

IBM Spectrum Protect Plus
10.1.8

vSnap Installation and User's Guide



Note:

Before you use this information and the product it supports, read the information in [“Notices”](#) on page [87](#).

Second edition (30th April 2021)

This edition applies to version 10, release 1, modification 8 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, 2021.**

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

Contents

About this publication.....	V
Who should read this publication.....	v
Publications	v
What's new in Version 10.1.8.....	vii
Chapter 1. Product overview.....	1
Product components.....	1
Overview of the serveradmin user account.....	2
Replicate backup-storage data.....	3
Copy snapshots to secondary backup storage.....	3
Configuration for copying or archiving data.....	5
Chapter 2. Installing and managing vSnap servers.....	9
Installing a vSnap server.....	9
Installing a physical vSnap server.....	9
Installing a virtual vSnap server in a VMware environment.....	10
Installing a virtual vSnap server in a Hyper-V environment.....	11
Start IBM Spectrum Protect Plus.....	12
Updating IBM Spectrum Protect Plus components.....	13
Updating vSnap servers.....	14
Uninstalling a vSnap server.....	16
Chapter 3. Initializing the vSnap server.....	17
Completing a simple initialization.....	17
Completing an advanced initialization.....	17
Chapter 4. Managing vSnap servers.....	19
Registering a vSnap server.....	19
Expanding a vSnap storage pool.....	20
Establishing a replication partnership for vSnap servers.....	21
Migrating onboard vSnap data to a stand-alone vSnap server.....	21
Registering a vSnap server as a VADP proxy.....	25
Editing settings for a vSnap server.....	25
Refreshing a vSnap server.....	26
Removing the Demo environment.....	26
Unregistering a vSnap server.....	28
Chapter 5. Configuring backup storage options.....	31
Adding new disks to backup storage.....	32
Rescanning a vSnap server after the storage is expanded.....	32
Refreshing the disk storage for a vSnap server.....	33
Configuring backup storage partners.....	33
Configuring network interface controllers.....	33
Configuring an Active Directory.....	35
Setting advanced storage options.....	35
Changing the throughput rate.....	37
Chapter 6. Managing secondary backup storage.....	39

Managing cloud storage.....	39
Configuration for copying data to cloud storage.....	39
Managing repository server storage.....	44
Configuration for copying or archiving data.....	44
Chapter 7. vSnap server administration reference	57
Storage management.....	57
Kernel headers and tools.....	60
User management.....	61
Chapter 8. Resetting the serveradmin password.....	63
Chapter 9. Troubleshooting vSnap servers.....	65
Synchronizing the vSnap Password.....	65
How do I tier data to tape or cloud storage?	65
Why is the vSnap server still offline?.....	66
How does SAN work with IBM Spectrum Protect Plus and a vSnap server?	66
How do I repair a failed source vSnap in an IBM Spectrum Protect Plus environment?.....	67
How do I repair a failed target vSnap in an IBM Spectrum Protect Plus environment?	71
How do I repair a failed dual-role vSnap in an IBM Spectrum Protect Plus environment?.....	74
How do I delete and recreate a vSnap storage pool?.....	78
Chapter 10. Product messages.....	81
Message prefixes.....	81
Appendix A. Search guidelines.....	83
Appendix B. Accessibility.....	85
Notices.....	87
Glossary.....	91
Index.....	93

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 [Component requirements](#).

Publications

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

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

What's new in Version 10.1.8

IBM Spectrum Protect Plus Version 10.1.8 introduces new features and updates.

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

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

Chapter 1. IBM Spectrum Protect Plus overview

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

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

Product components

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

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

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

IBM Spectrum Protect Plus server

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

Site

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

Example deployment

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

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

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

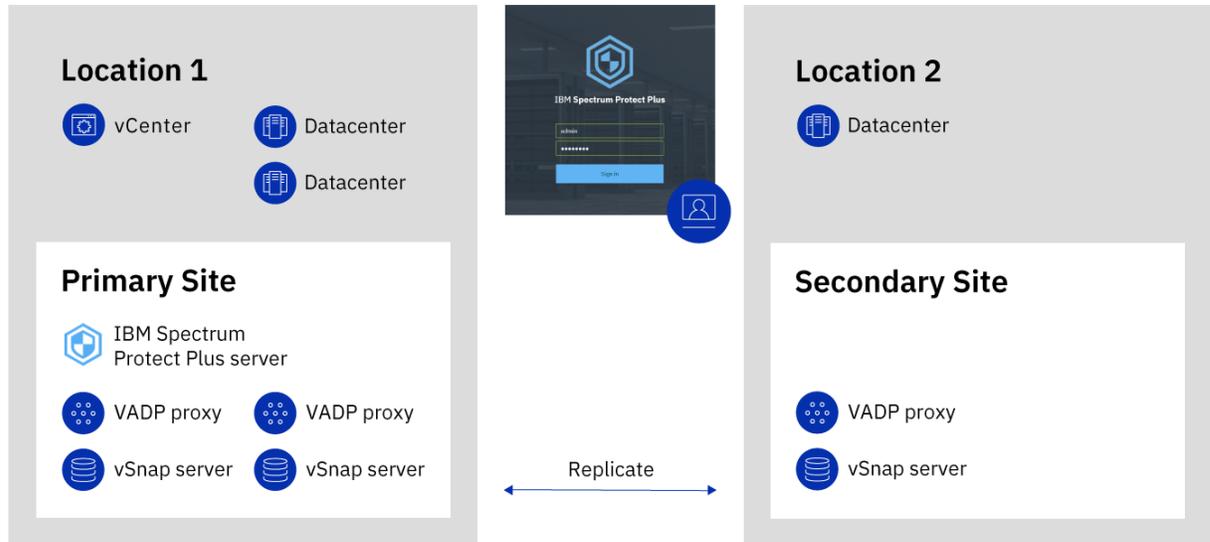


Figure 1. IBM Spectrum Protect Plus deployment across two geographical locations

Overview of the serveradmin user account

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

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

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

This may not be suitable for some environments. To harden the `serveradmin` user account password policy for IBM Spectrum Protect Plus and the vSnap server, the configuration for the account must be updated in the underlying CentOS operating system. To configure password aging, see [4.1.1.3. Configuring Password Aging](#). To configure account locking, see [4.1.2. Account Locking](#).

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

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

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

- In IBM Spectrum Protect Plus, edit the registration of the vSnap server to update the credentials to specify the new password.

