level.

For more information about the version level, see [“MongoDB requirements” on page 31](#). For information about setting sudo privileges, see [“Setting sudo privileges” on page 162](#).
4. If your MongoDB databases are protected by authentication, you must set up role-based access control. For more information, see [“Roles for MongoDB” on page 161](#).
5. Each MongoDB instance to be protected must be registered on IBM Spectrum Protect Plus. After the instances are registered, IBM Spectrum Protect Plus runs an inventory to detect MongoDB resources. Ensure that all instances that you want to protect are detected and listed correctly.
6. Ensure that the SSH service is running on port 22 on the server, and that firewalls are configured to allow IBM Spectrum Protect Plus to connect to the server with SSH. The SFTP subsystem for SSH must be enabled.
7. Ensure that you do not configure nested mount points.

Restrictions

The following restrictions apply to the MongoDB application server:

- MongoDB sharded cluster configurations are detected when you run an inventory, but these resources are not eligible for backup or restore operations.
- Unicode characters in MongoDB file path names cannot be handled by IBM Spectrum Protect Plus. All names must be in ASCII.

Virtualization

Protect your MongoDB environment with IBM Spectrum Protect Plus when it is running on one of the following guest operating systems:

- Red Hat Enterprise Linux
- SUSE Linux Enterprise Server Kernel-based Virtual Machine (KVM)

Roles for MongoDB

You must define role-based access control (RBAC) roles for the MongoDB agent users if authentication is enabled on the MongoDB database. When the roles are set up, users can protect and monitor MongoDB resources with IBM Spectrum Protect Plus in accordance with the users' defined roles.

Role-based access control for MongoDB

For each MongoDB user, specify access roles by using a command similar to the following example:

```
use admin
db.grantRolesToUser("<username>",
[ { role: "hostManager", db: "admin" },
  { role: "clusterManager", db: "admin" } ] )
```

The following roles are available:

hostManager

This role provides access to the **fsyncLock** command. This access is required for application-consistent backups of MongoDB databases where journaling is not enabled. This role also provides access to the shutdown command, which is used during a restore operation to shut down the MongoDB server instance that the restore is directed to.

clusterMonitor

This role provides access to commands for monitoring and reading the state of the MongoDB database. The following commands are available to users with this role:

- **getCmdLineOpts**
- **serverVersion**
- **replSetGetConfig**
- **replSetGetStatus**
- **isMaster**
- **listShards**

clusterManager

This role is only required only for running test restore operations of replica sets. Users who run the **replSetReconfig** command can create the restored instance of a single node replica set. This role enables read and write access during test restore operations of replica sets. Without this access, the node in the replica set would remain in the REMOVED state without read and write access. In addition, this role provides access to commands for reading the state of the MongoDB database. The following commands are available for this role:

- **replSetReconfig**
- **getCmdLineOpts**
- **serverVersion**
- **replSetGetConfig**
- **replSetGetStatus**
- **isMaster**
- **listShards**

Space prerequisites for MongoDB protection

Before you start backing up MongoDB data, ensure that you have enough free space on the target and source hosts, and in the vSnap repository. Extra space is required to store temporary Logical Volume Manager (LVM) backups of logical volumes where the MongoDB data is located. These temporary backups, that are known as LVM snapshots, are created automatically by the MongoDB agent.

LVM snapshots

LVM snapshots are point-in-time copies of LVM logical volumes. After the file copy operation finishes, earlier LVM snapshots are removed by the IBM Spectrum Protect Plus MongoDB agent in a cleanup operation.

For each LVM snapshot logical volume, you must allocate at least 10 percent free space in the volume group. If there is enough free space in the volume group, the IBM Spectrum Protect Plus MongoDB agent reserves up to 25 percent of the source logical volume size for the snapshot logical volume.

Linux LVM2

When you run a MongoDB backup operation, MongoDB requests a snapshot. This snapshot is created on a Logical Volume Management (LVM) system for each logical volume with data or logs for the selected database. On Linux systems, logical volumes are managed by LVM2.

A software-based LVM2 snapshot is taken as a new logical volume on the same volume group. The snapshot volumes are temporarily mounted on the same machine that runs the MongoDB instance so that they can be transferred to the vSnap repository.

On Linux, the LVM2 volume manager stores the snapshot of a logical volume within the same volume group. There must be enough space available to store the logical volume. The logical volume grows in size as the data changes on the source volume for the lifetime of the snapshot.

Setting sudo privileges

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

About this task

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

Procedure

1. Create an agent user by issuing the following command:

```
useradd -m agent
```

where *agent* specifies the name of the IBM Spectrum Protect Plus agent user.

2. Set a password for the new user by issuing the following command:

```
passwd mongodb_agent
```

3. To enable superuser privileges for the agent user, set the `!requiretty` setting. At the end of the sudo configuration file, add the following lines:

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

Alternatively, if your sudoers file is configured to import configurations from another directory, for example `/etc/sudoers.d`, you can add the lines in the appropriate file in that directory.

Adding a MongoDB application server

To start protecting MongoDB resources, you must add the server that hosts your MongoDB instances, and set credentials for the instances. Repeat the procedure to add all the servers that host MongoDB resources.

About this task

To add a MongoDB application server to IBM Spectrum Protect Plus, you must have the host address of the machine.

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > MongoDB > Backup**.
2. In the **Backup** window, click **Manage Application Servers**, and click **Add Application Server** to add the host machine.



3. In the **Application Properties** form, enter the host address.
4. Choose to register the host with a user or an SSH key.
If you select **User**, you can choose to enter a new user and password, or an existing user. If you select **SSH Key**, select the SSH key from the menu.

Restriction: Any user that is specified must have sudo privileges set up.

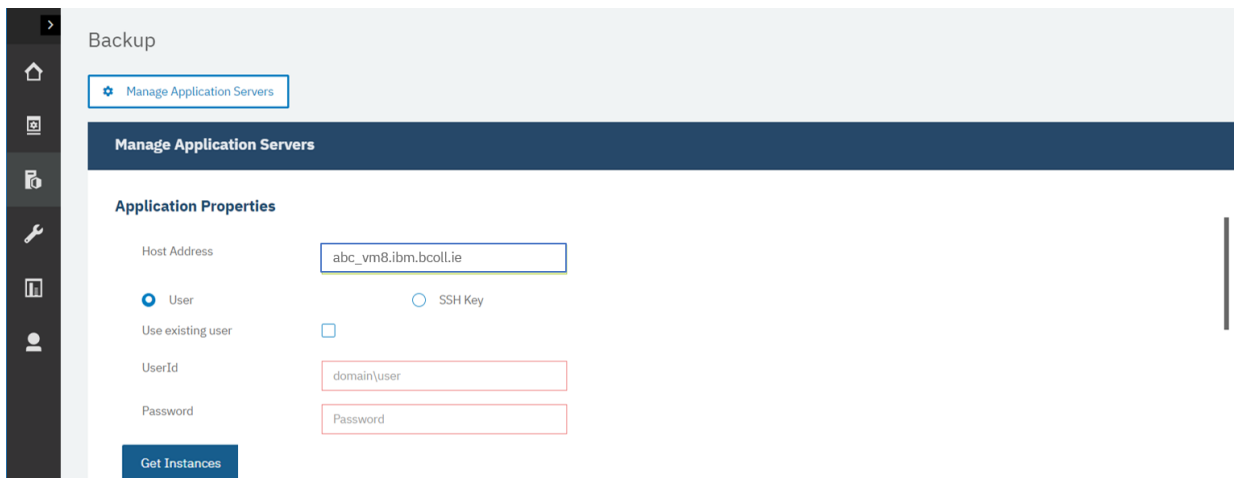


Figure 41. Adding a MongoDB agent

5. Click **Get Instances** to detect and list the MongoDB instances that are available on the host server that you are adding.

Each MongoDB instance is listed with its connection host address, status, and an indication of whether it is configured.

6. If you are using access control, configure an instance by setting credentials. Click **Set Credential**, and set the user ID, and password. Alternatively, you can select to use an existing user profile.

For more information about access control, see [Chapter 13, “Managing user access,”](#) on page 245.

When you set credentials, you assign MongoDB user roles for the backup and restore operations with access to role-protected MongoDB servers by using Salted Challenge Response Authentication Mechanism (SCRAM), or Challenge and response authentication. The MongoDB user that is assigned for the role-protected MongoDB server requires one of the following access levels to protect resources:

- **Host Manager:** manages the database as the administrator. This role is required for taking and managing snapshots.

- *Cluster Administrator*: retrieves configuration information and runs test mode restore operations of MongoDB replica sets. This role is required to reconfigure test mode restore operations of MongoDB replica sets for data queries.
 - *Cluster Monitor*: monitors the protection of MongoDB resources, and retrieves configuration information.
7. Optional: Set the option **Maximum concurrent databases** by entering a number in the field.
 8. Save the form, and repeat the steps to add other MongoDB application servers to IBM Spectrum Protect Plus.

What to do next

After you add MongoDB application servers to IBM Spectrum Protect Plus, an inventory is automatically run on each application server to detect the relevant databases in those instances. To verify that the databases are added, review the job log. Go to **Jobs and Operations**, click the **Policy and Job List** tab, and look for the newest Application Server Inventory log entry. Databases must be detected to ensure that they can be protected. For instructions about running a manual inventory, see [Detecting MongoDB resources](#).

Detecting MongoDB resources

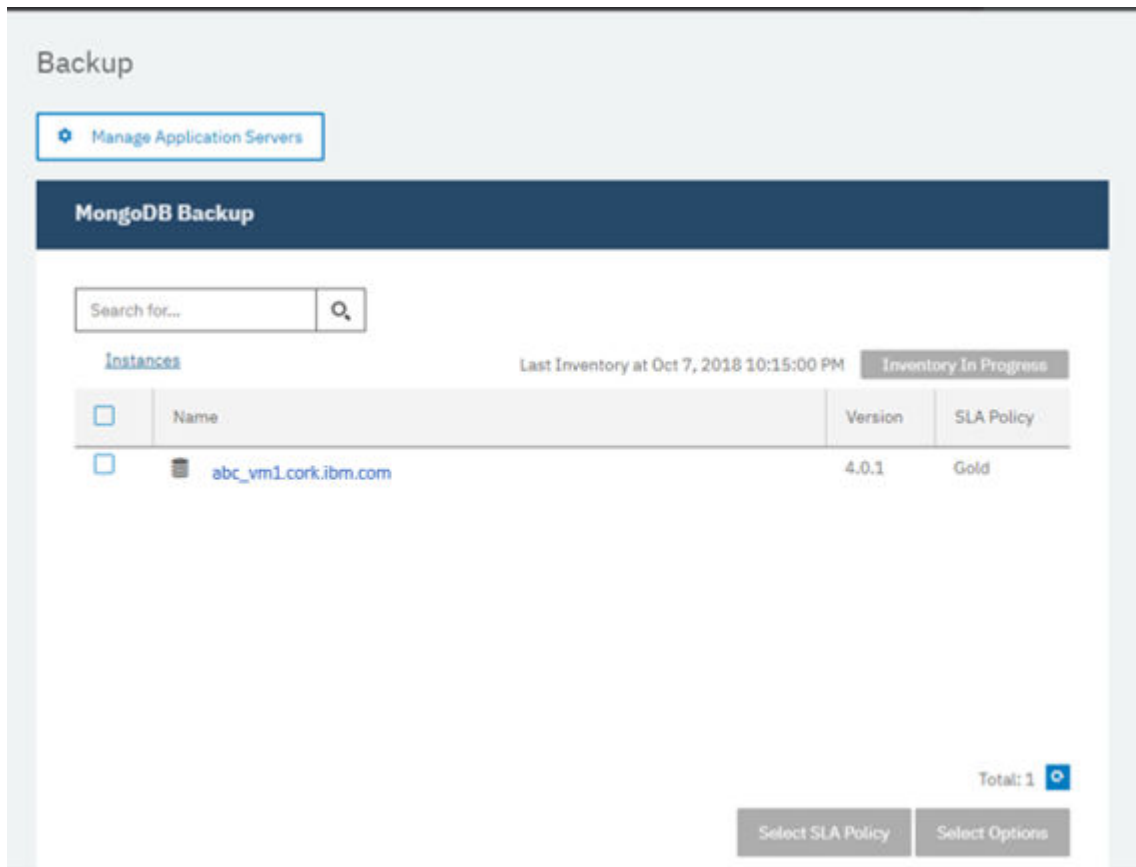
After you add your MongoDB application servers to IBM Spectrum Protect Plus, an inventory is run automatically to detect all MongoDB instances and databases. You can run a manual inventory on any application server to detect, list, and store all MongoDB databases for the selected host.

Before you begin

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

Procedure

1. In the navigation pane, expand **Manage Protection > Applications > MongoDB > Backup**.
Tip: To add more MongoDB instances to the **Instances** pane, follow the instructions in [Adding a MongoDB application server](#).
2. Click **Run Inventory**.



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

To monitor the inventory job, go to **Jobs and Operations**, click the **Policy and Job List** tab, and look for the newest Application Server Inventory log entry.

3. Click an instance to open a view that shows the databases that are detected for that instance. If any databases are missing from the **Instances** list, check your MongoDB application server and rerun the inventory. In some cases, certain databases are marked as ineligible for backup; hover over the database to reveal the reason why.

Tip: To return to the list of instances, click the **Instances** link in the **Backup MongoDB** pane.

What to do next

To start protecting MongoDB databases that are cataloged in the selected instance, apply a service level agreement (SLA) policy to the instance. For instructions about setting an SLA policy, see [Defining an SLA policy](#).

Testing the MongoDB connection

After you add a MongoDB application server, you can test the connection. The test verifies communication between IBM Spectrum Protect Plus and the MongoDB server. It also checks that the correct sudo permissions area available for the user who is running the test.

Procedure

1. In the navigation pane, click **Manage Protection > Applications > MongoDB > Backup**.
2. In the **Backup** window, click **Manage Application Servers**, and select the host address that you want to test.

A list of the MongoDB application servers that are available is shown.

- Click **Actions** and choose **Test** to start the verification tests for physical and remote system connections and settings.

1. Physical - Basic Test for physical host network configuration

| Name | Description | Status | Message |
|---------------------------|---|--------|---------|
| Host FQDN Resolvable Test | Host FQDN must be resolvable to an IPv4 address | ✓ | |
| Socket Connection Test | Must allow socket connection on port 22 for Linux | ✓ | |

2. Remote - Remote executor test for session creation and remote agent deployment

| Name | Description | Status | Message |
|---------------------------|--|--------|---------|
| Remote Session Test | Latest remote agent must be installed on host, SSH and SFTP service must be installed on Linux host, and port must be open to create session to SSH service. | ✓ | |
| Remote Agent Execute Test | Remote agent must be configured correctly using user credentials with sufficient privileges. | ✓ | |

3. LINUX - Basic Linux prerequisites for file and volume operations

| Name | Description | Status | Message |
|-----------------|--|--------|---------|
| Sudo Privileges | User must have password-less sudo privileges | ✓ | |

OK

The test report displays a list that includes tests for the physical host network configuration, and tests for the remote server installation on the host.

- Click **OK** to close the test report. If issues are reported, fix the issues and rerun the test to verify the fixes.

Backing up MongoDB data

Define regular MongoDB backup jobs with options to run and create backup copies to protect your data. To regularly back up your data, define a backup job that includes a service level agreement (SLA) policy.

Before you begin

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

Review the following prerequisites before you create a backup job definition:

- Add the application servers that you want to back up. For the procedure, see [Adding a MongoDB application server](#).
- Configure an SLA Policy. For the procedure, see [Defining a Service Level Agreement backup job](#).
- Before an IBM Spectrum Protect Plus user can set up backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources, and backup and restore operations, in the **Accounts** pane. For more information, see [Chapter 13, “Managing user access,” on page 245](#) and [“Roles for MongoDB” on page 161](#).

- **Restriction:** Do not run inventory jobs at the same time that backup jobs are scheduled.

Procedure

1. From the navigation menu, expand **Manage Protection > Applications > MongoDB > Backup**.
2. Select the check box for the instance that you want to back up.

Under each MongoDB instance, data to be backed up is listed as **ALL**. Each instance in the Instances pane is listed by instance name, version, and the applied SLA policy.

3. Click **Select Options** to specify the number of parallel streams for the backup operation. By selecting an appropriate number of parallel streams, you can minimize the time that is required for the backup job.
4. Click **Save**.

The saved options are used for all backup jobs for this instance as selected.

5. Select the instance again, and click **Select SLA Policy** to choose an SLA policy.
6. Save the SLA selection.

To define a new SLA or to edit an existing policy with custom retention and frequency rates, select **Manage Protection > Policy Overview**. In the **SLA Policies** pane, click **Add SLA Policy**, and define policy preferences.

What to do next

After the SLA policy is saved, you can run the policy at any time by clicking **Actions** beside the policy name and selecting **Start**. The status in the log changes to show that the backup job is in the Running state.

To cancel a job that is running, click **Actions** beside the policy name and select **Cancel**. A message asks whether you want to keep the data that is already backed up. Choose **Yes** to keep the backed up data, or **No** to discard the backup.

Defining a regular service level agreement job

After your MongoDB instances are listed, select and apply an SLA policy to start protecting your data.

Procedure

1. From the navigation menu, expand **Manage Protection > Applications > MongoDB > Backup**.
2. Select the MongoDB instance to back up all the data in that instance.