- On IBM Spectrum Protect Plus and vSnap servers, if you enable account locking for the `serveradmin` account, you may be unable to log in to the appliance if there are too many failed attempts. Depending on how you configure the account locking, the access should unlock after a certain amount of time has passed.
- You may want to log in and reset the `serveradmin` account password without waiting for the configured time to pass. This can be done through using the `root` account. By default on IBM Spectrum Protect Plus and vSnap server OVAs, the `root` account can only be accessed through the virtual console and the password for the account is unknown. The `root` account password must first be reset in order to reset the `serveradmin` account password. For more information, see [Chapter 8, “Resetting the serveradmin password,”](#) on page 63.

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.

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

Considerations for enabling replication of backup-storage data

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

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

Copy snapshots to secondary backup storage

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

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

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

- IBM Cloud® Object Storage (including IBM Cloud Object Storage Systems)

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

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

Standard object storage

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

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

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

Tape or cloud archive storage

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

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

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

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

When you are copying data to tape from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, it is not a good idea to use the IBM Spectrum Protect tiering function. If you are archiving data to tape, you must use a cold cache storage pool. For more information about tiering, see [“How do I tier data to tape or cloud storage?”](#) on page 65. For different scenarios and more information about how to set up storage, see [“Configuration for copying or archiving data to IBM Spectrum Protect”](#) on page 5.

Example deployments

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

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

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

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

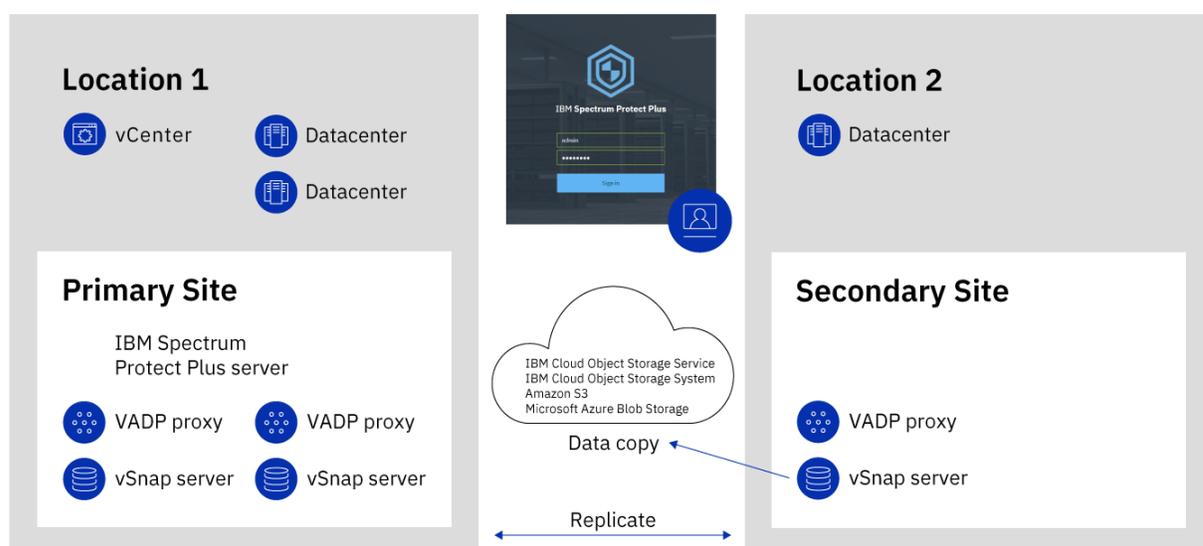


Figure 2. IBM Spectrum Protect Plus deployment across two geographical locations with copy to cloud storage

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

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

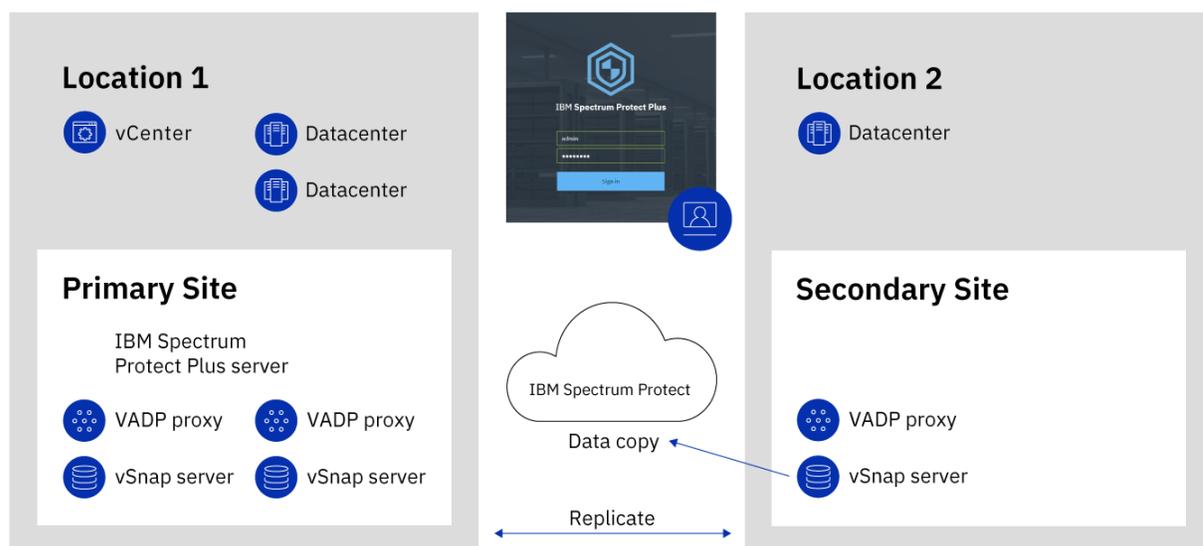


Figure 3. IBM Spectrum Protect Plus deployment across two geographical locations with copy to IBM Spectrum Protect

Configuration for copying or archiving data to IBM Spectrum Protect

If you are planning to copy or archive IBM Spectrum Protect Plus data to an IBM Spectrum Protect server, there are three possible configurations. Choosing which one to configure depends on which scenario applies to your data protection needs. For each scenario, there are steps that are required in both the IBM Spectrum Protect Plus and IBM Spectrum Protect server environments to complete the setup.

Tasks for configuring IBM Spectrum Protect

You must configure the IBM Spectrum Protect server to communicate with the IBM Spectrum Protect Plus server, and to enable process requests for backup and restore operations. The Amazon Simple Storage Service (S3) protocol enables communication between the two servers.