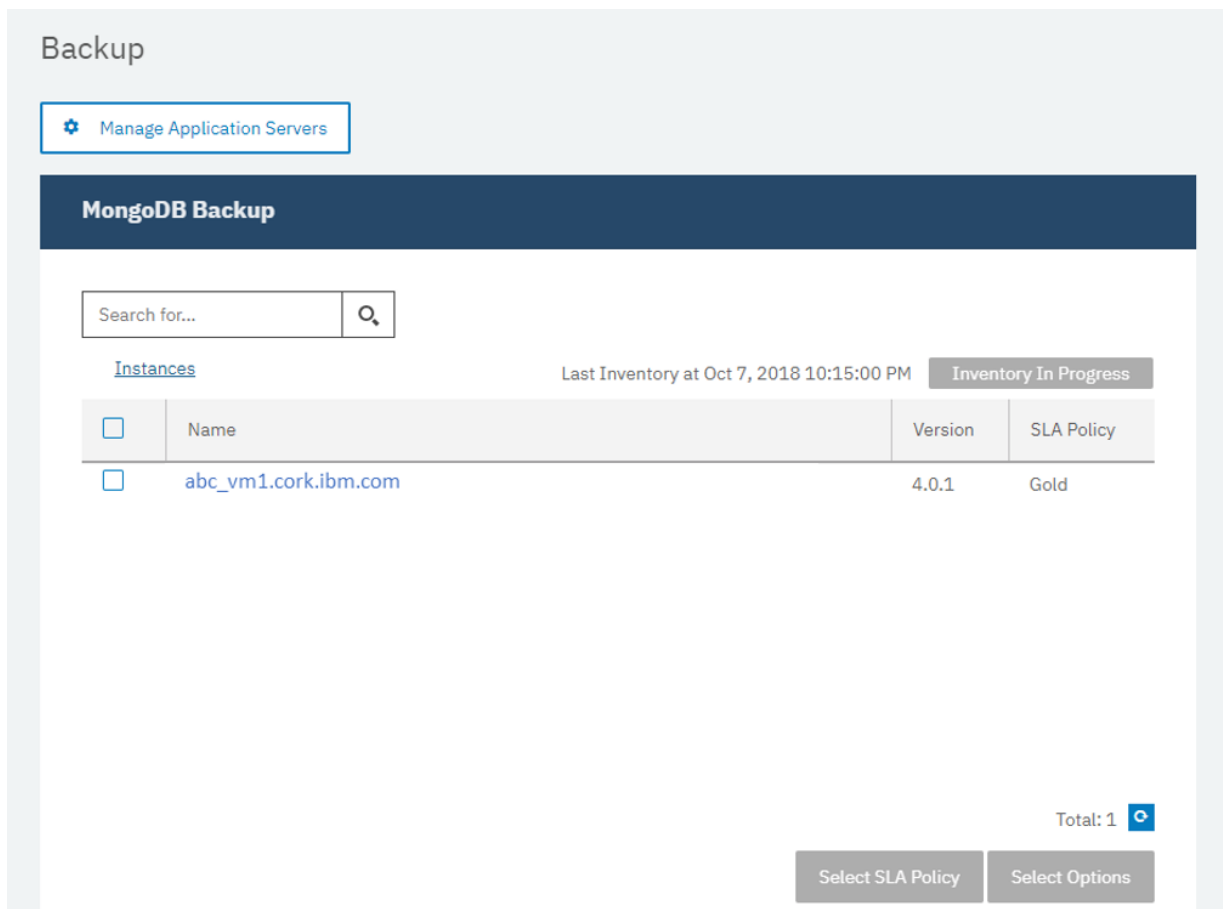


Figure 42. MongoDB Backup pane showing instances

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. Optional: To enable multiple backup streams to reduce the time that is taken to back up large databases, click **Select Options** and enter a number of parallel streams. Save your changes.

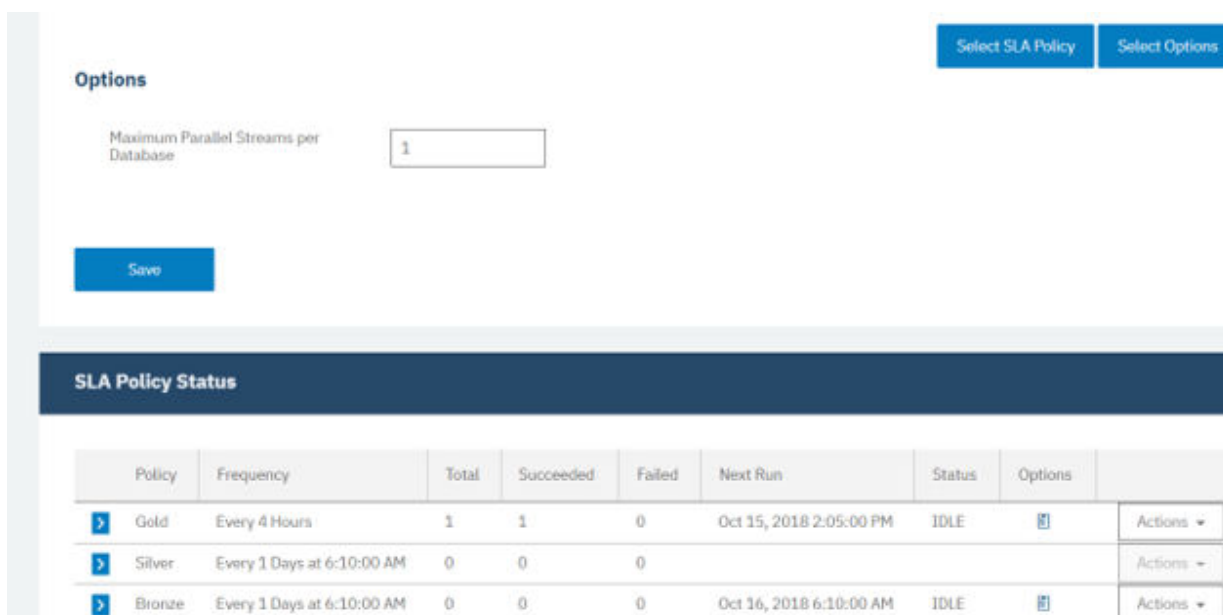


Figure 43. Backup options and SLA policies

5. Configure the SLA policy by clicking the icon in the **Options** column of the **SLA Policy Status** table.

For more information about SLA configuration options, see [“Setting SLA configuration options for your backup”](#) on page 169.

6. To run the policy outside of the scheduled job, select the instance. Click the **Actions** button and select **Start**. The status changes to **Running** for your chosen SLA and you can follow the progress of the job in the log shown.

What to do next


After the SLA policy is saved, you can run the policy at any time by clicking **Actions** beside the policy name and selecting **Start**. The status in the log changes to show that the backup job is in the Running state.

To cancel a job that is running, click **Actions** beside the policy name and select **Cancel**. A message asks whether you want to keep the data that is already backed up. Choose **Yes** to keep the backed up data, or **No** to discard the backup.

Setting SLA configuration options for your backup

After you set up a service level agreement (SLA) policy for your backup job, you can choose to configure extra options for that job. Additional SLA options include running scripts, and forcing a full base backup.

Procedure

1. In the **Policy Options** column of the **SLA Policy Status** table for the job that you are configuring, click the clipboard icon  to specify additional configuration options. If the job is already configured, click on the icon to edit the configuration.




Figure 44. Specifying additional SLA configuration options

2. Click **Pre-Script** and define the prescript configuration by choosing one of the following options:
 - Click **Use Script Server** and select an uploaded script from the menu.
 - Do not click **Use Script Server**. Select an application server from the list to run the script at that location.
3. Click **Post-Script** and define the PostScript configuration by choosing one of the following options:

- Click **Use Script Server** and select an uploaded script from the menu.
- Do not click **Use Script Server**. Select an application server from the list to run the script at that location.

Scripts and script servers are configured on the **System Configuration > Script** page. For more information about working with scripts, see [Configuring scripts](#).

4. To continue running the job when the script that is associated with the job fails, select **Continue job/task on script error**.

If this option is selected, the backup or restore operation is reattempted after an initial fail, and the script task status is reported as COMPLETED when the script completes processing with a nonzero return code. If this option is not selected, the backup or restore is not reattempted and the script task status is reported as FAILED.

5. Skip **Exclude Resources** for MongoDB SLA options, as you cannot specify resources to exclude. Instances are backed up rather than individual databases.
6. To create a full, new backup of a MongoDB instance, select **Force full backup of resources**.

A full new backup of that resource is created to replace the existing backup of that resource for one occurrence only. After that the resource is backed up incrementally as before.

Restoring MongoDB data

To restore data, define a job that restores data to the latest backup or select an earlier backup copy. Choose to restore data to the original instance or to an alternative instance on a different machine, creating a cloned copy. Define and save the restore job to run as an on-demand operation, or define a regularly run scheduled job.

Before you begin

In the **Restore** window, 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 instance, search by using the search field to find a specific instance name.

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

- At least one MongoDB backup job is set up and running successfully. For instructions about setting up a backup job, see [“Backing up MongoDB data” on page 166](#).
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For instructions about assigning roles, see [Chapter 13, “Managing user access,” on page 245](#), and [“Roles for MongoDB” on page 161](#).

Before you start, ensure that you allowed enough disk space at the target machine for the restore operation, and that you allocated dedicated volumes for file copying. Ensure that the same directory structure and layout are available on both the target and source machines. For restore operations to alternative instances, MongoDB must be at the same version level on the target and host machines. For more information about space requirements, see [Space prerequisites for MongoDB protection](#). For more information about prerequisites and setup, see [Prerequisites for MongoDB](#).

The IP address and port for the restored server can be found in the log file for the restore operation. Navigate to **Jobs and Monitoring > Monitoring** to find the logs for your restore operation.

Procedure

1. In the navigation menu, expand **Manage Protection > Applications > MongoDB > Restore**.
2. In the **MongoDB Restore** pane, click the instance that you want to restore.
3. Expand the data name **ALL** to show the restore points that are available for that instance.

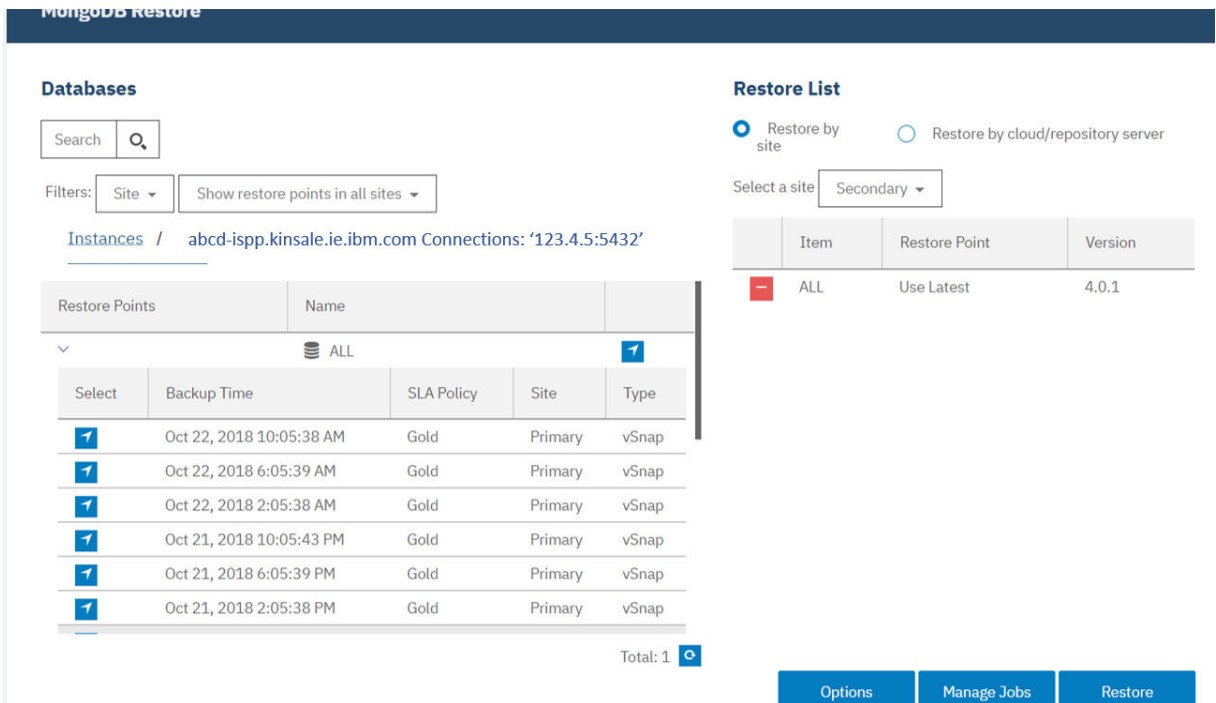


Figure 45. Restore panel.

Restore points are backup copies from a specific date and time. They are listed with the backup date and time, SLA policy, site information, and type.

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

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

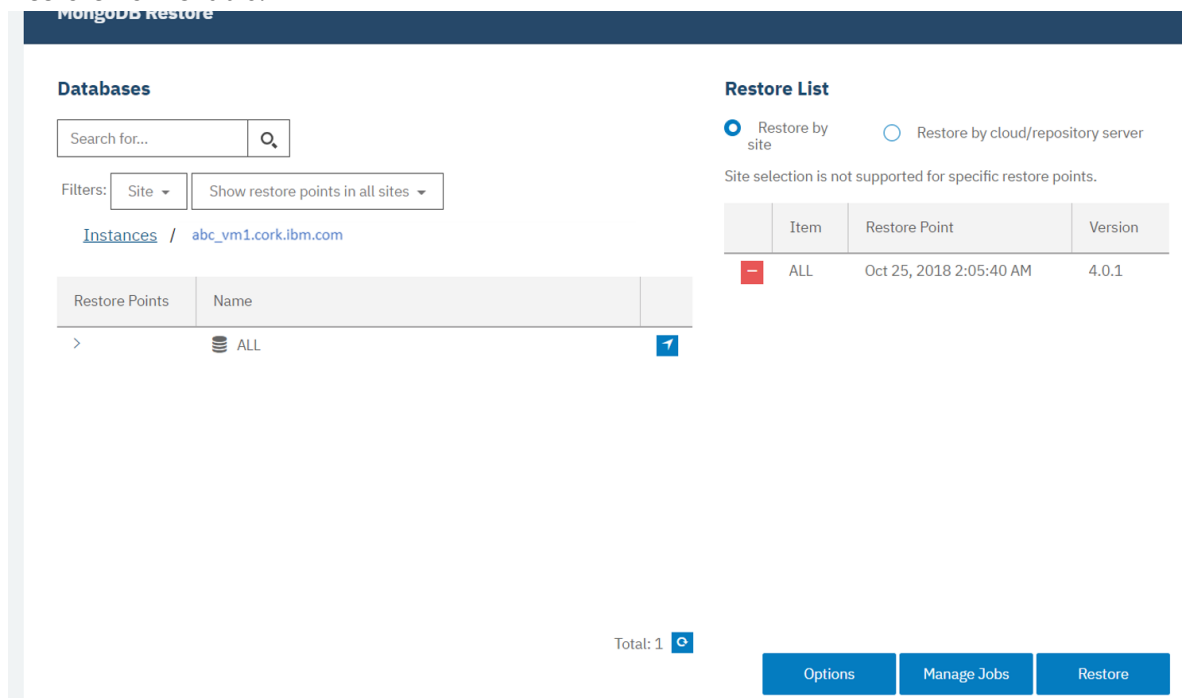


Figure 46. Selecting a database for a restore operation.

- 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 .

Restore

MongoDB Restore

Databases

Search for...

Filters: Site Show restore points in all sites

[Instances](#) / [abc_vm8.ibm.bcoll.ie](#)

| Restore Points | Name | | | |
|-------------------------------------|--------------------------|------------|---------|-------|
| ALL | 1 | | | |
| Select | Backup Time | SLA Policy | Site | Type |
| <input checked="" type="checkbox"/> | Oct 25, 2018 2:05:37 PM | Gold | Primary | vSnap |
| <input checked="" type="checkbox"/> | Oct 25, 2018 10:05:37 AM | Gold | Primary | vSnap |
| <input checked="" type="checkbox"/> | Oct 25, 2018 6:05:37 AM | Gold | Primary | vSnap |
| <input checked="" type="checkbox"/> | Oct 25, 2018 2:05:40 AM | Gold | Primary | vSnap |
| <input checked="" type="checkbox"/> | Oct 24, 2018 10:05:38 PM | Gold | Primary | vSnap |
| <input checked="" type="checkbox"/> | Oct 24, 2018 6:05:38 PM | Gold | Primary | vSnap |


Restore List

Restore by site Restore by cloud/repository server

Select a site

| Item | Restore Point | Version |
|---|---------------|---------|
| Select at least one resource on the left and add to the Restore list. | | |

Figure 47. Restore backup point.

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



Attention: Review the selected options before clicking **Restore** as data will be overwritten when the **Overwrite existing data** application option is selected. You can cancel a restore job when it is in progress, but when the **Overwrite existing data** option is selected, data is overwritten even if you cancel the job.