User scenario	Purpose	Steps
Copying to standard object storage when you are running daily or less frequent copies to standard object storage.	Copy data to standard object storage. In the first copy operation, a full backup copy is created. Subsequent copies are incremental. Copying data to standard object storage is useful if you want relatively fast backup and recovery times and do not require the longer-term protection, cost, and security benefits that are provided by tape storage.	To copy data to standard object storage to the IBM Spectrum Protect server, you must create a cloud-container or directory-container storage pool, and set up the object agent component of IBM Spectrum Protect. Adding the object agent is a mandatory step. In addition to setting up the required storage pool, follow steps 2-4 listed, here .
Copying to tape when you are creating a weekly or less frequent full-copy of your data to tape storage. Important: Archiving data to tape cannot be run more frequently than once a week. Recovery time objectives (RTO) should be considered when recovering data from archive copies in your disaster recovery action plan. Therefore, for disaster recovery, recovering from archive data should only be used as a last resort.	When you copy data to tape, a full copy of the data is created at the time of the copy process. Copying data to tape provides extra security benefits. By storing tape volumes at a secure, offsite location that is not connected to the internet, you can help to protect your data from online threats such as malware and hackers. However, because copying to these storage types requires a full data copy, the time that is required to copy data increases. In addition, the recovery time can be unpredictable and the data might take longer to process before it is usable.	To copy data to tape, you must create a cloud-container or directory-container storage pool for tape, and a cold-data-cache storage pool on the IBM Spectrum Protect server. Adding the object agent is a mandatory step. Follow steps 1-4 listed, here .
Mixture of both standard object storage and long-term copying to tape	Secure your data in incremental backups on the IBM Spectrum Protect server, as well as retaining data on tape for longer term security.	This is a combination of the previous cases: data is stored to tape and data is stored on standard object storage at the IBM Spectrum Protect server. As well as setting up the required data storage pools for both scenarios, the creation of an object agent is mandatory.

The four steps required to set up and configure the data transfer communication between IBM Spectrum Protect Plus and the IBM Spectrum Protect server are as follows:

1. If you are setting up storage pools for copying data to tape follow Step1. Create storage pools on the IBM Spectrum Protect server by using the IBM Spectrum Protect Operations Center. For instructions, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape” on page 46](#). This step is required only if you are setting IBM Spectrum Protect for archiving with copies run once a week or less frequently.

2. Create a policy domain that points to the storage pool or pools. The policy domain defines the rules that control the backup services for IBM Spectrum Protect Plus. For instructions, see [“Step 2: Configuring an object policy domain”](#) on page 48.
3. If you are copying data to a standard storage pool or to tape, you must add standard object storage on the IBM Spectrum Protect server. For instructions, see [“Step 3: Setting up standard object storage”](#) on page 49.
4. Add an object agent on the IBM Spectrum Protect server. The object agent provides a gateway between the IBM Spectrum Protect Plus server and the IBM Spectrum Protect server. For instructions, see [“Step 4: Adding an object agent for copying data”](#) on page 52.
5. To complete the setup, you must add an object client on the IBM Spectrum Protect server. The object client identifies the IBM Spectrum Protect Plus server and enables it to store objects at the IBM Spectrum Protect server. The same credentials as those that you used for IBM Spectrum Protect Plus are used for the object client, which is the object client that is associated with the policy domain as set up in Step 2. For instructions to set up an object client, see [“Step 5: Adding and configuring an object client for copying data”](#) on page 53.

Tip: Alternatively, enter the **DEFINE STGPOOL** command to create a storage pool as described in the following topics:

What to do next

1. After you complete the tasks required for IBM Spectrum Protect storage, you must add the IBM Spectrum Protect server to IBM Spectrum Protect Plus. For information about how to do this, follow the instructions in [“Registering a repository server as a backup storage provider”](#) on page 55.
2. When that is done, you can create an SLA policy that defines the IBM Spectrum Protect server as the backup storage target. For more information to help you choose which type of policy you need, see [Managing SLA policies for backup operations](#).

Chapter 2. Installing vSnap servers

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

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 virtual appliance is deployed. After deployment, some registration and configuration steps are required for these stand-alone vSnap servers.

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

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

Installing a vSnap server

You must have at least one vSnap server installed as part of your IBM Spectrum Protect Plus environment. This server is the primary backup destination. In larger enterprise environments, additional vSnap servers might be required. The Blueprints will help you determine how many vSnap servers are required.

Before you begin

Complete the following steps:

1. Review the vSnap system requirements. For more information, see [Component requirements](#).
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 6415111](#).

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

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

Installing a physical vSnap server

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

Procedure

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

See [Component requirements](#) for supported operating systems.

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

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

```
SELINUX=permissive
```

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

```
$ setenforce 0
```

4. Determine the vSnap installation file for the version of Linux on to which vSnap is being installed. The image name that contains *vsnap-dist-el7-<version-build>* is for CentOS and Red Hat Enterprise Linux 7.x. The image file that contains *vsnap-dist-el8-<version-build>* is for CentOS and Red Hat Enterprise Linux 8.x. Each image name has an associated **Passport Advantage Part Number**. Download the corresponding run file, *<part_number>.run*, for the image name from Passport Advantage Online. For information about downloading files, see [technote 6415111](#).

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

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

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

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

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

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

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

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

What to do next

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

Action	How to
Add the vSnap server to IBM Spectrum Protect Plus and configure the vSnap environment.	See Chapter 4, "Managing vSnap servers," on page 19 .

Installing a virtual vSnap server in a VMware environment

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

Before you begin

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

Procedure

1. Download the vSnap server template file *<part_number>.ova* from Passport Advantage Online. For information about downloading files, see [technote 6415111](#).
2. Deploy the vSnap server. Using the vSphere Client (HTML5) or the vSphere Web Client (FLEX), click the **Actions** menu and then click **Deploy OVF Template**.
3. Specify the location of the *<part_number>.ova* file and select it. Click **Next**.
4. Provide a meaningful name for the template, which becomes the name of your virtual machine. Identify an appropriate location to deploy the virtual machine. Click **Next**.
5. Select an appropriate destination to compute resource. Click **Next**.
6. Review the template details. Click **Next**.

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

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

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

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

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

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

What to do next

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

Action	How to
Add the vSnap server to IBM Spectrum Protect Plus and configure the vSnap environment.	See Chapter 4, “Managing vSnap servers,” on page 19.
For easier network administration, assign a static IP address for the virtual machine. Use the NetworkManager Text User Interface (nmtui) tool to assign the IP address.	For instructions, see Assigning a static IP address . Work with your network administrator when configuring network properties.

Installing a virtual vSnap server in a Hyper-V environment

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

Before you begin

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

Procedure

1. Download the vSnap installation file `<part_number>.exe` from Passport Advantage Online. For information about downloading files, see [technote 6415111](#).

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

What to do next

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

Action	How to
Add the vSnap server to IBM Spectrum Protect Plus and configure the vSnap environment.	See Chapter 4, “Managing vSnap servers,” on page 19 .