5. Click **Options** to edit options before you save the restore job definition.

- **Restore Type:** Choose one of the following options:
 - **Test:** In this mode, the agent creates a database with the MongoDB data files directly from the vSnap repository. This option is available only when you are restoring to an alternative instance. Members of replica sets will not be reconfigured after the MongoDB server is started. The server is started as a single-node replica set.
 - **Production:** In this mode, the MongoDB application server first copies the files from the vSnap repository to the target host whether that is 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 are restoring data to the original instance. MongoDB instances that are members of a replica set are not started during a production restore operation. This action prevents data from being overwritten when connecting to the replica set.
 - **Instant Access:** In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the share. Use the data for custom recovery from the files in the vSnap repository. This option is available only when you are restoring data to an alternative instance.
- **Destination:** Choose **Restore to original instance** to restore to the original server, or **Restore to alternate instance** to restore to a different location that you can select from the locations listed.

For more information about restoring data to the original instance, see [Restoring to the original instance](#). For more information about restoring your data to an alternative instance, see [Restoring to an alternate instance](#).

6. In the **Recovery Options** section, the **No Recovery** for MongoDB is selected by default, as recovery options from logs are not part of the backup or restore process.

7. Optional: Edit **Application Options**.

Choose **Overwrite existing databases** during a restore operation to replace existing data during the recovery. If this option is not selected, the restore job fails when data with the same name is found during the restore process.

Note: Ensure that no other data shares the same local database directory as the original data or it is overwritten when this choice is selected.

8. Edit **Advanced Options** 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 in case the recovery fails.

- **Allow session overwrite**

Select this option to replace existing databases with the same name during a restore operation. During an instance disk restore operation, the existing database is shut down and overwritten, and then the recovered database is restarted. If this option is not selected and a database with the same name is encountered, the restore operation fails with an error.

- **Continue with restores of other selected databases even if one fails**

If one database in the instance fails to restore successfully, the process continues for all other data being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.

- **Mount Point Prefix**

For **Instant Access** restore operations, specify a mount point prefix for the path where the mount is to be directed.

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

- Select **Pre-script** to select an uploaded script and an application or script server where the script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Scripts and script servers are configured through the **System Configuration > Script** page
- Select **Post-script** to select an uploaded script and an application or script server where the script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Scripts and script servers are configured through the **System Configuration > Script** page.
- 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 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. When all the options are chosen, click **Save**.

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


Figure 48. Restoring to alternative vSnap.

To cancel the job, navigate to **Jobs and Operations** and click the **Policy and Job List** tab. Find the restore job you want to cancel. Click **Actions**, and select **Cancel**.

When the test restore operation is listed in the **Active Resources** pane, select **Actions > Cancel** to cancel that process. If the status does not update, click **Refresh** to update the list.

Results

A few moments after you select **Restore**, the **onDemandRestore** is added to the **Job Sessions** pane. Expand the record to show the step-by-step details of the operation. You can also download the log file

by selecting the download icon . For any other jobs, navigate to **Jobs and Operations** and click the **Policy and Job List** tab. Find the restore job and expand the log to view the details.

For information about restoring data to the original instance, see [Restoring to the original instance](#). For information about restoring your data to an alternative instance, see [Restoring to an alternate instance](#).

Restoring MongoDB data to the original instance

Restore a MongoDB instance on the original host. Choose between restoring to the latest backup or an earlier MongoDB database backup version. When you restore data to its original instance, you cannot rename it. This restore option runs a full production restoration of data, and existing data is overwritten at the target site if the **Overwrite existing databases** application option is selected.

Before you begin

In the **Restore** window, 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 instance, search by using the search field to find a specific instance name.

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

- At least one MongoDB backup job is set up and running successfully. For instructions about setting up a backup job, see [“Backing up MongoDB data” on page 166](#).
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For instructions about assigning roles, see [Chapter 13, “Managing user access,” on page 245](#), and [“Roles for MongoDB” on page 161](#).

Before you start, ensure that you allowed enough disk space at the target machine for the restore operation, and that you allocated dedicated volumes for file copying. Ensure that the same directory structure and layout are available on both the target and source machines. For restore operations to alternative instances, MongoDB must be at the same version level on the target and host machines. For more information about space requirements, see [Space prerequisites for MongoDB protection](#). For more information about prerequisites and setup, see [Prerequisites for MongoDB](#).

The IP address and port for the restored server can be found in the log file for the restore operation. Navigate to **Jobs and Monitoring > Monitoring** to find the logs for your restore operation.

Procedure


1. In the navigation menu, expand **Manage Protection > Applications > MongoDB > Restore**.
2. In the **MongoDB Restore** pane, click the instance that you want to restore.
3. Expand the data name **ALL** to show the restore points that are available for that instance.

The screenshot shows the 'MongoDB Restore' interface. On the left, under 'Databases', there is a search box and filters for 'Site' and 'Show restore points in all sites'. Below this, the instance path is shown as 'Instances / abcd-isp.kinsale.ie.ibm.com Connections: '123.4.5:5432''. The main area displays a table of 'Restore Points' for the 'ALL' database. The table has columns for 'Select', 'Backup Time', 'SLA Policy', 'Site', and 'Type'. There are six rows of restore points, all with a 'Gold' SLA policy and 'Primary' site. On the right, the 'Restore List' section shows radio buttons for 'Restore by site' (selected) and 'Restore by cloud/repository server'. Below this is a 'Select a site' dropdown set to 'Secondary'. A table below shows a single entry: 'ALL' with 'Use Latest' as the restore point and '4.0.1' as the version. At the bottom right, there are buttons for 'Options', 'Manage Jobs', and 'Restore'. A 'Total: 1' indicator is visible at the bottom of the restore points table.

| Select | Backup Time | SLA Policy | Site | Type |
|--------------------------|--------------------------|------------|---------|-------|
| <input type="checkbox"/> | Oct 22, 2018 10:05:38 AM | Gold | Primary | vSnap |
| <input type="checkbox"/> | Oct 22, 2018 6:05:39 AM | Gold | Primary | vSnap |
| <input type="checkbox"/> | Oct 22, 2018 2:05:38 AM | Gold | Primary | vSnap |
| <input type="checkbox"/> | Oct 21, 2018 10:05:43 PM | Gold | Primary | vSnap |
| <input type="checkbox"/> | Oct 21, 2018 6:05:39 PM | Gold | Primary | vSnap |
| <input type="checkbox"/> | Oct 21, 2018 2:05:38 PM | Gold | Primary | vSnap |

Figure 49. Restore panel.

Restore points are backup copies from a specific date and time. They are listed with the backup date and time, SLA policy, site information, and type.

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





Figure 50. Selecting a database for a restore operation.

- 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 .

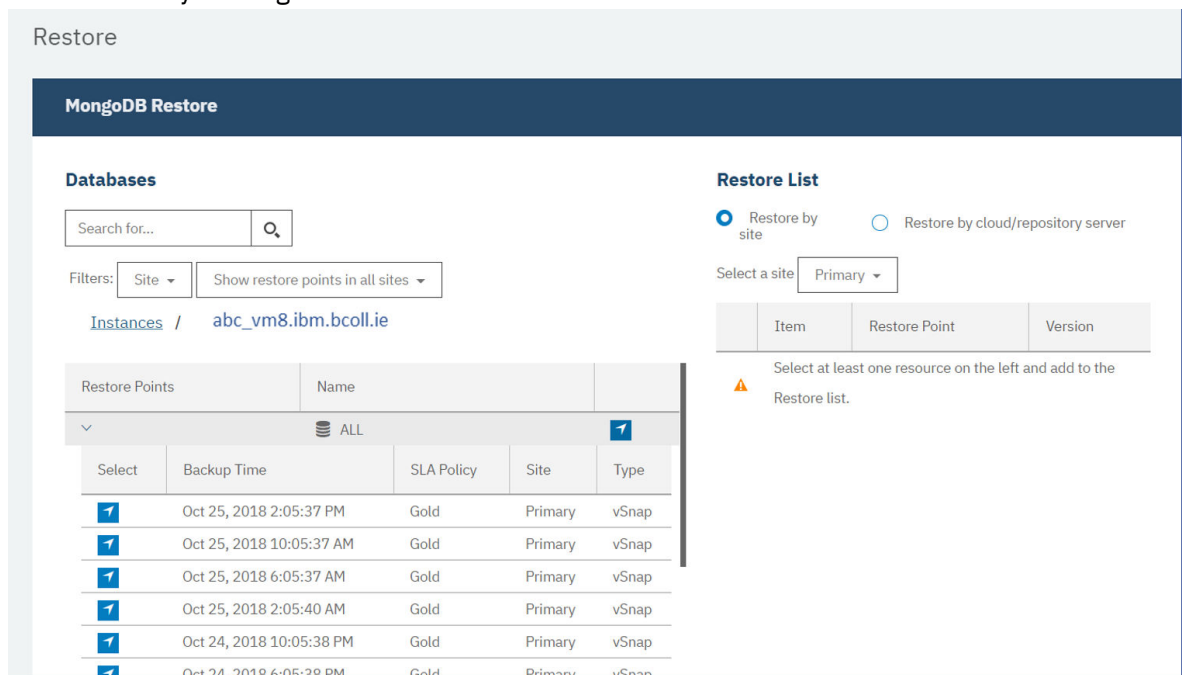



Figure 51. Restore backup point.

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



Attention: Review the selected options before clicking **Restore** as data will be overwritten when the **Overwrite existing data** application option is selected. You can cancel a restore job when it is in progress, but when the **Overwrite existing data** option is selected, data is overwritten even if you cancel the job.

- Click **Options** to edit options before you save the restore job definition.

6. Choose a **Restore Type** as follows:

- **Production**

For restore operations to the original instance when you need to recover the entire instance, it is best to choose this option with the overwrite application option chosen. MongoDB instances that are members of a replica set are not started during a production restore operation. This action prevents data from being overwritten when connecting to the replica set.

- **Test**

Choosing this option for a restore operation to the original instance, restores data to the same server but uses a different port.

- **Instant Access**

Choosing this option for a restore operation to the original instance, mounts the backup to the application server without restoring the data or overwriting the data.

7. Choose **Restore to original instance** to restore the data to the original server.

The screenshot shows the 'Restore Type' configuration page. On the left is a dark sidebar with navigation icons. The main content area is white with a blue border. It contains several sections: 'Restore Type' with a radio button for 'Production'; 'Destination' with radio buttons for 'Restore to original instance' and 'Restore to alternate instance'; 'Alternate vSnap' with a checked checkbox 'Use alternate vSnap server for the restore job' and a dropdown menu for 'Select alternate vSnap' showing 'localhost'; and 'Recovery Options' with radio buttons for 'No Recovery', 'Recover until end of backup', 'Recover until end of available logs', and 'Recover until specific point'. Below the 'Recovery Options' section are three input fields for time, each containing '00'.

8. In the **Recovery Options** section, the **No Recovery** for MongoDB is selected by default, as recovery options from logs are not part of the backup or restore process.

9. Optional: Edit **Application Options**.

Choose **Overwrite existing databases** during a restore operation to replace existing data during the recovery. If this option is not selected, the restore job fails when data with the same name is found during the restore process.

Note: Ensure that no other data shares the same local database directory as the original data or it is overwritten when this choice is selected.

10. Edit **Advanced Options** 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 in case the recovery fails.

- **Allow session overwrite**

Select this option to replace existing databases with the same name during a restore operation. During an instance disk restore operation, the existing database is shut down and overwritten, and then the recovered database is restarted. If this option is not selected and a database with the same name is encountered, the restore operation fails with an error.

- **Continue with restores of other selected databases even if one fails**

If one database in the instance fails to restore successfully, the process continues for all other data being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.

- **Mount Point Prefix**

For **Instant Access** restore operations, specify a mount point prefix for the path where the mount is to be directed.

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

- Select **Pre-script** to select an uploaded script and an application or script server where the script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Scripts and script servers are configured through the **System Configuration > Script** page
- Select **Post-script** to select an uploaded script and an application or script server where the script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Scripts and script servers are configured through the **System Configuration > Script** page.
- 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 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. When all the options are chosen, click **Save**.

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

The screenshot shows a configuration page for a restore job. On the left is a dark sidebar with icons for home, info, settings, jobs, and user. The main content area has the following sections:

- Name:** A text input field containing "DB1".
- New Database Name:** A text input field containing "COPYNEWDB".
- Alternate vSnap:**
 - A section header.
 - A checkbox labeled "Use alternate vSnap server for the restore job" which is checked.
 - A dropdown menu labeled "Select alternate vSnap" with "localhost" selected.
- Recovery Options:**
 - Four radio button options: "No Recovery", "Recover until end of backup" (selected), "Recover until end of available logs", and "Recover until specific point-in-time".
 - Below the radio buttons are three time input fields, each containing "00" and a dropdown arrow.

Figure 52. Restoring to alternative vSnap.

To cancel the job, navigate to **Jobs and Operations** and click the **Policy and Job List** tab. Find the restore job you want to cancel. Click **Actions**, and select **Cancel**.

Restoring MongoDB data to an alternative instance

Select a MongoDB instance backup for restoring to an alternative host. You can also choose to restore an instance to a different vSnap repository, or you can rename the resource. This process creates an exact copy of the instance on a different host.

Before you begin

Before you start, ensure that you allowed enough disk space and allocated dedicated volumes for the copying of files.

Tip: When you are restoring data to an alternative instance and the **Overwrite existing databases** option is not selected, the process fails when any data on the local database directory you are restoring to is found. No other data can share the local database target directory as the database to be restored. When the **Overwrite existing databases** option is selected, any existing data is wiped and the local database is cleaned up.

Ensure that the following requirements are in place before you start.

- The target database directory that you are restoring is clean and ready for the restore operation. If you selected the **Overwrite existing database** option, data in the target database directory is overwritten.
- At least one MongoDB backup job is set up and running. For instructions about setting up a backup job, see [Defining an SLA job](#).
- 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 245 and [“Roles for MongoDB”](#) on page 161.

About this task

Important: If your MongoDB database runs on a port other than the default port 27017, a restore operation to the original location with the overwrite option selected, restores the MongoDB database to port 27017. It does not restore the database to the same port that it was running on in the original instance. You can stop the restored instance and restart it with the MongoDB configuration file to run it on the original port.

Procedure

1. In the navigation menu, expand **Manage Protection > Applications > MongoDB > Restore**.
2. In the **MongoDB Restore** pane, click the instance that you want to restore.
3. Expand the data name **ALL** to show the restore points that are available for that instance.

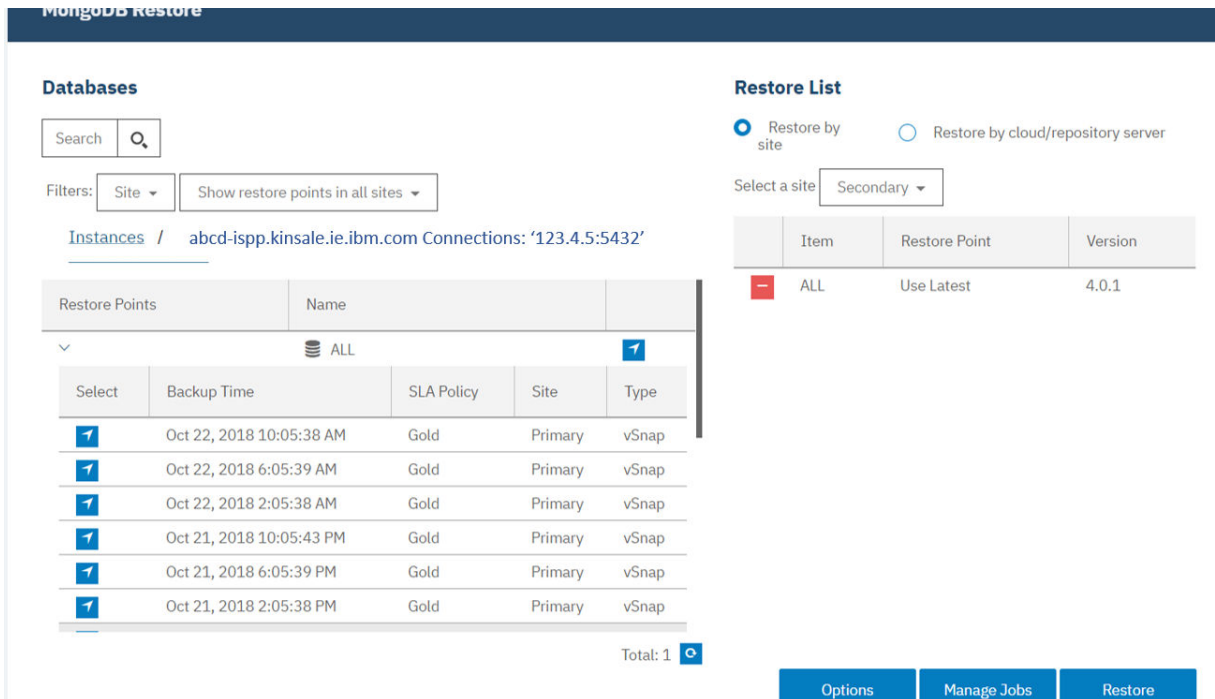


Figure 53. Restore panel.

Restore points are backup copies from a specific date and time. They 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 database with the MongoDB data files directly from the vSnap repository. This option is available only when you are restoring to an alternative instance. Members of replica sets will not be reconfigured after the MongoDB server is started. The server is started as a single-node replica set.
 - **Production:** In this mode, the MongoDB application server first copies the files from the vSnap repository to the target host whether that is 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 are restoring data to the original instance. MongoDB instances that are members of a replica set are not started during a production restore operation. This action prevents data from being overwritten when connecting to the replica set.
 - **Instant Access:** In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the share. Use the data for custom recovery from the files in the vSnap repository. This option is available only when you are restoring data to an alternative 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 versions levels cannot be selected. Other instances on the same host as the original instance, cannot be selected either.

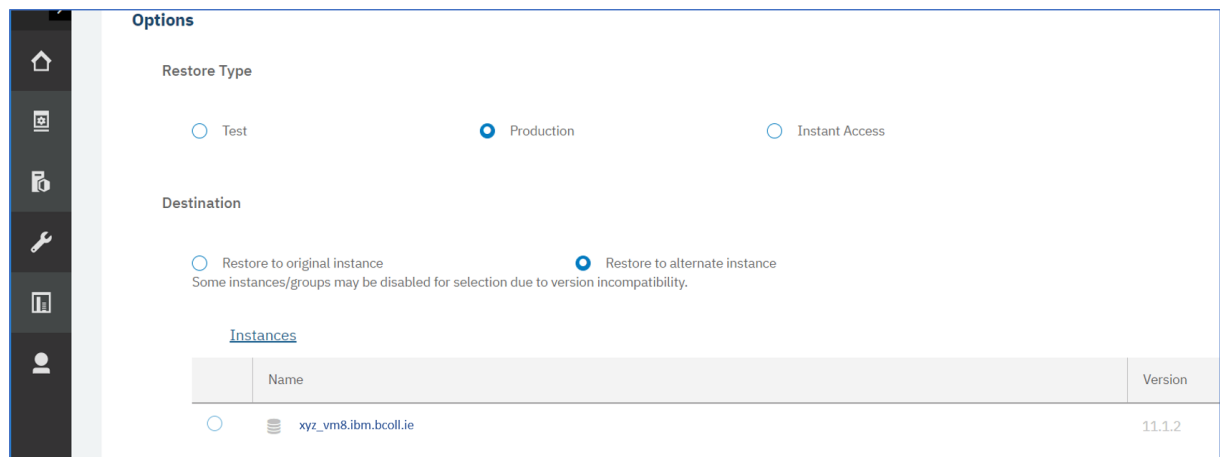


Figure 54. Restoring to alternate instance and Instances pane

7. Optional: To rename the restored database, enter a new name in the **New Database Name** field.
8. Choose an alternative vSnap server for the restore operation when you are restoring a restore point that is not the most recent copy. Select **Use alternate vSnap server for the restore job**, and choose a server from the menu.
9. Optional: Edit **Application Options**.

Choose **Overwrite existing databases** during a restore operation to replace existing data during the recovery. If this option is not selected, the restore job fails when data with the same name is found during the restore process.

Note: Ensure that no other data shares the same local database directory as the original data or it is overwritten when this choice is selected.

10. Edit **Advanced Options** 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 in case the recovery fails.
- **Allow session overwrite**
Select this option to replace existing databases with the same name during a restore operation. During an instance disk restore operation, the existing database is shut down and overwritten, and then the recovered database is restarted. If this option is not selected and a database with the same name is encountered, the restore operation fails with an error.

- **Continue with restores of other selected databases even if one fails**

If one database in the instance fails to restore successfully, the process continues for all other data being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.

- **Mount Point Prefix**

For **Instant Access** restore operations, specify a mount point prefix for the path where the mount is to be directed.

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

- Select **Pre-script** to select an uploaded script and an application or script server where the script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Scripts and script servers are configured through the **System Configuration > Script** page
- Select **Post-script** to select an uploaded script and an application or script server where the script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Scripts and script servers are configured through the **System Configuration > Script** page.
- 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 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. When all the options are chosen, click **Save**.

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

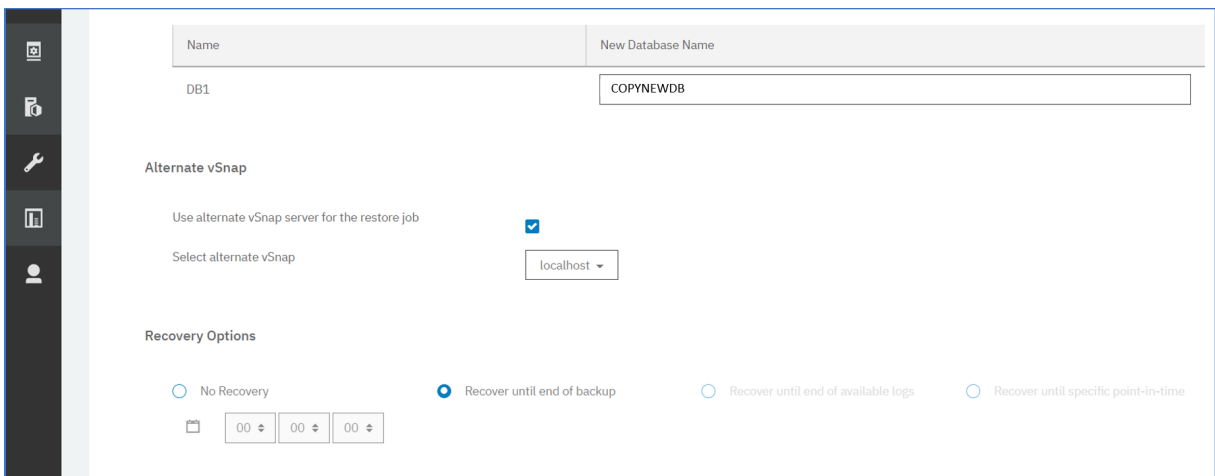



Figure 55. Restoring to alternative vSnap.

To cancel the job, navigate to **Jobs and Operations** and click the **Policy and Job List** tab. Find the restore job you want to cancel. Click **Actions**, and select **Cancel**.

When the test restore operation is listed in the **Active Resources** pane, select **Actions > Cancel** to cancel that process. If the status does not update, click **Refresh** to update the list.

Results

A few moments after you select **Restore**, the **onDemandRestore** is added to the **Job Sessions** pane. Expand the record to show the step-by-step details of the operation. You can also download the log file

by selecting the download icon . For any other jobs, navigate to **Jobs and Operations** and click the **Policy and Job List** tab. Find the restore job and expand the log to view the details.

For information about restoring data to the original instance, see [Restoring to the original instance](#). For information about restoring your data to an alternative instance, see [Restoring to an alternate instance](#).

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 37.

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 250 . |

Related concepts

[“Managing user access” on page 245](#)

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 184](#)

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

[“Restoring SQL Server data” on page 187](#)

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 LUN volume and creates an NTFS share on that iSCSI LUN. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus agent maps the LUN to the SQL Server server and mounts the NTFS volume to where the backup is completed. If log backups are enabled IBM Spectrum Protect Plus creates a separate vSnap volume and creates a CIFS on that volume. Log backup transaction files are copied to this share according to the schedule created for log backup.

When the backup 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 245](#).
- 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.
- Prior to copying logs to vSnap, SPP uses the backup folder configured for the SQL server instance as the staging area to collect logs. The volume where this folder is located must have sufficient space available to hold all of the transaction logs between backup jobs. The staging area can be modified by changing the backup folder configuration using SQL Server Management Studio (SSMS).
- 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 SQL Servers that you want to back up. For more information, see [“Adding a SQL Server application server”](#) on page 182.
- Configure SLA policies. For more information, see [“Create backup policies”](#) on page 62.
- 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 an 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** and then clicking the **Policy and Job List** tab.
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 a 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.

SQL log files are temporarily stored in a local staging area before they are copied to a CIFS share. The SQL server default backup destination serves as the staging area and must have enough free space to temporarily store the transaction log files before they can be copied to the CIFS share.

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 **Policy Options** field that is associated with the job in the **SLA Policy Status** section. Set the additional policy options:

Pre-scripts and post-scripts

Pre-scripts and post-scripts are scripts that can be run before or after a job runs. Batch and PowerShell scripts are supported.

In the **Pre-script** or **Post-script** section, select an uploaded script and an application or script server where the script 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 on the **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

Exclude Resources

Exclude specific resources from the backup job through single or multiple exclusion patterns. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *.

Separate multiple filters with a semicolon.

Force Full Backup of Resources

Force base 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 187 . |

Related concepts

[“Configuring scripts for backup and restore operations” on page 208](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and Batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Starting jobs” on page 206](#)

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 an SQL backup job. For more information, see [“Backing up SQL Server data”](#) on page 184.
- Review the SQL Server system requirements in [“Microsoft SQL Server requirements”](#) on page 37.
- 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 245.
- When completing a production restore to an 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 an 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 the iSCSI protocol 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 an 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, click the add to restore list icon  at the resource level. Select **Restore by site** or **Restore by cloud/repository server**. If restoring from a Site, click the **Select a site** drop-down menu to choose a site associated with the backup storage server you want to restore from. If restoring from a cloud or repository server, the restore source will be automatically selected.

To select a restore point that is not the latest, expand a resource in the **Restore** pane, and then click the add to restore list icon  that is associated with the restore point. Adding a combination of latest restore points and non-latest restore points to the Restore List is not supported. Click the delete icon  to remove restore points from the **Restore List**.

Additional filtering options are available when viewing non-latest restore points. To view available restore points from sites, cloud resources, or repository servers, expand a resource in the **Restore** pane, then select the source type through the **Filters** menu. Once a source type is selected, all available restore points associated with the source type display. To view restore points on a specific resource, select it from the drop-down menu adjacent to the source type drop-down menu. For example, if the filter is set to **Sites**, click **Show restore points in all sites** to select a specific site.

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 an 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 an SQL Always On environment using test mode, the source availability database is restored to the selected target instance.

When performing restore operations in an 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.

Alternate vSnap

When restoring from a restore point that was offloaded to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server used to complete the restore is the same vSnap server used to complete the backup and offload operations. To reduce load, an alternate vSnap server can be selected to serve as the gateway to complete the restore. To select an alternate vSnap server when restoring a specific, non-latest restore point from a cloud resource or repository server, select **Use alternate vSnap server for the restore job**, then select a server from the **Select alternate vSnap** menu.

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 an SQL backup job definition, point-in-time restore options will be available when creating an 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 an SQL Always On environment using test mode, the restored database will join the instance where the availability group resides.

When performing restore operations in an 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.

Note: Before performing restore operations in a SQL Always On environment using **production** mode with the **Overwrite existing database** option, ensure that the database is not present on the replicas of the target availability group. To do so, you must manually clean up the original databases (to be overwritten) from all replicas of the target availability group.

Maximum Parallel Streams per Database

Set the maximum 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.

Mount Point Prefix

For instant access restore operations, specify the prefix for the path where the mount point is to be directed.

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 **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

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 208

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and Batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Adding a SQL Server application server” on page 182](#)

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 184](#)

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 34](#).

Adding an Oracle application server

When an Oracle application server is added, an inventory of the instances and databases that are associated with the application server is captured and added to IBM Spectrum Protect Plus. This process enables you to complete backup and restore jobs, as well as run reports.

Procedure

To register an Oracle application server, complete the following steps.

1. In the navigation pane, click **Manage Protection > Applications > Oracle > Backup**.
2. Click **Manage Application Servers**.
3. Click **Add Application Server** to add the host machine.
4. In the **Application Properties** pane, enter the host address.

The host address is a resolvable IP address, or a resolvable path and machine name.

5. Select **User of SSH key**.

When you choose to use a user, select an existing user or enter a user ID and password. When you choose to use an SSH key, pick the key from the menu. The user must have sudo privileges set up.

6. Populate the fields in the **Applications Properties** pane as follows:

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.

7. 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 250 . |

Related concepts

[“Managing user access” on page 245](#)

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 193](#)

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

[“Restoring Oracle data” on page 195](#)

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 245](#).
- 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 191](#).
- Configure SLA policies. For more information, see [“Create backup policies” on page 62](#).

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** and then clicking the **Policy and Job List** tab.
5. To edit options before you create the job definition, click **Select Options**. Set the job definition options.

Enable Log Backup

Enable Log Backup must be selected to allow for Oracle point-in-time restore.

Select **Enable Log Backup** to permit IBM Spectrum Protect Plus to automatically create a log backup volume and mount it to the application server. IBM Spectrum Protect Plus then automatically

discovers the location of the existing primary archived log and uses cron to configure a scheduled job. The scheduled job completes a transaction log backup from the primary location to that log backup volume at the frequency specified through the **Frequency** setting.

The **Frequency** can be set to a value independent of the database backup frequency specified in the SLA Policy settings. For example, the SLA Policy may be configured to back up the database once per day while the log backup frequency could be set to once per 30 minutes.

For Oracle 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 based on the retention settings in the SLA policy.

Select **Truncate source logs after successful backup** to automatically delete older archived logs from the database's primary archived log location. If the option is cleared, archived logs on the primary log destination are not deleted, and Database Administrators must continue to manage those logs using their existing log retention policies. If the option is selected, IBM Spectrum Protect Plus deletes older unneeded archived logs from the primary log location at the end of every successful database backup.

When the option **Truncate source logs after successful backup** is selected, set the retention of primary logs through the **Primary log retention in days** setting. This setting controls the quantity of archived logs that are retained in the primary archived log locations. For example, if **Primary log retention in days** is set to **3**, IBM Spectrum Protect Plus deletes all archived logs older than three days from the primary archived log location at the end of every successful database backup.

Maximum Parallel Streams per Database

Set the maximum data stream per database to the backup storage. This setting applies to each database in the job definition. Multiple databases can be backed up in parallel if the value of the option is set to **1**. Multiple parallel streams might improve backup speed, but high bandwidth consumption might affect overall system performance.

6. When you are satisfied that the job-specific information is correct, click **Save**.
7. To configure additional options, click the **Policy Options** field that is associated with the job in the **SLA Policy Status** section. Set the additional policy options:

Pre-scripts and Post-scripts

Pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job level. Windows-based machines support Batch and PowerShell scripts while Linux-based machines support shell scripts.

In the **Pre-script** or **Post-script** section, select an uploaded script and an application or script server where the script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Scripts and script servers are configured through the **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

Exclude Resources

Exclude specific resources from the backup job through single or multiple exclusion patterns. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *.

Separate multiple filters with a semicolon.

Force Full Backup of Resources

Force base backup operations for specific virtual machines or databases in the backup job definition. Separate multiple resources with a semicolon.

What to do next

After you create the backup job definition, complete the following action:

| Action | How to |
|--|---|
| Create an Oracle Restore job definition. | See “Restoring Oracle data” on page 195 . |

Related concepts

[“Configuring scripts for backup and restore operations” on page 208](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and Batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Restoring Oracle data

Use a restore job to restore 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 193](#).
- Review the Oracle system requirements in [“Oracle requirements” on page 34](#)
- 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 245](#).
- 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, click the add to restore list icon  at the resource level. Select **Restore by site** or **Restore by cloud/repository server**. If restoring from a Site, click the **Select a site** drop-down menu to choose a site associated with the backup storage server you want to restore from. If restoring from a cloud or repository server, the restore source will be automatically selected.

To select a restore point that is not the latest, expand a resource in the **Restore** pane, and then click the add to restore list icon  that is associated with the restore point. Adding a combination of latest restore points and non-latest restore points to the Restore List is not supported. Click the delete icon  to remove restore points from the **Restore List**.

Additional filtering options are available when viewing non-latest restore points. To view available restore points from sites, cloud resources, or repository servers, expand a resource in the **Restore** pane, then select the source type through the **Filters** menu. Once a source type is selected, all available restore points associated with the source type display. To view restore points on a specific resource, select it from the drop-down menu adjacent to the source type drop-down menu. For example, if the filter is set to **Sites**, click **Show restore points in all sites** to select a specific site.

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.

Alternate vSnap

When restoring from a restore point that was offloaded to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server used to complete the restore is the same vSnap server used to complete the backup and offload operations. To reduce load, an alternate vSnap server can be selected to serve as the gateway to complete the restore. To select an alternate vSnap server when restoring a specific, non-latest restore point from a cloud resource or repository server, select **Use alternate vSnap server for the restore job**, then select a server from the **Select alternate vSnap** menu.

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.

Mount Point Prefix

For instant access restore operations, specify the prefix for the path where the mount point is to be directed.

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 **System Configuration > Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

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 208](#)

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and Batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

[“Adding an Oracle application server” on page 191](#)

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 reduce system load, ensure other jobs are not scheduled to run during the IBM Spectrum Protect Plus backup job. To create an SLA policy, see [“Creating an SLA policy” on page 79](#).

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. Catalog backup data can also be offloaded to cloud resources and repository servers.
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** and then clicking the **Policy and Job List** tab. For instructions, see [“Start a job” on page 69](#).

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.

If restoring an offloaded cloud backup, the cloud resource or repository server must be registered on the alternate IBM Spectrum Protect Plus location.

Procedure

To restore IBM Spectrum Protect Plus data:

1. In the navigation pane, click **Manage Protection > IBM Spectrum Protect Plus > Restore**.
2. Select a vSnap server, cloud resource, or repository server.

Data can be restored to the same location, or an alternate location in disaster recovery scenarios.

Available snapshots for the server are displayed.

3. Click **Restore** for the catalog snapshot that you want to restore.
4. Select one of the following restore modes:

Restore the catalog and suspend all scheduled jobs

The catalog is restored and all scheduled jobs are left in a suspended state. No scheduled jobs are started, which allows for the validation and testing of catalog entries and the creation of new jobs. Typically, this option is used in DevOps use cases.

Restore the catalog

The catalog is restored and all scheduled jobs continue to run as captured in the catalog backup. Typically, this option is used in disaster recovery.