Start IBM Spectrum Protect Plus

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

Procedure

To start IBM Spectrum Protect Plus, complete the following steps:

1. In a supported web browser, enter the following URL:

```
https://hostname
```

The *hostname* value depends on whether IBM Spectrum Protect Plus installed as a set of OpenShift® containers or as a virtual appliance.

Hostname for a container installation

The hostname must be in the following format:

```
instancename-spp.routerCanonicalHostname
```

where *instancename* is the name of the IBM Spectrum Protect Plus instance and *routerCanonicalHostname* is the external host name for the OpenShift router.

Hostname for a virtual appliance installation

The hostname is the IP address of the virtual machine where the application is deployed.

2. Enter your username and password to log on to IBM Spectrum Protect Plus.

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

This user account is the superuser account and is assigned the `SUPERUSER` role. This role is assigned to only one IBM Spectrum Protect Plus user. The `SUPERUSER` role provides the user with access to all IBM Spectrum Protect Plus functions.

3. Click **Sign In**.

4. If IBM Spectrum Protect Plus is installed on a virtual appliance and you are logging in for the first time, you are prompted to change the `serveradmin` password. The initial password is `sppDP758-SysXyz`. The `serveradmin` user is used to access the administrative console and the IBM Spectrum Protect Plus virtual appliance. The password for `serveradmin` must be changed before accessing the administrative console and IBM Spectrum Protect Plus virtual appliance.

The following rules are enforced when creating a new password:

- The minimum acceptable password length is 15 characters.
- There must be eight characters in the new password that are not present in the previous password.
- The new password must contain at least one character from each of the classes (numbers, uppercase letters, lowercase letters, and other).
- The maximum number of identical consecutive characters that are allowed in the new password is three characters.
- The maximum number of identical consecutive class of characters that are allowed in the new password is four characters.

Updating IBM Spectrum Protect Plus components

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

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

Review the following restrictions and tips:

- The update process through the IBM Spectrum Protect Plus user interface updates IBM Spectrum Protect Plus features and the underlying infrastructure components including the operating system and file system. Do not use another method to update these components.
- Do not update any of the underlying components for IBM Spectrum Protect Plus unless the component is provided in an IBM Spectrum Protect Plus update package. Infrastructure updates are managed by IBM update facilities. The IBM Spectrum Protect Plus user interface is the primary means for updating

IBM Spectrum Protect Plus features and underlying infrastructure components including the operating system and file system.

Before you update components, it is important that you back up your IBM Spectrum Protect Plus environment as described in [Backing up the IBM Spectrum Protect Plus application](#).

Updating vSnap servers

vSnap servers, both virtually deployed or physically installed, must occasionally be updated.

Before you begin

You can update the IBM Spectrum Protect Plus and vSnap servers directly from two previous versions ($n-2$) to the current version (n). If you are using an older version, you must update at least to ($n-2$) version and then update to the current version.

Test restore jobs need to complete prior to initiating an update to vSnap. During a vSnap upgrade, a reboot will occur and any connected clients will experience a temporary disconnection. This disconnection may result in errors for any virtual machines or applications with active test mode restore. Additionally, jobs that are not completed or canceled when an update is initiated will not be visible once the update has completed. If jobs are not visible once the update has completed, re-run test restore jobs.

Review the system requirements before you update the vSnap servers.

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

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

```
$ vsnap system info
```

Ensure that no jobs that use the vSnap server are running during the update procedure. Pause the schedule for any jobs that do not have a status of IDLE or COMPLETED.

Updating the operating system for a physical vSnap server

If you have installed the vSnap server on a machine that is running Red Hat Enterprise Linux, you must update the operating system to version 7.5 or 7.6 before you update the vSnap server. For instructions about how to update the operating system, see the Red Hat Enterprise Linux documentation.

Related tasks

“Updating a vSnap server” on page 15

vSnap servers, both virtually deployed or physically installed, must occasionally be updated.

Updating the operating system for a virtual vSnap server

Updating the vSnap server operating system with the ISO file, provides you with the latest available patches and security updates. If the operating system is CentOS Linux version 7.4 or earlier, you must update the operating system before you update the vSnap server software. Updating the operating system is optional for version 7.5 or 7.6. An ISO file is downloaded and used to upgrade virtual vSnap servers.

Procedure

1. Download the ISO file `<part_number>.iso`. Move the ISO file to the `/tmp` directory on the vSnap server and rename the file to `spp_with_os.iso`.

```
$mv <part_number>.iso /tmp/spp_with_os.iso
```

Important: It is critical to rename the downloaded ISO file as described in this step and move it to the `/tmp` directory on the vSnap server if you wish to update the operating system.

2. Next, follow the instructions found in the [“Updating a vSnap server” on page 15](#) topic. When the `<part_number>.run` file is executed, the installer will optionally update the operating system if `/tmp/spp_with_os.iso` is present.

Note: Verify that the correct installation file for Linux is downloaded. For upgrades, the image file that contains `vsnap-dist-el7-<version-build>` should be used for both CentOS and Red Hat Enterprise Linux 7.x. Download the associated **Passport Advantage Part Number** file that is named `<part_number>.run`.

One of the two following scenarios will occur depending on the presence of the ISO file.

- If the file is present, operating system packages are upgraded, then vSnap software is upgraded.
- If the file is not present, a message is displayed:

```
File /tmp/spp_with_os.iso is not present, skipping update of OS packages.  
To update OS packages, download the ISO file to /tmp/spp_with_os.iso and rerun this  
installer.
```

Then vSnap software is then is upgraded.

Once the installer completes, `/tmp/spp_with_os.iso` can be deleted.

Related tasks

[“Updating a vSnap server” on page 15](#)

vSnap servers, both virtually deployed or physically installed, must occasionally be updated.

Updating a vSnap server

vSnap servers, both virtually deployed or physically installed, must occasionally be updated.

Procedure

To update a vSnap server, complete the following steps:

1. Log on to the vSnap server as the `serveradmin` user.
2. From the directory where the `<part_number>.run` file is located, make the file executable by issuing the following command:

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

3. Run the installer by issuing the following command:

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

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

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

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

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

The vSnap packages are installed.

4. After the vSnap packages are installed, start the updated version of the vSnap server.
5. In the navigation pane, click **Jobs and Operations**, and then click the **Schedule** tab.
Find the jobs that you paused.
6. From the **Actions** menu for the paused jobs, select **Release Schedule**.

Uninstalling a vSnap server

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

Before you begin

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

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

To view the SLA policies that are associated with jobs, see the **Backup** page for the hypervisor or application that is scheduled for backup. For example, for VMware backup jobs, click **Manage Protection > Hypervisors > VMware**. You must unregister the vSnap server from the IBM Spectrum Protect Plus server. See “Unregistering a vSnap server” on page 28 for more information.



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

Procedure

1. Log on to the vSnap server console with the user ID `serveradmin`. The initial password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in “[Start IBM Spectrum Protect Plus](#)” on page 12.

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 [Chapter 7, “vSnap server administration reference,”](#) on page 57.

2. Run the following commands:

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

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

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

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

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

Results

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

Chapter 3. Initializing the vSnap server

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

About this task

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

For information about how to expand, create, and administer storage pools, see [“Storage management” on page 57](#).

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

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

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

Completing a simple initialization

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

Procedure

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

1. In the navigation pane, click **System Configuration > Backup Storage > Disk**.
2. From the actions menu icon  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 your environment. Initializing by using the vSnap server console offers more options for initializing the server, including the ability to create a storage pool by using advanced redundancy options and a specific list of disks.

Procedure

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

1. Log in to the vSnap server console with the user ID `serveradmin` by using SSH. When deployed virtually, the initial password is `sppDP758-SysXyz`. You will be prompted to change this password

during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus”](#) on page 12. If deployed physically, use the password that you created for the `serveradmin` account during installation.

You can also use a user ID that has vSnap privileges that was previously created using the **vsnap user create** command. For more information about using console commands, see [Chapter 7, “vSnap server administration reference,”](#) on page 57.

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

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

What to do next

After you complete the initialization, complete the following action:

Action	How to
Create a storage pool	See “Storage management” on page 57.

Chapter 4. Managing vSnap servers

To enable backup and restore jobs, IBM Spectrum Protect Plus requires at least one vSnap server. The vSnap server is its own appliance, either deployed virtually or installed physically on a system that meets the minimum requirements. Each vSnap server in the environment must be registered in IBM Spectrum Protect Plus so that it is recognized.

Registering a vSnap server as a backup storage provider

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

Before you begin

After you add and register a vSnap server as a backup storage provider, you can choose to configure and administer certain aspects of the vSnap, such as network configuration or storage pool management. For more information, see [Chapter 5, “Configuring backup storage options,” on page 31](#).

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

Procedure

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

1. Log on to the vSnap server console with the user ID `serveradmin`. The initial password is `sppDP758-SysXyz`.
You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus” on page 12](#).
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.

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

Site

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

Username

Enter the user name for the vSnap server that you created in step “2” on [page 19](#).

Password

Enter the password for the user.

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

7. 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 Chapter 3, “Initializing the vSnap server,” on page 17.
Expand the vSnap storage pool.	See “Configuring backup storage partners” on page 33.
If necessary, configure and administer certain aspects of vSnap, such as network configuration or storage pool management.	See Chapter 5, “Configuring backup storage options,” on page 31

Related tasks

[“Start IBM Spectrum Protect Plus” on page 12](#)

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

Expanding a vSnap storage pool

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

Before you begin

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

Note:

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

Procedure

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

1. In the navigation pane, click **System Configuration > Backup Storage > Disk**.
2. Click the **Manage** icon that is associated with the vSnap server .
3. Click **Add New Disks to Backup Storage**, and click **Rescan**.

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

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

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

What to do next

After you expand the storage pool, rescan the disk or vSnap server to pick up the new disks. For instructions on how to run a rescan operation, see [“Rescanning a vSnap server after the storage is expanded” on page 32.](#)

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**.

Migrating onboard vSnap data to a stand-alone vSnap server

Beginning with IBM Spectrum Protect Plus Version 10.1.7, the onboard vSnap server is no longer included. If you upgrade your system to IBM Spectrum Protect Plus V10.1.7 or later, but data remains in an onboard vSnap server from a previous release prior to V10.1.7, you must migrate the data to a new, stand-alone vSnap server.

Before you begin

Beginning with IBM Spectrum Protect Plus V10.1.7, new deployments will no longer contain an onboard vSnap server. Systems upgraded from a previous version of IBM Spectrum Protect Plus still contain an onboard vSnap server which can be part of the Demo site. The onboard vSnap server will no longer be upgraded as part of general updates to IBM Spectrum Protect Plus.

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

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

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

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



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

About this task

In previous releases, an onboard vSnap server was included for proof-of-concept (POC) and demo purposes. The vSnap server was named localhost and was part of the Primary site by default. Beginning with IBM Spectrum Protect Plus V10.1.5, the onboard vSnap server was part of a Demo site that provided limited functionality. Users were able to manually remove the onboard vSnap from the Demo site and then register it with another site at which point the vSnap server was no longer limited in functionality.

Determine whether an onboard vSnap server was used in the previous release. Users who did not unregister the onboard vSnap from the Demo site will follow a different procedure from users who

unregistered the onboard vSnap server from the Demo site and assigned the server to another site. Consider the two scenarios below:

Scenario 1: If the onboard vSnap was unused or previously used only in the Demo site, stop using the onboard vSnap. Unregister the vSnap from IBM Spectrum Protect Plus, for more information, see [“Unregistering a vSnap server” on page 28](#). After completing those steps, uninstall the vSnap software from the IBM Spectrum Protect Plus server. Skip the steps and begin with Step 9 in this procedure.

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

- Manually upgrade the onboard vSnap server using the appropriate `.run` file. Follow to the general procedure for upgrading an external vSnap server as described in the [“Updating a vSnap server” on page 15](#) topic.
- Deploy a new stand-alone vSnap server using the most recent OVA. For more information, see [“Installing a virtual vSnap server in a VMware environment” on page 10](#).
- Upon completing the migration of data, uninstall the vSnap software from the IBM Spectrum Protect Plus server. These steps are detailed in this procedure beginning with Step 9.

Procedure

1. Update the onboard vSnap server and collect the vSnap pool information.
 - a) Using secure shell (SSH), log in to the onboard vSnap as the **serveradmin** user.
 - b) Upgrade the vSnap server to the most recent release. For more information, see [“Updating a vSnap server” on page 15](#).
 - c) Determine the version level of the vSnap server. At the command prompt enter the **vsnap system info** command:

```
$ vsnap system info
```
 - d) Determine all the disks labeled `vsnap_pool1`. The storage pool is comprised of these disks which will be detached from the onboard vSnap server and attached to the new vSnap server later in this procedure. At the command prompt, issue the **vsnap disk show** command to identify the disks:

```
$ vsnap disk show
```
2. Deploy a new, stand-alone vSnap server using the most recent `.ova`, apply custom settings, and verify the version level.
 - a) Log in to the vSphere Client.
 - b) Deploy a new stand-alone vSnap server using the most recent version of the vSnap `.ova`. For more information, see [“Installing a virtual vSnap server in a VMware environment” on page 10](#).
 - c) The new, stand-alone vSnap server will contain an unused 100GB disk that is used as the initial disk for creating a new storage pool. Detach this disk from the stand-alone vSnap server and delete it.
 - d) Configure the network properties as appropriate to your environment on the newly created vSnap server. Document the IP address or hostname for later use in this procedure.
 - e) Using secure shell (SSH), log in to the newly created vSnap as the **serveradmin** user.
 - f) Determine the version level of the newly created vSnap server. At the command prompt enter the **vsnap system info** command:

```
$ vsnap system info
```