5. Click **Restore**.
6. To run the restore job, in the dialog box, click **Yes**.

Managing IBM Spectrum Protect Plus restore points


You can use the **Restore Point Retention** pane to search for restore points in the IBM Spectrum Protect Plus catalog by backup job name, view their creation and expiration dates, and override the assigned retention.

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 job session to expire:

1. In the navigation pane, click **Manage Protection > IBM Spectrum Protect Plus > Restore Point Retention**.
2. In the Backup Sessions tab, search for the desired job session or restore point. For more information about using the search function, see [Appendix A, "Search guidelines,"](#) on page 261. Alternatively, in the Virtual Machines / Databases tab, select either Applications or Hypervisors to search for the desired catalog entry by entering the name. Names can be searched by entering partial text, using the asterisk (*) as a wildcard character, or using the question mark (?) for pattern matching.
3. If searching from the Backup Sessions tab, use filters to fine-tune your search across job types and date range when the associated backup job started.
4. Click the search icon .
5. Select the job sessions you want to expire.
6. From the **Actions** list, select one of the following options:
 - **Expire** is used to expire a single job session.
 - **Expire All Job Sessions** is used to expire all unexpired job sessions for the selected job.

7. To confirm the expiration, in the dialog box, click **Yes**.

Results

The job session is removed during the next run of the maintenance job.

Related concepts

[“Job types” on page 205](#)

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

Deleting IBM Spectrum Protect Plus resources from the catalog



You can use the **Virtual Machines/Databases** tab in the **Restore Point Retention** pane to expire catalog metadata that is associated with a resource from the IBM Spectrum Protect Plus catalog. Resources are added to the catalog through inventory jobs. Expiring a resource removes the metadata that is associated with a restore point from the catalog, which frees up space in the catalog and removes the restore point from recovery screens.

About this task

Expiring a resource from the catalog does not remove associated snapshots from a vSnap server or secondary backup storage.

Procedure

To expire a resource from the catalog:

1. In the navigation pane, click **Manage Protection > IBM Spectrum Protect Plus > Restore Point Retention**.
2. Click the **Virtual Machines/Databases** tab.
3. Use the filter to search by resource type, then enter a search string to search for a resource by name. For more information about using the search function, see [Appendix A, “Search guidelines,” on page 261](#).
4. Click the search icon .
5. Click the delete icon  that is associated with a resource.
6. To confirm the expiration, in the dialog box, click **Yes**.

Results

The catalog metadata associated with the resource is removed from the catalog.

Related concepts

[“Job types” on page 205](#)

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** and then click the **Policy and Job List** tab.

The name, type, status, last run time, and next scheduled run time (if applicable) are shown for each job.

To view logs for a job, expand the job.

To download a .zip file that contains logs, 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** and then click the **Policy and Job List** tab.
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** and then click the **Policy and Job List** tab.
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** and then click the **Policy and Job List** tab.
2. To cancel a running job session, click the **Actions** menu that is associated with the job, and then click **Cancel**.

Rerunning partially completed backup jobs

If the last instance of a backup job was partially completed, you can rerun the job to back up virtual machines and databases that were skipped.

About this task

A backup job can be rerun only in the same session ID as the original partially completed backup job. No successful backup of the same resource can have completed since the partial backup job you choose to rerun.

Note: Backup jobs can be rerun only in response to a hypervisor or database backup failure. The following events do not qualify for backup job rerun operations:

- A VM backup was completed with an FLI failure.
- A snapshot condense failure occurred for a storage system.
- A backup job failed with an unknown issue such as a cataloging error.
- A resource is missing from the vCenter.

For applications for which log backups are supported, log backups are not disabled when using the rerun feature. Log backups will be disabled for the applicable databases when the job is next started without using the on-demand backup or rerun feature.

Procedure

Complete the following steps to rerun a partially completed backup operation:

1. In the navigation pane, click **Jobs and Operations** and then click the **Job History** tab.
2. Use the search function and filters to find the last instance of the backup job that was partially completed.
3. Select the job instance, then click **Rerun**.

Note:

If the backup job cannot be rerun, the **Rerun** option is not available.

All SLA options and any exclusions that are associated with the original job are included in the rerun operation. No option or exclusion changes are applied since the partial backup completed. If the rerun job is completed successfully, the job summary is updated to show success.

Backing up a single resource

If a hypervisor or application server is associated with an SLA policy, a single virtual machine or application can be backed up immediately by running an on-demand backup operation. Select **Run** in a hypervisor or application server backup screen to run an on-demand backup operation. This option is enabled when an existing SLA policy is associated with the resource.

About this task

Rerunning a backup job for a single resource is applicable only for backup operations, not replication or offload operations.

For applications for which log backups are supported, log backups are not disabled when using the on-demand backup or rerun feature. Log backups will be disabled for the applicable databases when the job is next started without using the on-demand backup or rerun feature.

Procedure

Complete the following steps to run an on-demand backup job of a single virtual machine or application server:

1. In the navigation pane, click **Manage Protection**. Depending on the type of backup operation, select **Hypervisors > Backup**, or **Applications > Backup**.
2. Click one of the listed instances to show the associated virtual machine or application resources.
The hypervisor or application server must be associated with an existing SLA policy.
3. Click **Run**.
If the virtual machine or application is a member of multiple SLA policies, select the SLA policy to run for the on-demand job.
4. To confirm the backup job, in the dialog box, click **OK**.

Configuring scripts for backup and restore operations

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and Batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Before you begin

Review the following considerations for using scripts with hypervisors:

- The user who is running the script must have the **Log on as a service** right enabled, which is required for running prescripts and postscripts. For more information about this right, see [Add the Log on as a service Right to an Account](#).
- Windows Remote Shell (WinRM) must be enabled.

Uploading a script

Supported scripts include shell scripts for Linux-based machines and Batch and PowerShell scripts for Windows-based machines. Scripts must be created using the associated file format for the operating system.

Procedure

Complete the following steps to upload a script:

1. In the navigation pane, click **System Configuration > Script**.
2. In the **Scripts** section, click **Upload Script**.
The **Upload Script** pane 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 208 . |

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, click **System Configuration > Script**.

2. In the **Script Servers** section, click **Add Script Server**.

The **Script Server Properties** pane displays.

3. Set the server options.

Host Address

Enter the resolvable IP address or a resolvable path and machine name.

Use existing user

Enable to select a previously entered user name and password for the provider.

Username

Enter your username for the provider. If entering a SQL server, the user identity follows the default *domain\name* format if the virtual machine is attached to a domain. The format *local_administrator* is used if the user is a local administrator.

Password

Enter your password for the provider.

OS Type

Select the operating system of the application server.

4. Click **Save**.

Chapter 11. Configuring and maintaining the IBM Spectrum Protect Plus system environment

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

Maintenance tasks include reviewing the configuration of the IBM Spectrum Protect Plus virtual appliance, collecting log files for troubleshooting, and managing Secure Sockets Layer (SSL) certificates.

In most cases, IBM Spectrum Protect Plus is installed on a virtual appliance. The virtual appliance contains the application and the inventory. Maintenance tasks are completed in vSphere Client, by using the IBM Spectrum Protect Plus command line, or in a web-based management console.

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

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

The vSnap server is the primary backup location for snapshots and all IBM Spectrum Protect Plus environments have at least one vSnap server. You can also optionally offload snapshots to cloud storage or a repository server for longer-term protection.

Managing cloud storage

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

Adding Amazon S3 cloud storage as a backup storage provider

Add Amazon S3 cloud storage to enable IBM Spectrum Protect Plus to offload data to S3.

Before you begin

Configure the key that is required for the cloud object. For instructions, see [“Adding an access key” on page 221](#).

Ensure that there are cloud storage buckets created for the IBM Spectrum Protect Plus data before you add the cloud storage in the following steps. For information how to create buckets, see [Amazon Simple Storage Service Documentation](#).

Procedure

To add Amazon S3 cloud storage as backup storage provider, complete the following steps:

1. In the navigation menu, click **System Configuration > Backup Storage > Cloud**.
2. Click **Add Cloud**.
3. From the **Provider** list, select **Amazon S3**.
4. Complete the fields in the **Cloud Registration** pane:

Name

Enter a meaningful name to help identify the cloud storage.

Region

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

Use existing key

Enable to select a previously entered key for the storage, and then select the key from the **Select a key** list.

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

Key name

Enter a meaningful name to help to identify the key.

Access key

Enter the AWS access key. Access keys are created through the AWS Management Console.

Secret key

Enter the AWS secret key. Secret keys are created through the AWS Management Console.

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

6. Click **Register**.

The cloud storage is added to the cloud servers table.

What to do next

After you add the S3 storage, complete the following action:

| Action | How to |
|--|--|
| Associate the cloud storage with the SLA policy that is used for the backup job. | To create an SLA policy, see “Creating an SLA policy” on page 79. To modify an existing SLA policy, see “Editing an SLA policy” on page 81. |

Adding IBM Cloud Object Storage as a backup storage provider

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

Before you begin

Configure the key and certificate that are required for the cloud object. For instructions, see [“Adding an access key”](#) on page 221 and [“Adding a certificate”](#) on page 222.

Ensure that there are cloud storage buckets created for the IBM Spectrum Protect Plus data before you add the cloud storage in the following steps. For information how to create buckets, see [About IBM Cloud Object Storage](#).

When adding IBM Cloud Object Storage (COS), the method for obtaining the access and secret key will depend on the deployment model. If on-premise, keys can be obtained from the IBM COS Manager Console. For IBM COS IaaS, keys are created when a service account is created and can be obtained from the softlayer portal. If using IBM COS (COS as a Service), the access and secret key are not created by default; when a service account is created, check the **Include HMAC Credential** box, and add `{“HMAC”: true}` to the **Add Inline Configuration Parameters** text area.

Procedure

To add IBM Cloud Object Storage as a backup storage provider, complete the following steps:

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

Name

Enter a meaningful name to help identify the cloud storage.

Endpoint

Select the endpoint of the cloud storage.

Use existing key

Enable to select a previously entered key for the storage, and then select the key from the **Select a key** list.

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

Key name

Enter a meaningful name to help to identify the key.

Access key

Enter the access key.

Secret key

Enter the secret key.

Certificate

Select a method of associating a certificate with the resource:

Upload

Select to browse for the certificate locally, then click **Upload**.

Copy and paste

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

Use existing

Select to use a previously uploaded certificate.

A certificate is not required if you are adding public IBM Cloud Object Storage.

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

6. Click **Register**.

The cloud resource is added to the cloud servers table.

What to do next

After you add the IBM Cloud Object Storage, complete the following action:

| Action | How to |
|--|--|
| Associate the cloud storage with the SLA policy that is used for the backup job. | To create an SLA policy, see “Creating an SLA policy” on page 79. To modify an existing SLA policy, see “Editing an SLA policy” on page 81. |

Adding Microsoft Azure cloud storage as a backup storage provider

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

Before you begin

Ensure that there are cloud storage buckets created for the IBM Spectrum Protect Plus data before you add the cloud storage in the following steps. For information how to create buckets, see Azure documentation.

Procedure

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

1. In the navigation pane, click **System Configuration > Backup Storage > Cloud**.

2. Click **Add Cloud**.
3. From the **Provider** list, select **Microsoft Azure Blob Storage**.
4. Complete the fields in the **Cloud Registration** pane:

Name

Enter a meaningful name to help identify the cloud storage.

Endpoint

Select the endpoint of the cloud storage.

Use existing key

Enable to select a previously entered key for the storage, and then select the key from the **Select a key** list.

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

Key name

Enter a meaningful name to help identify the key.

Storage Account Name

Enter the Microsoft Azure access storage account name. This is from the Azure Management Portal.

Storage Account Shared Key

Enter the Microsoft Azure key from any one of the key fields in the Azure Management Portal, either key1 or key2.

5. Click **Get Buckets**, and then select a bucket to serve as the offload target.
6. Click **Register**.
The cloud storage is added to the cloud servers table.

What to do next

After you add the Microsoft Azure storage, complete the following action:


| Action | How to |
|--|--|
| Associate the cloud storage with the SLA policy that is used for the backup job. | To create an SLA policy, see “Creating an SLA policy” on page 79 . To modify an existing SLA policy, see “Editing an SLA policy” on page 81 . |

Editing settings for cloud storage

Edit the settings for a cloud storage provider to reflect changes in your cloud environment.

Procedure

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

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


Deleting cloud storage

Delete a cloud storage provider to reflect changes in your cloud environment. Ensure that the provider is not associated with any SLA policies before deleting the provider.

Procedure

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

1. In the navigation menu, click **System Configuration > Backup Storage > Cloud**.

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

Managing repository server storage

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

Configuring IBM Spectrum Protect server as an offload target

To offload to an IBM Spectrum Protect server, you must first set up IBM Spectrum Protect Plus as an object client to the server.

About this task

You must add the object client by using IBM Spectrum Protect Operations Center. After you complete the object client set up, keys and a certificate are provided to enable secure connection to the IBM Spectrum Protect server. This information is required to add the repository server in IBM Spectrum Protect Plus.

To add the object client, you must be familiar with the IBM Spectrum Protect server environment and have experience working with the Operations Center and the IBM Spectrum Protect server administrative commands. If you require assistance, contact your IBM Spectrum Protect administrator.

IBM Spectrum Protect Plus is aware of offloads to IBM Spectrum Protect server, but is not aware of subsequent IBM Spectrum Protect server replication operations.

Related tasks

[“Adding a repository server as a backup storage provider” on page 219](#)

Add a repository server to enable IBM Spectrum Protect Plus to offload data to the server.

Preparing to offload data from IBM Spectrum Protect Plus

Before you offload data from IBM Spectrum Protect Plus to IBM Spectrum Protect, complete the preparation steps in the IBM Spectrum Protect environment.

Procedure

1. Verify that you can open an IBM Spectrum Protect server port to the IBM Spectrum Protect Plus object client that you plan to use for data offload operations. The default port number is 9000. If there are any firewalls between the object client and agent, configure the object agent to access the appropriate port through the firewall.
2. Verify the settings for the policy domain that you plan to use for data offload operations. An object client node is associated with this policy domain when the node is registered or updated by using the IBM Spectrum Protect server administrative commands **REGISTER NODE** or **UPDATE NODE**.

Considerations for specifying policy domains for IBM Spectrum Protect Plus offload operations include the following:

- The domain that the node is assigned to must have a backup copy group. Objects that are stored to an object client node are always backup objects. An archive copy group is not required.
- You must use a container storage pool. The storage pool that is specified in the copy group Copy Destination must be either a directory-container or cloud-container storage pool.
- All objects are uniquely named. There are no inactive versions of objects, so you can set the Versions Data Exists field to 1.
- Backup copy groups contain only active versions, so you can set the Retain Extra Versions and Retain Only Version fields to 0.
- The IBM Spectrum Protect server controls the time when objects are deleted. Ensure that the object client node is enabled to allow backup copy group deletion.

Example: Display detailed information about a policy domain for an IBM Spectrum Protect Plus offload operation

Display settings for a copy group for an object client node.

```
query copygroup format=detailed
```

```
Policy Domain Name: TAPSRV03_OBJECT
Policy Set Name: SET1
Mgmt Class Name: BACK_DISK
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 1
Versions Data Deleted: 0
Retain Extra Versions: 0
Retain Only Version: 0
Copy Mode: Modified
Copy Serialization: Shared Static
Copy Frequency: 0
Copy Destination: DEDUPPOOL
Table of Contents (TOC) Destination:
Last Update by (administrator): JBASIL
Last Update Date/Time: 01/17/2019 14:38:05
Managing profile:
Changes Pending: No
```

Offloading data to AIX systems

You can offload data from IBM Spectrum Protect Plus to an IBM Spectrum Protect server on AIX.

About this task

An IBM Spectrum Protect object agent cannot run directly on an IBM AIX operating system. However, you can offload IBM Spectrum Protect Plus data to an IBM Spectrum Protect object client on an AIX system by first setting up an object agent on a Linux x86_64 operating system. The standalone object agent is available only on the Linux x86_64 operating system.

After the IBM Spectrum Protect Plus object client sends data to the IBM Spectrum Protect object agent on Linux x86_64, the object agent transfers data to an IBM Spectrum Protect object client on AIX.

Procedure

To offload data from IBM Spectrum Protect Plus to an IBM Spectrum Protect server on AIX, complete the following steps:

1. On the AIX server, issue the following IBM Spectrum Protect server administrative command:

```
setopt EnableAIXS3Interface Yes
```

2. On the AIX server, define an object agent by issuing the following IBM Spectrum Protect server administrative command. To set the high-level address (HLA) and low-level address (LLA), use the IP address of the host system and port that the object agent will use.

```
define server object_agent_name
hla=object_agent_host_system_ip_address
lla=object_agent_port objectagent=yes
```

Tip: The default value for the object agent port is 9000. If a local object agent is already running on the system, the object agent that is being configured for the AIX server must use a different port number from that of the existing object agent.

3. Download the following scripts to the object agent host system:

- ftp://public.dhe.ibm.com/storage/tivoli-storage-management/maintenance/objectagent/v8r1/Linux/8.1.7.000/x86_64/spObjectAgent
- ftp://public.dhe.ibm.com/storage/tivoli-storage-management/maintenance/objectagent/v8r1/Linux/8.1.7.000/x86_64/startObjectAgent.sh

- ftp://public.dhe.ibm.com/storage/tivoli-storage-management/maintenance/objectagent/v8r1/Linux/8.1.7.000/x86_64/spObjectAgent.rc
- ftp://public.dhe.ibm.com/storage/tivoli-storage-management/maintenance/objectagent/v8r1/Linux/8.1.7.000/x86_64/spObjectAgent.rc.u
- ftp://public.dhe.ibm.com/storage/tivoli-storage-management/maintenance/objectagent/v8r1/Linux/8.1.7.000/x86_64/delObjectAgentSvc.sh

Either IBM Spectrum Protect Plus or an IBM Spectrum Protect server can be installed on the object agent host system.