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

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


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


```
$ scp vsnap_config_backup.tar.gz serveradmin@<ip_address_new_vsnap>:/home/serveradmin
```
 - d) Enter the password for the **serveradmin** account on the stand-alone vSnap server. The file will begin transferring.
 - e) Disable the vSnap services for the onboard vSnap server using the **systemctl stop** and **systemctl disable** commands:


```
$ sudo systemctl stop vsnap
```

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


```
$ vsnap system config restore --file /home/serveradmin/vsnap_config_backup.tar.gz
```
6. Power off the onboard vSnap server and the stand-alone vSnap server, detach the disks from the onboard vSnap and attach the disks to the stand-alone vSnap server. Power on both vSnap servers.
 - a) Log in to the vSphere Client.
 - b) Power off the onboard vSnap and the stand-alone vSnap virtual machines and edit the settings of the virtual machine that has the onboard vSnap.
 - c) Detach the disks associated with the vSnap pool that is to be migrated as identified in Step 1d.
 - d) Edit the settings of the stand-alone vSnap virtual machine and attach the disks that were detached from the onboard vSnap server in Step 6c.

- e) Power on the onboard vSnap and the stand-alone vSnap virtual machines.
- 7. Verify the status of both the onboard vSnap server and the newly deployed stand-alone vSnap server.
 - a) Using secure shell (SSH), log in to the onboard vSnap server as the **serveradmin** user.
 - b) Run the **vsnap_status** command to determine the status of the vSnap services on the onboard vSnap server. It is expected that the services will no longer be running since the **systemctl stop** and **systemctl disable** commands were previously executed in Step 4.

```
$ vsnap_status
```

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

```
$ vsnap_status
```

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

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

```
$ vsnap pool show
```

- 8. Update the vSnap server registration, the associated credentials, re-add the replication partners, and release the job schedules.
 - a) Log on to the IBM Spectrum Protect Plus server.
 - b) Click on **System Configuration > Backup Storage > Disk** and click on the edit icon  beside the onboard vSnap server.
 - c) Enter the IP address or the hostname in the **Hostname/IP** field of the newly created stand-alone vSnap server.
 - d) The existing user may display as **LocalvSnapAdmin** or as another identity. Deselect **Use existing user**. Enter **serveradmin** in the **User ID** field and the associated password for the stand-alone vSnap server in the **Password** field.
 - e) Click **Save**.
 - f) On the **Disk** screen locate the vSnap server that was just edited and click on the actions menu icon . Select **Refresh**.
 - g) After the refresh completes verify that the information for the vSnap server is accurate.
 - h) Click on settings icon  and then click on the **Partners** tab.
 - i) Re-enter the replication partners that were removed in Step 3. For more information, see [“Configuring backup storage partners”](#) on page 33.
 - j) Release schedules for all jobs that were paused in Step 3. Navigate to **Jobs and Operations > Schedule** and then click **Release All Schedules**.
- 9. Remove the vSnap software from the IBM Spectrum Protect Plus server.
 - a) Using secure shell (SSH), log in to the IBM Spectrum Protect Plus server as the **serveradmin** user.
 - b) Execute the **yum remove** commands to remove the vSnap server software from the IBM Spectrum Protect Plus server:

```
$ sudo yum remove vsnap
```

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

Results

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

Registering a VADP proxy on a vSnap server

You can install and register a VADP proxy on a physical or virtual vSnap server. When you install and register a VADP proxy locally on a vSnap server, no NFS mount is needed. Data movement is optimized because the file system is on the same machine and can be referenced directly for both backup and restore jobs. VADP proxies that are not installed and registered on a vSnap server still require an NFS mount.

Before you begin

One or more stand-alone vSnap servers must be properly deployed and configured in your environment and added to IBM Spectrum Protect Plus backup storage providers. For instructions, see [“Registering a vSnap server as a backup storage provider”](#) on page 19.

For the combined system requirements of a vSnap server and the VADP proxy, see [VADP proxy on vSnap server requirements](#).

Ensure that you have the required user permissions to work with VADP proxies. For instructions about managing VADP proxy permissions, see [Permission types](#).

The identity associated with a vSnap server is the account that is used to register the VADP proxy on the vSnap server. When you register a VADP proxy on a vSnap server, an installer is pushed and requires sudo privileges to successfully install the VADP proxy software. The identity associated with a vSnap server must have sudo privileges.

Tip: Use the `serveradmin` User ID when adding a vSnap server to IBM Spectrum Protect Plus. When you deploy a VADP proxy to a vSnap server, this account is used which already has all of the necessary privileges.

Procedure

1. In the navigation pane, on **System Configuration > Backup Storage > Disk**. Available vSnap servers are displayed in the table in the Disk Storage pane.
2. Select the vSnap server on which the VADP proxy is to be installed and registered.
3. Click the actions menu icon . Select **Register as VADP Proxy**.
4. In the Confirm dialog box, click **Yes**.

Results

When the process is complete, a green checkmark will appear in the **VADP Proxy** column in the table of the Disk Storage pane.

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**.

Refreshing a vSnap server

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

Procedure

1. In the navigation pane, click **System Configuration > Backup Storage > Disk**.
2. Click the actions icon  for the disk that you want to refresh.
3. Click **Refresh** to refresh the details of the disk. For example, if the **Status/Capacity** percentage has changed due to usage, the update is refreshed in the table.

Removing the Demo environment

The IBM Spectrum Protect Plus appliance includes an onboard vSnap server that is named localhost, a site for demonstration purposes that is named Demo, and an associated SLA policy that is named Demo. For larger production environments, do not use the onboard vSnap server. Instead use, one or more stand-alone vSnap servers. The Demo SLA policy, Demo site, and onboard vSnap server, collectively as the Demo environment, can be safely removed to conserve disk space.

Before you begin

For IBM Spectrum Protect Plus appliances that are in production, back up the IBM Spectrum Protect Plus application. For instructions, see [Backing up the IBM Spectrum Protect Plus application](#). For new deployments, backing up the application is not necessary.

Verify that the data on the localhost vSnap server is not needed.

Ensure that at least one stand-alone vSnap server is deployed as a backup destination.

About this task

When deployed, an IBM Spectrum Protect Plus appliance has six virtual hard disks. When you remove the Demo configuration and localhost vSnap server from the IBM Spectrum Protect Plus appliance, you can free storage through the removal of two of the associated virtual hark disks.

The procedure in this topic must be followed in order to remove the Demo environment from IBM Spectrum Protect Plus.

Procedure

1. Disable SLA policies that are assigned to the Demo environment by completing the following steps:
 - a) From a supported browser, log in to the IBM Spectrum Protect Plus user interface.
 - b) View any jobs that are assigned to the Demo SLA. In the navigation pane, click **Jobs and Operations**, and then click the **Schedule** tab. Locate any jobs that follow the naming pattern *Job_Name_Demo*, where *Job_Name* is the name of the job. This naming pattern indicates that the Demo SLA is used.
 - c) Pause the schedule for every Demo job. Click the actions menu icon  and select **Pause Schedule** for each job that ends in *_Demo*.
2. Delete the Demo SLA by completing the following steps:
 - a) In the navigation pane, click **Manage Protection > Policy Overview**. Scroll down to the table in the SLA Policies pane and locate the Demo policy.
 - b) Click the delete icon  beside the Demo SLA.
 - c) Enter the code in the **Confirm** dialog box and click **OK**.

3. Delete the localhost vSnap disk storage by completing the following steps:
 - a) In the navigation pane, click **System Configuration > Backup Storage > Disk**. Locate the localhost vSnap storage that is assigned to the Demo site.
 - b) Click the delete icon  beside the localhost vSnap storage.
 - c) Enter the code in the **Confirm** dialog box and click **DELETE**.
4. Delete the Demo site by completing the following steps:
 - a) In the navigation pane, click **System Configuration > Site**. Locate the site that is named Demo.
 - b) Click the delete icon  beside the Demo site.
 - c) Click **Yes** in the **Confirm** dialog box to complete the removal of the Demo site.
5. Remove the LocalvSnapAdmin identity by completing the following steps:
 - a) In the navigation panel, click on **Accounts > Identity**.
 - b) Click the delete icon  beside the LocalvSnapAdmin identity.
 - c) Click on **Yes** in the **Confirm** dialog box to remove the identity.
6. Clean up the file system and LVM configurations by completing the following steps:
 - a) Log in to the IBM Spectrum Protect Plus by using the Secure Shell (SSH) protocol or through the hypervisor console by using the `serveradmin` account.
 - b) Obtain the ID of the localhost vSnap storage pool. Issue the following command:


```
$ vsnap pool show
```

 **Attention:** To ensure that no data is lost, verify that the ID obtained is the ID of the localhost vSnap storage pool.
 - c) Delete the localhost vSnap storage pool. Issue the following command where `<ID>` is the ID obtained in the previous step:


```
$ vsnap pool delete --id <ID>
```
 - d) Unmount the localhost vSnap storage cloud cache. Issue the following command:


```
$ sudo umount -f /opt/vsnap-data
```
 - e) Edit the `fstab` file to disable the cloud cache from starting. Using `sudo` and a text editor, comment out the line starting with `/dev/mapper/vsnapdata-vsnapdata1v`.
 - f) Deactivate the LVM volume group that is associated with the cloud cache. Issue the following command:


```
$ sudo vgchange -an vsnapdata
```
7. By using vSphere or Hyper-V Manager, detach the virtual hard disks that are no longer needed from the IBM Spectrum Protect Plus appliance. Proceed with caution to ensure that the correct disks are detached. The localhost vSnap server has two associated virtual hard disks, which are 100 GB and 128 GB in size. For detailed instructions about detaching or removing virtual hard disks, see the appropriate hypervisor documentation.

 **Attention:** Power off the IBM Spectrum Protect Plus appliance before you detach the virtual hard disks. Do not delete the virtual hard disks until proper functionality has been confirmed after powering on the appliance and running a maintenance job.
8. Rescan the SCSI bus and disable the vSnap service by completing the following steps:
 - a) Log in to the IBM Spectrum Protect Plus by using the Secure Shell (SSH) protocol or through the hypervisor console by using the `serveradmin` account.
 - b) Rescan the SCSI bus by issuing the following command:

```
$ sudo rescan-scsi-bus.sh
```

c) Stop the vSnap service by issuing the following command:

```
$ sudo systemctl stop vsnap
```

d) Disable the vSnap service by issuing the following command:

```
$ sudo systemctl disable vsnap
```

Unregistering a vSnap server

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

Before you begin

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



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

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

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

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

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

- Do not unregister the vSnap server. If you unregister the vSnap server, recovery points that are associated with the server will be purged and backups will be rebased.
- Edit the settings for the vSnap server to specify the new host name or IP address. To edit the settings for a vSnap server, follow the instructions [“Editing settings for a vSnap server”](#) on page 25.

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

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

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

1. Unregister the vSnap server and run a maintenance job to ensure that recovery points that are associated with the old vSnap pool are purged from IBM Spectrum Protect Plus.
 - Incremental backups of the data that was present in the old pool will no longer be possible.
 - Recovering data that was present in the old pool will no longer be possible.
2. On the vSnap server, create a pool.
3. Add the vSnap server back into IBM Spectrum Protect Plus. To add a vSnap server to IBM Spectrum Protect Plus, see [“Registering a vSnap server as a backup storage provider”](#) on page 19.
 - Subsequent runs of backup jobs will automatically create volumes on this or another vSnap server in the same site and will perform new base backups.

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

- Do not unregister the vSnap server from IBM Spectrum Protect Plus. The deletion process will cause backups to be rebased.
- Replace the vSnap server. For information about replacing a failed, primary vSnap server, see this section [Chapter 9, “Troubleshooting vSnap servers,”](#) on page 65.

Procedure

To unregister 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. Confirm removal of the vSnap server by entering the code in the text box. Click **DELETE** to delete the server from IBM Spectrum Protect Plus.