If the IBM Spectrum Protect server is installed, you can use the `spObjectAgent` file in the server directory and you do not have to download the agent and its scripts again.

4. Ensure that the following files have executable permissions:

- `spObjectAgent`
- `startObjectAgent.sh`
- `spObjectAgent.rc`
- `spObjectAgent.rc.u`
- `delObjectAgentSvc.sh`

5. From the AIX server system, copy the following two items to a directory on the object agent host system on Linux:

- Object agent server directory
- Server public certificate

The object agent server directory was created when you ran the **DEFINE SERVER** command. The directory includes the following file and certificates:

- A configuration file for creating and starting an object agent service
- Certificates for communication between the object agent and the server

The object agent server directory is created in the server instance directory: `/server_instance_home_dir/object_agent_name`. For example,

`/home/tsminst1/OBJAGENT1`

The server public certificate (`cert256.arm`) is typically located in the server instance directory.

6. In the object agent server directory that you copied in the previous step, locate the object agent configuration file (`spObjectAgent_objectagentname_serverport.config`).

For example: `spObjectAgent_OBJAGENT1_1500.config`

In the configuration file, update the locations of the following files. For example:

```
objagentexe="/opt/tivoli/tsm/server/bin/spObjectAgent\"
keystore="/home/tsminst1/OBJAGENT1/agentcert.p12"
pwdfile="/home/tsminst1/OBJAGENT1/agentcert.pwd"
serverkeypub="/home/tsminst1/OBJAGENT1/cert256.arm"
agentconfig="/home/tsminst1/OBJAGENT1/spObjectAgent_OBJAGENT1_1500.config"
```

7. Override the **SERVERHLA** parameter in the object agent configuration file by using the AIX server IP address:

```
serverhla=aix_server_ip_address
```

Tip: The object agent uses this value to locate the IBM Spectrum Protect server.

8. To create and start the object agent on the host system, run the `startObjectAgent.sh` script with the configuration file:

```
startObjectAgent.sh spObjectAgent_objectagentname_serverport.config
```

9. Register an object agent client on the AIX server by issuing the following IBM Spectrum Protect server command:

```
register node nodename type=objectclient
```

Important: Record the login user ID and password that are automatically generated. You will need the credentials to connect to the object agent.

10. To connect the IBM Spectrum Protect Plus object client to the object agent, go to the IBM Spectrum Protect Plus online documentation and follow the instructions in [Adding a repository server as a backup storage provider](#).

Offloading data to Linux and Windows systems

You can offload data from IBM Spectrum Protect Plus to an IBM Spectrum Protect server on Linux or Windows.

Procedure

To offload data from IBM Spectrum Protect Plus to an IBM Spectrum Protect server on Linux or Windows, complete the following steps:

1. Set up an object agent.
 - a) On the Operations Center menu bar, click **Servers**.
 - b) Select a server row and click **Details**.
 - c) Select **Object Agent** in the left navigation pane and complete the steps to create an object agent and start an object agent service. To authenticate to the object agent, use the certificate that is generated.

Tip: Alternatively, use the IBM Spectrum Protect server administrative command **DEFINE SERVER** to create an object agent. Specify OBJECTAGENT=YES. Complete the configuration by starting an object agent service on the system that is hosting the IBM Spectrum Protect server.

2. Set up an object client.

Tip: If you create an object client before creating the corresponding object agent, the Add Client wizard forces the creation of the object agent.

- a) On the Operations Center menu bar, click **Clients**.
- b) In the Clients table, click **+Client**.
- c) Select Object Client and follow the instructions in the **Add Client** wizard.

After you complete the wizard, it provides you with the endpoint for communicating with the object agent on the server, and the access key ID and secret access key for connecting securely. When IBM Spectrum Protect Plus is used as an object client, it must direct its requests to the endpoint, and must use the access key ID and secret access key.

Tip: Alternatively, use the command **REGISTER NODE** to create an object client. Specify TYPE=OBJECTCLIENT.

Deleting an object agent service

When an object agent is deleted from the IBM Spectrum Protect server, the object agent service must be deleted from the host system. To complete the deletion process for an object agent, delete the corresponding service.

Before you begin

To delete the object agent service on a Linux operating system, you must run the `delObjectAgentSvc.sh` script with the object agent configuration file. Ensure that you can log in to the object agent host system with the root user ID.

To delete the object agent service on a Windows operating system, you must run the `delObjectAgentSvc.cmd` batch file with the object agent configuration file. Ensure that you have Windows administrator privileges to log on to the object agent host system.

Procedure

1. Verify that the object agent is deleted from the IBM Spectrum Protect server by issuing the server administrative command **QUERY SERVER**.
2. Open a command line.
3. Issue the following command on one line. The default server directories are used in the examples.

Linux

```
/opt/tivoli/tsm/server/bin/delObjectAgentSvc.sh  
/object_agent_config_path/spObjectAgent_objectagentname_server_port.config
```

Windows

```
"C:\Program Files\Tivoli\TSM\server\delObjectAgentSvc.cmd"  
"object_agent_config_path\spObjectAgent_objectagentname_server_port.config"
```

where

object_agent_config_path

Specifies the configuration path for the object agent.

objectagentname

Specifies the name of the object agent.

server_port

Specifies the port number of the IBM Spectrum Protect server.

Adding a repository server as a backup storage provider

Add a repository server to enable IBM Spectrum Protect Plus to offload data to the server.

Before you begin

Configure the key and certificate that are required for the repository server. For instructions, see [“Adding an access key”](#) on page 221 and [“Adding a certificate”](#) on page 222.

For the current release of IBM Spectrum Protect Plus, the repository server must be an IBM Spectrum Protect server.

Configure IBM Spectrum Protect Plus as an object client to the IBM Spectrum Protect server. The object client node transfers and stores offloaded data. After you complete the setup procedure, the wizard provides you with the endpoint for communicating with the object agent on the server, and the access ID, secret key, and certificate for connecting securely. [“Configuring IBM Spectrum Protect server as an offload target”](#) on page 215.

Certificates can be obtained from the IBM Spectrum Protect server Operations Center by navigating to the following pane: **Server > Object Agent > Agent Certificate**. Alternatively, the certificate can be obtained from the IBM Spectrum Protect Plus appliance by running the following command: `openssl s_client -showcerts -connect <ip-address>:9000 </dev/null 2>/dev/null | openssl x509`

Offload retention settings are fully controlled through associated SLA policies in IBM Spectrum Protect Plus. IBM Spectrum Protect server copygroup retention settings are not used for offload operations.

Procedure

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

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

Name

Enter a meaningful name to help identify the repository server.

Hostname

Enter the resolvable host name of the repository server. The host name must match the object agent name found in the associated certificate.

Port

Enter the communications port of the repository server.

Use existing key

Enable to select a previously entered key for the repository, and then select the key from the **Select a key** list.

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

Key name

Enter a meaningful name to help to identify the key.

Access key

Enter the access key.

Secret key

Enter the secret key.

Certificate

Select a method of associating a certificate with the resource. If copying the certificate, the BEGIN and END lines of text must be included.

Upload

Select to browse for the certificate locally, then click **Upload**.

Copy and paste

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

Use existing

Select to use a previously uploaded certificate.

4. Click **Register**.

The IBM Spectrum Protect server is added to the repository servers table.

What to do next

After you add a repository server, complete the following action:


| Action | How to |
|--|--|
| Associate the repository server with the SLA policy that is used for the backup job. | To create an SLA policy, see “Creating an SLA policy” on page 79. To modify an existing SLA policy, see “Editing an SLA policy” on page 81. |

Editing settings for a repository server

Edit the settings for a repository server provider to reflect changes in your cloud environment.

Procedure

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

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

The **Update Repository Server** pane is displayed.


3. Revise the settings for the repository server provider, and then click **Update**.

Deleting a repository server

Delete a repository server provider to reflect changes in your environment. Ensure that the provider is not associated with any SLA policies before deleting the provider.

Procedure

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

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

Managing keys and certificates

Cloud resources and repository servers require credentials to serve as offloading destinations. Access keys and secret keys are provided by your cloud resource or repository server interface. These keys serve as the user name and password of your offload destinations and allow them to be accessed by IBM Spectrum Protect Plus. Some offload destinations also require certificates for additional data security.

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

Adding an access key

Add an access key to provide cloud resource or repository server credentials.

Procedure

To add a key, complete the following steps:

1. Create your access key and secret key through the interface of the cloud resource or repository server. Make note of the access key and secret key.
2. In the navigation menu, click **System Configuration > Keys and Certificates**.
3. From the **Access Keys** section, click **Add Access Key**.
4. Complete the fields in the **Key Properties** pane:

Name

Enter a meaningful name to help identify the access key.

Access Key

Enter the access key of the cloud resource or repository server. For Microsoft Azure, enter the storage account name.

Secret Key

Enter the secret key of the cloud resource or repository server. For Microsoft Azure, enter the key from one of the key fields, either key1 or key2.

5. Click **Save**.


The key displays in the **Access Keys** table and can be selected when utilizing a feature that requires credentials to access a resource through the **Use existing key** option.

Deleting an access key

Delete an access key when it becomes obsolete. Ensure that you reassign a new access key to your cloud resource or repository server.

Procedure

To delete an access key, complete the following steps:

1. In the navigation menu, click **System Configuration > Keys and Certificates**.
2. Click the delete icon  that is associated with an access key.
3. Click **Yes** to delete the access key.

Adding a certificate

Add a certificate to provide cloud resource or repository server credentials.

Procedure

To add a certificate, complete the following steps:

1. Export a certificate from your cloud resource or repository server.
2. In the navigation menu, click **System Configuration > Keys and Certificates**.
3. In the **Certificates** section, click **Add Certificate**.
4. Complete the fields in the **Certificate Properties** pane:

Type

Select the cloud resource or repository server type.

Certificate

Select a method to add the certificate:

Upload

Select to browse for the certificate locally.

Copy and paste

Select to enter the name of the certificate and copy and paste the contents of the certificate.

5. Click **Save**.


The key displays in the **Certificates** table and can be selected when utilizing a feature that requires credentials to access a resource through the **Use existing certificate** option.

Deleting a certificate

Delete a certificate when it becomes obsolete. Ensure that you reassign a new certificate to your cloud resource or repository server.

Procedure

To delete a certificate, complete the following steps:

1. In the navigation menu, click **System Configuration > Keys and Certificates**.
2. Click the delete icon  that is associated with a certificate.
3. Click **Yes** to delete the certificate.

Adding an SSH key

Add an SSH key to provide credentials for Linux-based resources, including file indexing and restore operations on virtual machines under vCenter and Hyper-V, as well as Oracle, Db2, and MongoDB

application servers. SSH keys provide a secure connection between your resources and IBM Spectrum Protect Plus.

Before you begin

- The SSH service must be running on port 22 on the server and any firewalls must be configured to allow IBM Spectrum Protect Plus to connect to the server using SSH. The SFTP subsystem for SSH must also be enabled.
- Ensure the public SSH key is placed in the appropriate `authorized_keys` file for the IBM Spectrum Protect Plus agent user. Typically, the file is located at `/home/<username>/.ssh/authorized_keys`. The `.ssh` directory and all files under it must have their permissions set to 600.

Procedure

To add a key, complete the following steps:

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

Name

Enter a meaningful name to help identify the SSH key.

User

Enter the user associated with the resource and SSH key.

Private key

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

7. Click **Save**.


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

Deleting an SSH key

Delete an SSH key when it becomes obsolete. Ensure that you reassign a new SSH key to your resources.

Procedure

To delete an SSH key, complete the following steps:

1. In the navigation menu, click **System Configuration > Keys and Certificates**.
2. Click the delete icon  that is associated with an SSH key.
3. Click **Yes** to delete the access key.

Managing sites

A site is 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 253](#)

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

[“Scheduling a report” on page 242](#)

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.

About this task

You can add a Microsoft Active Directory or OpenLDAP server. Note that OpenLDAP does not support the `sAMAccountName` user filter that is commonly used with Active Directory. Additionally, the **memberOf** option must be enabled on the OpenLDAP server.

Procedure

To register an LDAP server, complete the following steps:

1. In the navigation pane, click **System Configuration > LDAP/SMTP**.
2. In the **LDAP Servers** pane, click **Add LDAP Server**.
3. Populate the following fields in the **LDAP Servers** pane:

Host Address

The IP address of the host or logical name of the LDAP server.

Port

The port on which the LDAP server is listening. The typical default port is 389 for non SSL connections or 636 for SSL connections.

SSL

Enable the SSL option to establish a secure connection to the LDAP server.

Use existing user

Enable to select a previously entered user name and password for the LDAP server.

Bind Name

The bind distinguished name that is used for authenticating the connection to the LDAP server. IBM Spectrum Protect Plus supports simple bind.

Password

The password that is associated with the Bind Distinguished Name.

Base DN

The location where users and groups can be found.

User Filter

A filter to select only those users in the Base DN that match certain criteria. An example of a valid default user filter is `cn={0}`.

Tips:

- To enable authentication by using the **sAMAccountName** Windows user naming attribute, set the filter to `samaccountname={0}`. When this filter is set, users log in to IBM Spectrum Protect Plus by using only a user name. A domain is not included.
- To enable authentication using the user principal name (UPN) naming attribute, set the filter to `userprincipalname={0}`. When this filter is set, users log in to IBM Spectrum Protect Plus by using the `username@domain` format.
- To enable authentication by using an email address that is associated with LDAP, set the filter to `mail={0}`.

The **User Filter** setting also controls the type of user name that appears in the IBM Spectrum Protect Plus display of users.

User RDN

The relative distinguished path for the user. Specify the path where user records can be found. An example of a valid default RDN is `cn=Users`.

Group RDN

The relative distinguished path for the group. If the group is at a different level than the user path, specify the path where group records can be found.

4. Click **Save**.

Results

IBM Spectrum Protect Plus completes the following actions:

1. Confirms that a network connection is made.
2. Adds the LDAP server to the database.

After the SMTP server is added, the **Add LDAP Server** button is no longer available.

What to do next

If a message is returned indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connections.

Related tasks

[“Creating a user account for an LDAP group” on page 253](#)

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

Adding an SMTP server

You must add an SMTP server to send scheduled reports to email recipients. Only one SMTP server can be associated with a IBM Spectrum Protect Plus virtual appliance.

Procedure

To add an SMTP server, complete the following steps:

1. In the navigation pane, click **System Configuration > LDAP/SMTP**.
2. In the **SMTP Servers** pane, click **Add SMTP Server**.
3. Populate the following fields in the **SMTP Servers** pane:

Host Address

The IP address of the host, or the path and host name of the SMTP server.

Port

The communications port of the server that you are adding. The typical default port is 25 for non-SSL connections or 443 for SSL connections.

Username

The name that is used to access the SMTP server.

Password

The password that is associated with the user name.

Timeout

The email timeout value in milliseconds.

From Address

The address that is associated with email communications from IBM Spectrum Protect Plus.

Subject Prefix

The prefix to add to the email subject lines sent from IBM Spectrum Protect Plus.

4. Click **Save**.

Results

IBM Spectrum Protect Plus completes the following actions:

1. Confirms that a network connection is made.
2. Adds the server to the database.

If a message is returned indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connections.

To test the SMTP connection, click the **Test SMTP Server** button, then enter an e-mail address. Click **Send**. A test e-mail message is sent to the e-mail address to verify the connection.

After the SMTP server is added, the **Add SMTP Server** button is no longer available.

What to do next

Related tasks

[“Scheduling a report” on page 242](#)


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

Editing settings for an LDAP or SMTP server

Edit the settings for an LDAP or SMTP server to reflect changes in your IBM Spectrum Protect Plus environment.

Procedure

To edit the settings for an LDAP or SMTP server, complete the following steps:


1. From the navigation menu, click **System Configuration > LDAP/SMTP**.
2. Click the edit icon  that is associated with the server.
The edit pane is displayed.
3. Revise the settings for the server, and then click **Save**.

Deleting an LDAP or SMTP server

Delete an LDAP or SMTP server when it becomes obsolete. Ensure that the server is not in use by IBM Spectrum Protect Plus before deleting the server.

Procedure

To delete an LDAP or SMTP server, complete the following steps:

1. From the navigation menu, click **System Configuration > LDAP/SMTP**.
2. Click the delete icon  that is associated with the server.
3. Click **Yes** to delete server.

Logging on to the administrative console

Log on to the administrative console to review the configuration of the IBM Spectrum Protect Plus virtual appliance. Available information includes general system settings, network, and proxy settings.

Procedure

To log on to the administrative console, complete the following steps:

1. From a supported browser, enter the following URL:

```
https://HOSTNAME:8090/
```

Where *HOSTNAME* is the IP address of the virtual machine where the application is deployed.

2. In the login window, select one of the following authentication types in the **Authentication Type** list:

| Authentication Type | Logon information |
|----------------------------------|--|
| IBM Spectrum Protect Plus | To log on as an IBM Spectrum Protect Plus user with SYSADMIN privileges, enter your administrator user name and password. |
| System | To log on as a system user, enter the serveradmin password. The default password is sppDP758. 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 11](#)

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

[“Managing roles” on page 249](#)

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. |

| Authentication Type | Login information |
|---------------------|--|
| 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. Restart the IBM Spectrum Protect Plus appliance from the Administrative Console.
6. Once the IBM Spectrum Protect Plus appliance has restarted, view the current time zone. Select **Product Information** from the main page of the Administrative Console and verify the updated time zone.

Uploading an SSL certificate from the administrative console

To establish secure connections in IBM Spectrum Protect Plus, you can upload an SSL certificate such as an HTTPS or LDAP certificate by using the administrative console.

About this task

For HTTPS certificates, PEM encoded certificates with `.cer` or `.crt` extensions are supported.

For LDAP/Hyper-V certificates, DER encoded certificates with `.cer` or `.crt` extensions are supported. If you are uploading an LDAP SSL certificate, ensure that IBM Spectrum Protect Plus has connectivity to the LDAP server and that the LDAP server is running.

ASCII and binary format certificates are accepted with the standard `.pem`, `.cer`, and `.crt` file extensions. However, the administrative console certificate import function cannot be used to update the appliance SSL web server communications. To upload ASCII and binary format certificates, use the command line as described in [“Uploading an SSL certificate from the command line” on page 230](#)

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 |

| Authentication Type | Logon information |
|---------------------|---|
| | is sppDP758. 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/ngxdd.jar
```

4. To connect to the tool, enter the following URL in a browser:

```
http://hostname:9000
```

where *hostname* specifies the IP address of the virtual machine where the application is deployed.

5. To specify the node to test, 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. Launch the tool from a command-line interface. Java is only required on the system where the tool will be launched. Endpoints or target systems that are tested by the tool do not require Java.

The following command launches the tool in a Linux environment:

```
# java -jar -Dserver.port=9000 /<tool path >/ngxdd.jar
```

3. To connect to the tool, enter the following URL in a browser:

```
http://hostname:9000
```

where *hostname* specifies the IP address of the virtual machine where the application is deployed.

4. To specify the node to test, populate the following fields:

Host

The host name or IP address of the node that you want to test.

Port

The connection port to test.

5. Click **Save**.

6. To run the tool, hover the cursor over the tool, and then click the green **Run** button.

If a connection cannot be established, the error condition is displayed, along with possible causes and actions.

7. Stop the tool by issuing the following command on the command line:

Adding virtual disks

You can add new virtual disks (hard disks) to your IBM Spectrum Protect Plus virtual appliance by using vCenter.

When you deploy the IBM Spectrum Protect Plus virtual appliance, you can deploy all virtual disks to one datastore that you specify at the time of deployment. You can add a disk within the virtual appliance and configure it as a Logical Volume Manager (LVM). You can then mount the new disk as a new volume or attach the new disk to the existing volumes within the virtual appliance.

You can review the disk partitions by using the **fdisk -l** command. You can review the physical volumes and the volume groups on the IBM Spectrum Protect Plus virtual appliance by using the **pvdisk** and **vgdisplay** commands.

Adding a disk to the virtual appliance

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

Before you begin

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

```
sudo bash
```

Press Enter.

```
echo "- - -" > /sys/class/scsi_host/host#/scan
```

Where # is the latest host number.

Adding storage capacity from a new disk to the appliance volume

After you add a disk to the virtual appliance, you can attach the new disk to the existing volumes within the virtual appliance.

Before you begin

To run commands, you must connect to the console of the IBM Spectrum Protect Plus virtual appliance by using SSH and log in with the user ID **serveradmin**. The default initial password is sppDP758 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 237](#)

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

Protection reports

IBM Spectrum Protect Plus provides reports that display the protection status of your resources. By viewing the reports and taking any necessary action, you can help to ensure that your data is protected through user-defined recovery point objective parameters.

To view protection reports, complete the following steps:

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

The following reports are available:

Protected and Unprotected VMs report

Run the Protected and Unprotected VMs report to view the protection status of your virtual machines. The report displays the total number of virtual machines added to the IBM Spectrum Protect Plus inventory before backup jobs are started.

Use the report options to filter by hypervisor type and to select specific hypervisors to display.

To exclude unprotected virtual machines in the report, select **Hide Unprotected VMs**.

To exclude virtual machines that are not backed up to secondary backup storage, select **Show only the VMs with Offloaded Backups**.

The **Summary View** displays an overview of your virtual machine protection status, including the number of unprotected and protected virtual machines and the managed capacity of the protected virtual machines. The managed capacity is the used capacity of a virtual machine. The **Detail View** provides further information about the protected and unprotected virtual machines, including names and location.

Protected and Unprotected Databases report

Run the Protected and Unprotected Databases report to view the protection status of your databases. The report displays the total number of databases added to the IBM Spectrum Protect Plus inventory before backup jobs are started.

Use the report options to filter by application type, application server, and application server type to display.


To exclude databases that are protected through hypervisor-based backup jobs, select **Hide Databases Protected as part of Hypervisor Backup**.

To exclude unprotected databases in the report, select **Hide Unprotected Databases**.

The **Summary View** displays an overview of your application server protection status, including the number of unprotected and protected databases, as well as the front end capacity of the protected databases. The front end capacity is the used capacity of a database. The **Detail View** provides further information about the protected and unprotected databases, including their names and location.


VM Backup History report

Run the VM Backup History report to review the protection history of specific virtual machines. To run the report, at least one virtual machine must be specified in the **VMs** option. You can select multiple virtual machine names.

Use the report options to filter by failed or successful jobs and time of the last backup. The report can be further filtered by specific service level agreement (SLA) policies. In the **Detail View**, click the plus icon  next to an associated job to view job details, such as the reason why a job failed or the size of a successful backup.

Database Backup History report

Run the Database Backup History report to review the protection history of specific databases. To run the report, at least one database must be specified in the **Databases** option. You can select multiple databases.

Use the report options to filter by failed or successful jobs and time of the last backup. The report can be further filtered by specific SLA policies. In **Detail View**, click the plus icon  next to an associated job to view further job details, such as the reason why a job failed or the size of a successful backup.

VM SLA Policy RPO Compliance report

The VM SLA Policy RPO Compliance report displays virtual machines in relation to recovery point objectives as defined in SLA policies. The report displays the following information:

- Virtual machines in compliance
- Virtual machines not in compliance
- Virtual machines in which the last backup job session failed

Use the report options to filter by hypervisor type and to select specific hypervisors to display. The report can be further filtered by virtual machines that are in compliance or not in compliance with the defined RPO.

Database SLA Policy RPO Compliance report

The Database SLA Policy RPO Compliance report displays databases in relation to recovery point objectives as defined in SLA policies. The report displays the following information:

- Databases in compliance
- Databases not in compliance
- Databases in which the last backup job session failed

Use the report options to filter by application type and to select specific application servers to display. The report can be further filtered by databases that are in compliance or not in compliance with the defined RPO, or by protection type, including data that was backed up to vSnap or by using replication.

Related concepts

[“Types of reports” on page 237](#)

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 237](#)

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 237](#)

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

Report actions

You can run, save, or schedule reports in IBM Spectrum Protect Plus.

Running a report

You can run IBM Spectrum Protect Plus reports with default parameters or run customized reports with custom parameters.

Before you begin

Custom roles that are assigned to users that run reports require that the appropriate permissions be set on that role so that the report can be viewed. For more information about roles, permission types, and permissions, see [“Managing roles” on page 249](#).

Procedure

To run a report, complete the following steps:

1. In the navigation pane, click **Reports and Logs > Reports**.
2. Expand a report type and select a report to run.
3. Run the report either with custom parameters or default parameters:
 - To run the report with custom parameters, set the parameters in the **Options** section, and click **Run**. Parameters are unique to each report.
 - To run the report with default parameters, click **Run**.

What to do next

Review the report in the **Reports** pane.

Related concepts

[“Managing reports and logs” on page 237](#)

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 237](#)

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. Edit the report parameters in the **Options** section,
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 237](#)


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 253](#)

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

Enabling the display of alerts in the system log

Alerts that are generated by IBM Spectrum Protect Plus can be logged to the system log. After you enable this feature, you can search the system log to find alerts.

About this task

An alert list in the IBM Spectrum Protect Plus user interface displays errors and warnings, which can help you to identify and resolve issues. Each item in the alert list includes a status icon, a summary of the issue, and a link to associated logs. These alerts can also be logged to the system log on the IBM Spectrum Protect Plus appliance.

Procedure

To enable logging of IBM Spectrum Protect Plus alerts to the system log, complete the following steps:

1. Ensure that there are no running jobs on the IBM Spectrum Protect Plus appliance.
2. Stop the Virgo service by issuing the following command:

```
sudo systemctl stop virgo
```

3. Log in to the IBM Spectrum Protect Plus appliance by using the Secure Shell (SSH) protocol.
4. Open the persistence properties file in the following directory:

```
/opt/virgo/repository/ecx-usr/com.syncsort.dp.xsb.persistence.properties
```

Tip: You can edit the persistence properties file by using the `vi` text editor, which is packaged along with the IBM Spectrum Protect Plus appliance.

5. Set the `alertInSystemLogMessages` variable to **yes**:

```
mongodb.alertInSystemLogMessages=yes
```

6. Save the properties file.

7. Restart the Virgo service by issuing the following command:

```
sudo systemctl start virgo
```

Results

Alerts are logged to the system log.

Example

```
[root@veguardian-cb3 log]# grep SPPAlert /var/log/messages
Jan 25 23:53:03 veguardian-cb3 virgo[993]: ERROR [SPPAlert]
vmware_Offload_Protect:1548460344227 Job vmware_Offload_Protect (id=1006,
session=1,548,460,344,227) failed.
```

Note: The following format is used for alerts that are logged by IBM Spectrum Protect Plus:

Message_type: The type can be ERROR, WARNING, and other types.

Message_source: Indicates the source of the message. This appears as SPPAlert for IBM Spectrum Protect Plus messages.

Message_text: Contains the IBM Spectrum Protect Plus alert details.

What to do next

Optionally, to search for an alert, log in to the IBM Spectrum Protect Plus appliance and search for SPPAlert.

Chapter 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 that are defined in a resource group, but not delete them. Similarly, permissions can be set to create administrator accounts, allowing a user to create and edit other accounts, set up sites and resources, and interact with all of the available IBM Spectrum Protect Plus features.

User accounts

A user account associates a resource group with a role. To enable a user to log in to IBM Spectrum Protect Plus and use its functions, you must first add the user as an individual user (referred to as a native user) or as part of an imported group of LDAP users, and then assign resource groups and roles to the user account. The account will have access to the resources and features that are defined in the resource group as well as the permissions to interact with the resources and features that are defined in the role.

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.

Before you begin


You may not assign more than one application per machine as an application server to a resource group. For example, if SQL and Exchange occupy the same machine and both are registered in SPP, only one of those can be added as an application server to a given resource group.

Procedure

To create a resource group, complete the following steps:

1. In the navigation pane, click **Accounts > Resource Group**.
2. Click **Create Resource Group**. The **Create Resource Group** pane displays.
3. Enter a name for the resource group.
4. From the **I would like to create a resource group** menu, select one of the following options:

| Option | Actions |
|----------------------|---|
| New | <ol style="list-style-type: none">a. Select a resource type from the Choose a resource type menu.b. Select resource subtypes, and then click Add Resources. Resources are added to the Selected Resources view. |
| From template | <ol style="list-style-type: none">a. Select a resource group from the Which resource group would you like to use as a template? list. Resources from the selected template are added to the Selected Resources view.b. You can add resources by using the Choose a resource type list and its associated lists. <p>To view available resource types and their usage, see “Resource types ” on page 247.</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 250 . |

Resource types

Resource types are selected when resource groups are created and determine the resources that are available to a user assigned to a group.

The following resource types and subtypes are available:

| Resource Type | Subtype | Description |
|--------------------|--|--|
| Accounts | <ul style="list-style-type: none"> • Role • User • Identity | Used to grant access to roles and users through the Accounts pane. |
| Application | <ul style="list-style-type: none"> • Db2 • Oracle • SQL Standalone/Failover Cluster • SQL Always On | Used to grant access to viewing individual application databases on an application server in IBM Spectrum Protect Plus. |
| Application Server | <ul style="list-style-type: none"> • Db2 • SQL • Oracle | Used to grant access to application servers in IBM Spectrum Protect Plus without access to individual databases. |
| Hypervisor | <ul style="list-style-type: none"> • VMware • Hyper-V | Used to grant access to hypervisor resources. |
| Job | None | Used to grant access to Inventory, Backup, and Restore jobs. The Job resource group is mandatory for all Backup and Restore operations, including assigning SLA Policies to resources. |
| Report | <ul style="list-style-type: none"> • Backup Storage Utilization • Protection • System • VE Environment | Used to grant access to report types and individual reports. |
| Screen | None | Used to grant or deny access to screens in the IBM Spectrum Protect Plus interface. If certain screens are not included in a resource group for a user, the user will not be able to access the functionality provided on the screen, regardless of the permissions granted to the user. |

| Resource Type | Subtype | Description |
|----------------------|---------------|--|
| SLA Policy | None | Used to grant access to SLA Policies for Backup operations. |
| System | Identity | Used to grant access to the credentials required to access your resources. Identity functionality is available through the System > Identity pane. |
| System Configuration | Disk | Used to grant access to vSnap backup storage servers. |
| System Configuration | LDAP | Used to grant access to LDAP servers for user registration. |
| System Configuration | Logs | Used to grant access to viewing and downloading Audit and System logs. |
| System Configuration | Script | Used to grant access to uploaded prescripts and postscripts. |
| System Configuration | Script Server | Used to grant access to script servers, where scripts are run during a Backup or Restore job. |
| System Configuration | Site | Used to grant access to sites, which are assigned to vSnap backup storage servers. |
| System Configuration | SMTP | Used to grant access to SMTP servers for job notifications. |
| System Configuration | VADP Proxy | Used to grant access to VADP proxy servers. |

Editing a resource group

You can edit a resource group to change the resources and features that are assigned to the group. Updated resource group settings take affect when user accounts that are associated with the resource group log in to IBM Spectrum Protect Plus. You may not assign more than one application per machine as an application server to a resource group.

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**.

You may not assign more than one application per machine as an application server to a resource group. For example, if SQL and Exchange occupy the same machine and both are registered in SPP, only one of those can be added as an application server to a given resource group.

Procedure

To edit a resource group, complete the following steps:

1. In the navigation pane, click **Accounts > Resource Group**.
2. Select a resource group and click the options icon ******* for the resource group. Click **Modify resources**.

3. Revise the resource group name, resources, or both.
4. Click **Update Resource Group**.

Deleting a resource group

You can delete any resource group that is not designated as **Cannot be modified**.

Procedure

To delete a resource group, complete the following steps:

1. In the navigation pane, click **Accounts > Resource Group**.
2. Select a resource group and click the options icon **☰** for the resource group. Click **Delete resource group**.
3. Click **Yes**.

Managing roles

Roles define the actions that can be completed for the resources that are defined in a resource group. While a resource group defines the resources that are available to an account, a role sets the permissions to interact with the resources.

For example, if a resource group is created that includes backup and restore jobs, the role determines how a user can interact with the jobs. Permissions can be set to allow a user to create, view, and run the backup and restore jobs that are defined in a resource group, but not delete them.

Similarly, permissions can be set to create administrator accounts, allowing a user to create and edit other accounts, set up sites and resources, and interact with all of the available IBM Spectrum Protect Plus features.

The functionality of a role is dependent on a properly configured resource group. When selecting a predefined role or configuring a custom role, you must ensure that access to necessary IBM Spectrum Protect Plus operations, screens, and resources align with the proposed usage of the role.

The following user account roles are available:

Application Admin

The Application Admin role allows users to complete the following actions:

- Register and modify application database resources that are delegated by an administrator.
- Associate application databases to assigned SLA policies.
- Complete backup and restore operations.
- Run and schedule reports to which the user has access.

Access to resources must be granted by an administrator through the **Accounts > Resource Groups** pane.

Backup Only

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

- Create, view, and run backup operations
- View, create, and edit SLA policies to which the user has access

Access to resources, including specific backup jobs, must be granted by an administrator by clicking **Accounts > Resource Groups**.

Restore Only

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

- Run, edit, and monitor restore operations.
- View, create, and edit SLA Policies to which the user has access.

Access to resources, including specific restore jobs, must be granted by an administrator through the **Accounts > Resource Groups** pane.

Self Service

The Self Service role allows users to monitor existing backup and restore operations that are delegated by an administrator.

Access to resources, including specific jobs, must be granted by an administrator through the **Accounts > Resource Groups** pane.

SYSADMIN

The SYSADMIN role is the administrator role. This role provides access to all resources and privileges.

Users with this role can add users and complete the following actions for all users other than the admin user:

- Modify and delete user accounts
- Change user passwords
- Assign user roles

An administrator can also access the administrative console by selecting **IBM Spectrum Protect Plus** from the **Authentication Type** list in the console login window and entering administrator credentials.

From the administrative console, the administrator can apply software updates, restart the IBM Spectrum Protect Plus appliance, and set the local time zone.

For more information about using the Administrative Console, see [“Logging on to the administrative console” on page 228](#).

VM Admin

The VM Admin role allows a users to complete the following actions:

- Register and modify hypervisor resources to which the user has access.
- Associate hypervisors to SLA policies.
- Complete backup and restore operations.
- Run and schedule reports to which the user has access.

Access to resources must be granted by an administrator through the **Accounts > Resource Groups** pane.

Creating a role

Create roles to define the actions that can be completed by the user of an account that is associated with a resource group. Roles are used to define permissions to interact with the resources that are defined in the resource group.

Procedure

To create a user role, complete the following steps:

1. In the navigation pane, click **Accounts > Role**.
2. Click **Create Role**. The **Create Role** pane displays.
3. From the **I would like to create a role** list, select one of the following options:

| Option | Actions |
|---------------|---|
| New | Select permissions to apply to the role. By default, none of the permissions are pre-selected. |
| From template | <ol style="list-style-type: none">a. Select a role from the Which role would you like to use as a template? menu. Permissions that are associated with the template role are selected by default.b. Select additional permissions to apply to the role, and delete permissions that are not required. |

| Option | Actions |
|--------|--|
| | To view available permissions and their usage, see “Permission types” on page 251. |

4. Enter a name for the role, and then click **Create Role**.

Results

The new role is displayed in the roles table and can be applied to new and existing user accounts.

Permission types

Permission types are selected when user accounts are created and determine the permissions that are available to the user.

The following permissions are available:

| Name | Permissions | Description |
|--------------------|---|---|
| Application | View | Used to view individual application databases on an application server in IBM Spectrum Protect Plus. |
| Application Server | Register, view, edit, deregister | Used to interact with application servers, such as SQL or Oracle servers, without access to individual databases. |
| Certificate | Create, view, edit, delete | Used to interact with SSL certificates to access Cloud Servers. |
| Cloud | Register, view, edit, deregister | Used to interact with Cloud Servers that are defined as backup storage for offloads. |
| Hypervisor | Register, view, edit, deregister, options | Used to interact with hypervisor virtual machines, such as VMware or Hyper-V virtual machines. |
| Identity and Keys | Create, view, edit, delete | Used to interact with the credentials required to access your resources. Identity functionality is available through the Accounts > Identities pane. |
| LDAP | Register, view, edit, deregister | Used to interact with LDAP servers for user registration. |
| Log | View | Used to view Audit and System logs. |
| Job | Create, view, edit, run, delete | Used to interact with Inventory, Backup, and Restore jobs. Note: If the user has permission to Run a job, then they also can Hold , Release , and Perform custom restore actions for the job |
| VADP Proxy | Register, view, edit, deregister | Used to interact with VADP |
| Report | Create, view, edit, delete | Used to interact with reports. |

| Name | Permissions | Description |
|----------------|----------------------------------|--|
| Resource Group | Create, view, edit, delete | Used to interact with resource groups, which define the IBM Spectrum Protect Plus resources that are made available to a user. |
| Role | Create, view, edit, delete | Used to interact with roles, which define the actions that can be performed on the resources defined in a resource group. |
| Script | Upload, view, replace, delete | Used to interact with prescripts and postscripts that are added to IBM Spectrum Protect Plus and run before or after a job. |
| Site | Create, view, edit, delete | Used to interact with sites, which are assigned to vSnap backup storage servers. |
| SMTP | Register, view, edit, deregister | Used to interact with SMTP servers for job notifications. |
| Backup Storage | Register, view, edit, deregister | Used to interact with vSnap backup storage servers. |
| SLA Policy | Create, view, edit, delete | Used to interact with SLA Policies, which allow users to create customized templates for Backup jobs. |
| User | Create, view, edit, delete | Used to interact with users, 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 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 246 and [“Creating a role”](#) on page 250.

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 225.

- If you want to use custom roles and resource groups, create them before you create a user. See [“Creating a resource group” on page 246](#) and [“Creating a role” on page 250](#).

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. If an identity is associated with a registered application server, it must be removed from the application server before it can be deleted. To remove the

association, navigate to the **Backup > Manage Application Servers** page associated with the application server type, then edit the settings of the application server.

Procedure

To delete an identity, complete the following steps:

1. In the navigation pane, click **Accounts > Identity**.
2. Click the delete icon  that is associated with an identity.
3. Click **Yes** to delete the identity.

Chapter 14. Licensing

License auditing in IBM Spectrum Protect Plus is enabled by default to determine if the current usage is within license entitlement levels and to prevent potential license violations.

IBM Spectrum Protect Plus generates entitlement audit logs as IBM® Software License Metric Tag (.slmtag) files. IBM® License Metric Tool (ILMT) is then used to translate the file and generate License Consumption Reports. Use the information in this section to interpret your .slmtag files.

Software License Metric (SLM) tags

IBM Spectrum Protect Plus generates entitlement audit logs as IBM® Software License Metric Tag (.slmtag) files. IBM® License Metric Tool (ILMT) is then used to translate the file and generate License Consumption Reports. Use the provided information to interpret your .slmtag files.

The .slmtag files can store information up to a maximum file size of 1 MB, after which the file is archived and a new log file is created. A maximum of 10 log files are kept.

Upgrade requirements: If you are upgrading to IBM Spectrum Protect Plus 10.1.3 from a prior release, you must run the maintenance job to generate the .slmtag files. For future upgrades, you must run the maintenance job to update existing .slmtag files.

Log format

The .slmtag files are stored in XML format, with new metric records appended to the end of the file.

The following is a sample .slmtag file:

```
<SchemaVersion>2.1.1</SchemaVersion>
<SoftwareIdentity>
  <Name>SPP</Name>
  <InstanceId>/opt/virgo</InstanceId>
</SoftwareIdentity>
<Metric logTime="2018-11-05T16:05:09+00:00">
  <Type>HYPERVISOR_SERVER_COUNT</Type>
  <SubType>HYPERVISOR_SERVER_COUNT</SubType>
  <Value>0</Value>
  <Period>
    <StartTime>2018-11-05T16:05:09+00:00</StartTime>
    <EndTime>2018-11-05T16:05:09+00:00</EndTime>
  </Period>
</Metric>
<Metric logTime="2018-11-05T16:05:09+00:00">
  <Type>APPLICATION_INSTANCE_COUNT</Type>
  <SubType>APPLICATION_INSTANCE_COUNT</SubType>
  <Value>0</Value>
  <Period>
    <StartTime>2018-11-05T16:05:09+00:00</StartTime>
    <EndTime>2018-11-05T16:05:09+00:00</EndTime>
  </Period>
</Metric>
```

where the Value element displays the number of hosts in all the resource groups with packages deployed for an instance group, at the specified time in the EndTime element.

The file grows over time and may be edited to remove older metric elements. Ensure that you retain elements long enough for ILMT scanning; the scanning frequency is determined by the ILMT administrator, but generally it should be sufficient to keep elements for a month.

Log location

The .slmtag file is located in the /data/slmtag directory.

Related concepts

[“Job types” on page 205](#)

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

Related tasks

[“Starting jobs” on page 206](#)

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

Integration with IBM License Metric Tool (ILMT)

Use IBM License Metric Tool (ILMT) to help determine whether your system environment is compliant with licensing requirements.

ILMT provides useful features for managing virtualized environments and measuring license utilization. ILMT discovers the software that is installed in your infrastructure, helps you to analyze the consumption data, and allows you to generate audit reports. Each report provides you with different information about your infrastructure, for example the computer groups, software installations, and the content of your software catalog.

By default, every ILMT audit report presents data from the previous 90 days. You can customize the type and amount of information displayed in a report by using filters, and save your personal settings for future use. You can also export the reports to .csv or .pdf format, and schedule report emails so that specified recipients are notified when important events occur.

For more information, see the [IBM License Metric Tool](#) product documentation.

Chapter 15. 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 the IBM Spectrum Protect product family

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

Overview

The IBM Spectrum Protect family of products includes the following major accessibility features:

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

The IBM Spectrum Protect family of products uses the latest W3C Standard, [WAI-ARIA 1.0](http://www.w3.org/TR/wai-aria/) (www.w3.org/TR/wai-aria/), to ensure compliance with [US Section 508](http://www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and [Web Content Accessibility Guidelines \(WCAG\) 2.0](http://www.w3.org/TR/WCAG20/) (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

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

Keyboard navigation

This product uses standard navigation keys.

Interface information

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

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

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

Vendor software

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

Related accessibility information

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

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

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

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

Trademarks

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

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

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

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

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

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

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

Terms and conditions for product documentation

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

Applicability

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

Personal use

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

Commercial use

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

Rights

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

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

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

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

Privacy policy considerations

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

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

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

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

Glossary

A glossary is available with terms and definitions for the IBM Spectrum Protect family of products.
See the [IBM Spectrum Protect glossary](#).

Index

A

- Access control
 - MongoDB [161](#)
- accessibility features [263](#)
- adding
 - Hyper-V servers [107](#)
 - identities [255](#)
 - LDAP server [225](#)
 - Oracle application servers [191](#)
 - sites [224](#)
 - SMTP server [226](#)
 - SQL Server application servers [182](#)
 - vCenter Server instances [83](#)
 - virtual disks to a vCenter virtual machine [233](#)
 - vSnap servers [50](#)
- Adding Db2 [122](#)
- Adding MongoDB [163](#)
- Administrative Console, logging on to [228](#)
- application server
 - Db2 [119](#)

B

- backup
 - jobs
 - on demand [207](#)
- backup jobs
 - ad hoc
 - on demand [207](#)
 - creating
 - Hyper-V [109](#)
 - IBM Spectrum Protect Plus [201](#)
 - Oracle [193](#)
 - SQL Server [184](#)
 - VMware [91](#)
 - excluding VMDKs from [94](#)
 - rerunning
 - on demand [207](#)
 - starting
 - on demand [206](#)
 - on schedule [79](#)
- backup policies, *See* SLA policies

C

- certificate
 - adding [222](#)
 - deleting [222](#)
- cloud provider
 - deleting [214](#)
 - editing [214](#)
- cloud server
 - adding a Microsoft azure cloud resource [213](#)
 - adding an amazon s3 cloud resource [211](#)
 - adding an IBM Cloud Object Storage resource [212](#)
- creating

- creating (*continued*)
 - reports [242](#)
 - resource groups [246](#)
 - roles [250](#)
 - SLA policies [79](#)
 - users
 - individual [253](#)
 - LDAP group [253](#)
 - VADP proxies [96](#)

D

- data protection [218](#)
- Db2
 - system requirements [29](#)
- deleting
 - identities [255](#)
 - LDAP server [227](#)
 - resource groups [249](#)
 - roles [253](#)
 - sites [224](#)
 - SLA policies [82](#)
 - SMTP server [227](#)
 - users [254](#)
- disability [263](#)

E

- early availability updates, obtaining and applying [77](#)
- editing
 - identities [255](#)
 - LDAP server [227](#)
 - resource groups [248](#)
 - roles [252](#)
 - settings [227](#)
 - sites [224](#)
 - SLA policies [81](#)
 - SMTP server [227](#)
 - users [254](#)
- Exchange Server
 - system requirements [27](#)

F

- fenced network, creating [104](#)
- files
 - restoring [116](#)
 - searching for [261](#)

H

- Hyper-V
 - adding [107](#)
 - backup job, creating [109](#)
 - installing on virtual appliance [43](#)
 - restore job, creating [112](#)

Hyper-V (continued)

- servers
 - detecting resources for [109](#)
 - enabling WinRM [108](#)
 - testing connection to [109](#)
- virtual appliance
 - accessing [231](#)

I

- IBM spectrum protect server
 - adding a repository server [219](#)
- identities
 - adding [255](#)
 - deleting [255](#)
 - editing [255](#)
- installing
 - download packages, obtaining [41](#)
 - virtual appliance
 - on Hyper-V [43](#)
 - on VMware [41](#)
 - vSnap servers
 - Hyper-V environment [49](#)
 - physical environment [47](#)
 - VMware environment [48](#)
- iSCSI utilities
 - installing [45](#)

J

- jobs
 - backing up single resource [207](#)
 - canceling [206](#)
 - names of [205](#)
 - pausing [206](#)
 - releasing [206](#)
 - rerunning [207](#)
 - starting
 - on demand [206](#)
 - on schedule [79](#)
 - types of [205](#)

K

- key
 - adding [221](#), [222](#)
 - deleting [222](#), [223](#)
- keyboard [263](#)
- keys [221](#)

L

- LDAP
 - group, creating a user account for [253](#)
 - server
 - adding [225](#)
 - deleting [227](#)
 - settings, editing [227](#)
- Linux-based vCenter virtual appliance, backing up [95](#)
- logs
 - audit
 - downloading [243](#)
 - viewing [243](#)

logs (continued)

- system
 - downloading [243](#), [259](#)
 - viewing [243](#), [259](#)

M

- MongoDB
 - system requirements [31](#)
- MongoDB application server [160](#)

N

- network
 - testing [231](#), [232](#)
- New in IBM Spectrum Protect Plus Version 10.1.2 [ix](#)

O

- Oracle
 - application servers
 - adding [191](#)
 - detecting resources for [192](#)
 - testing connection to [192](#)
 - backup job, creating [193](#)
 - restore job, creating [195](#)
 - system requirements [34](#)

P

- prerequisites
 - Db2 [119](#)
 - MongoDB [160](#)
- Prerequisites
 - MongoDB [161](#)

Q

- quick start [61](#)

R

- RBAC
 - MongoDB [161](#)
- reports
 - custom, creating [242](#)
 - running
 - on demand [241](#)
 - on schedule [242](#)
 - types of
 - backup storage utilization [237](#)
 - protection [238](#)
 - system [239](#)
 - VM environment [240](#)
- repository server provider
 - deleting [221](#)
 - editing [220](#)
- rerunning
 - jobs
 - on demand [207](#)
- resource groups
 - creating [246](#)

resource groups (*continued*)

- deleting [249](#)
- editing [248](#)
- types of [247](#)

restore jobs

- creating
 - Hyper-V [112](#)
 - IBM Spectrum Protect Plus [201](#)
 - Oracle [195](#)
 - SQL Server [187](#)
 - VMware [99](#)
- running
 - Hyper-V [112](#)
 - Oracle [195](#)
 - SQL Server [187](#)
 - VMware [99](#)

restore points, deleting [203](#)

restore points, managing [202](#)

roles

- creating [250](#)
- deleting [253](#)
- editing [252](#)
- permission types [251](#)

S

Schedule jobs

- Backup [127](#), [147](#), [167](#)

scripts for backup and restore operations

- uploading [208](#)

service level agreement, *See* SLA policies

sites

- adding [224](#)
- deleting [224](#)
- editing [224](#)

SLA [127](#), [147](#), [167](#)

SLA policies

- adding [79](#)
- deleting [82](#)
- editing [81](#)

SMTP

server

- adding [226](#)
- deleting [227](#)
- settings, editing [227](#)

SQL Server

- application servers
 - adding [182](#)
 - detecting resources for [183](#)
 - testing connection to [183](#)
- backup job, creating [184](#)
- requirements for data protection [182](#)
- restore job, creating [187](#)
- system requirements [37](#)

SSL certificate, uploading

- from administrative console [229](#)
- from command line [230](#)

starting

- IBM Spectrum Protect Plus [62](#)
- jobs

- on demand [206](#)
- on schedule [79](#)

system requirements

- components [11](#)

system requirements (*continued*)

- Db2 [29](#)
- Exchange Server [27](#)
- file index and restore [23](#)
- hypervisors [23](#)
- MongoDB [31](#)
- Oracle [34](#)
- SQL Server [37](#)
- system(
 - logs [243](#)

T

t_object_agent_client_sppIBM Spectrum Protect Plus [218](#)

time zone, setting [228](#)

U

user access [5](#), [245](#)

users

- deleting [254](#)
- editing [254](#)
- individual, creating [253](#)
- LDAP group, creating [253](#)
- resource groups
 - creating [246](#)
 - deleting [249](#)
 - editing [248](#)
 - types of [247](#)
- roles
 - creating [250](#)
 - deleting [253](#)
 - editing [252](#)
 - permission types [251](#)

V

VADP proxies

- creating [96](#)
- options, setting [98](#)
- uninstalling [99](#)
- updating [76](#)

virtual appliance

- accessing
 - in Hyper-V [231](#)
 - in VMware [230](#)
- adding a disk to [233](#)
- adding storage capacity [234](#)
- installing
 - on Hyper-V [43](#)
 - on VMware [41](#)

virtual environments [218](#)

VMware

- backup job, creating [91](#)
- backup job, excluding VMDKs from SLA policy [94](#)
- installing on virtual appliance [41](#)
- restore job
 - creating a fenced network [104](#)
- restore job, creating [99](#)
- vCenter Server instances
 - adding [83](#)
- vCenter Server, detecting resources [90](#)
- vCenter Server, testing connection to [90](#)

- VMware (*continued*)
 - virtual appliance
 - accessing [230](#)
 - virtual machine privileges, required [84](#)
- vSnap server
 - administering
 - network administration [59](#)
 - storage administration [57](#)
 - change throughput [55](#)
 - deleting [52](#)
 - editing [51](#)
 - initializing
 - advanced [53](#)
 - simple [53](#)
 - replication partnership, establishing [55](#)
 - storage options, managing [54](#)
 - storage pools, expanding [54](#)
- vSnap servers
 - adding [50](#)
 - installing
 - Hyper-V environment [49](#)
 - physical environment [47](#)
 - VMware environment [48](#)
 - uninstalling [60](#)

W

- WinRM, enabling for connection to Hyper-V servers [108](#)



Product Number: 5737-F11

Printed in USA