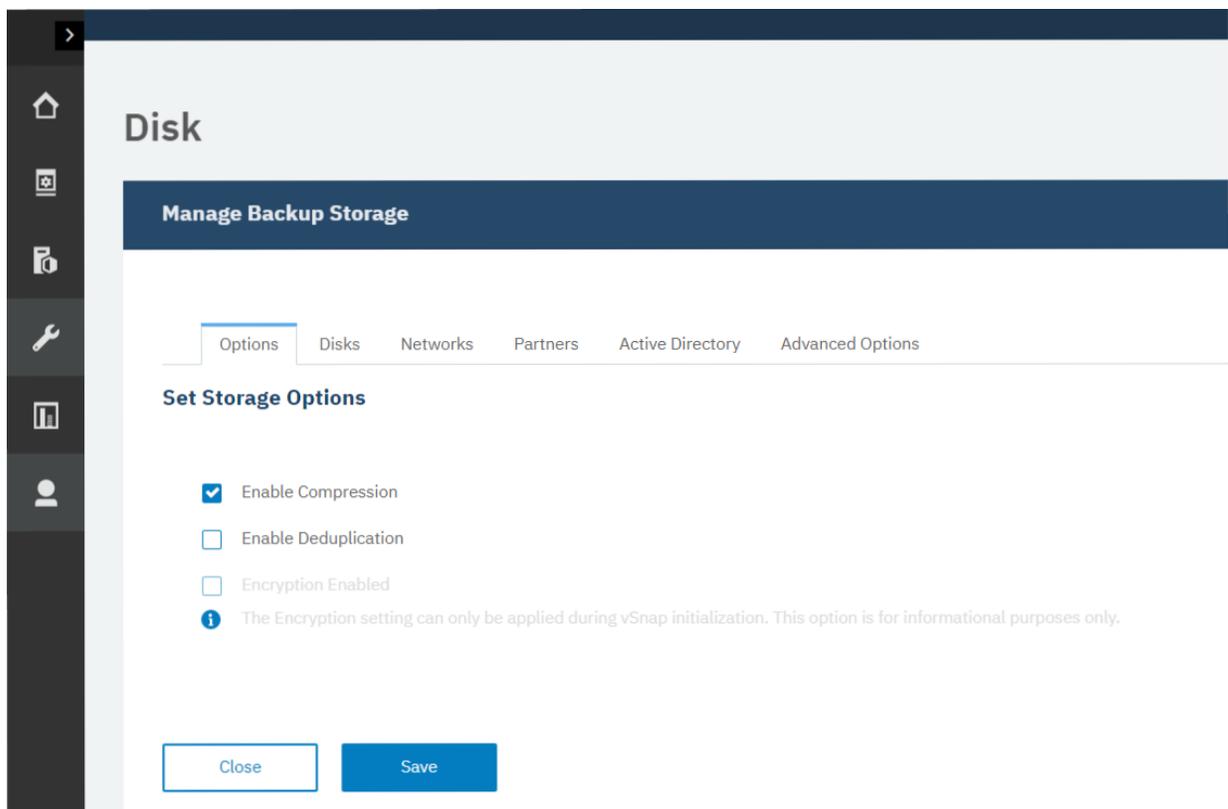
Chapter 5. Configuring backup storage options

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

Procedure

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

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



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

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

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

4. Click **Save**.

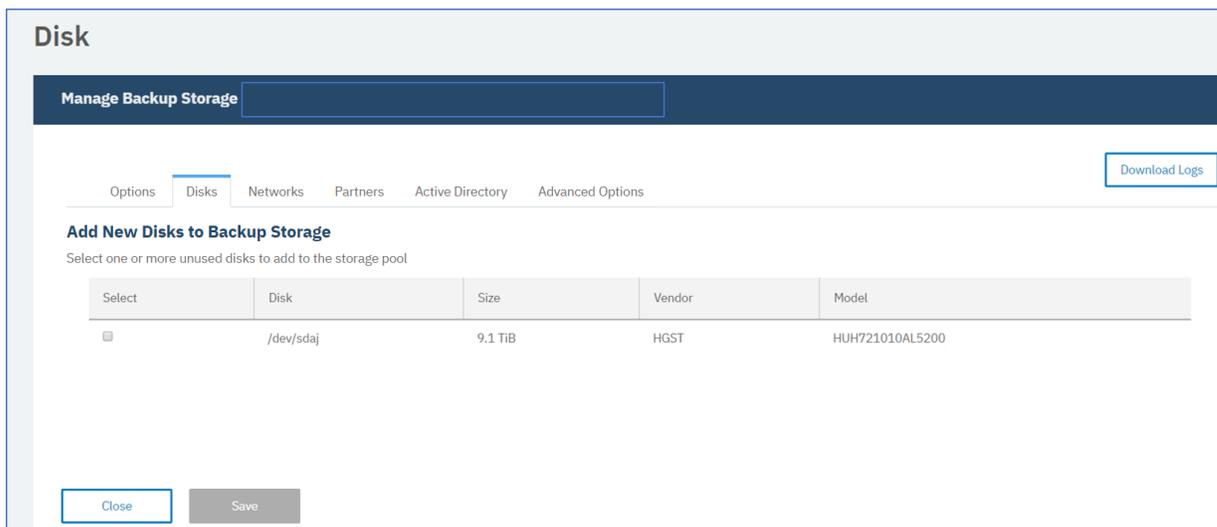
Adding new disks to backup storage

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

Procedure

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

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



The screenshot shows the 'Disk' management interface. At the top, there is a 'Manage Backup Storage' header. Below it, there are tabs for 'Options', 'Disks', 'Networks', 'Partners', 'Active Directory', and 'Advanced Options'. A 'Download Logs' button is located in the top right corner. The main section is titled 'Add New Disks to Backup Storage' and includes the instruction 'Select one or more unused disks to add to the storage pool'. Below this is a table with the following columns: 'Select', 'Disk', 'Size', 'Vendor', and 'Model'. The table contains one row with a checkbox in the 'Select' column, the path '/dev/sdaj' in the 'Disk' column, '9.1 TiB' in the 'Size' column, 'HGST' in the 'Vendor' column, and 'HUH721010AL5200' in the 'Model' column. At the bottom of the interface, there are 'Close' and 'Save' buttons.

Select	Disk	Size	Vendor	Model
<input type="checkbox"/>	/dev/sdaj	9.1 TiB	HGST	HUH721010AL5200

4. Click **Save**.

Rescanning a vSnap server after the storage is expanded

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

Procedure

1. In the navigation pane, click **System Configuration** > **Backup Storage** > **Disk**.
2. Click the settings icon  for the vSnap server that you want to rescan.
3. Open the **Disks** tab.
Any added disks are listed in the table.
4. Click **Rescan** to scan the vSnap server for any storage pool expansion or changes.
This operation can several minutes to finish. The disk remains fully operational during the scanning process.
5. Optional: Select **Refresh** to refresh the details of the disk. For example, if the **Status/Capacity** figure has changed due to usage, the update is refreshed in the table.

Refreshing the disk storage for a vSnap server

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

Procedure

1. In the navigation pane, click **System Configuration > Backup Storage > Disk**.
2. Click the actions icon  for the disk that you want to refresh.
3. Click **Refresh** to refresh the details of the disk.
For example, if the **Status/Capacity** percentage has changed due to usage or because you recently expanded the storage pool, the information is refreshed in the table.

Configuring backup storage partners

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

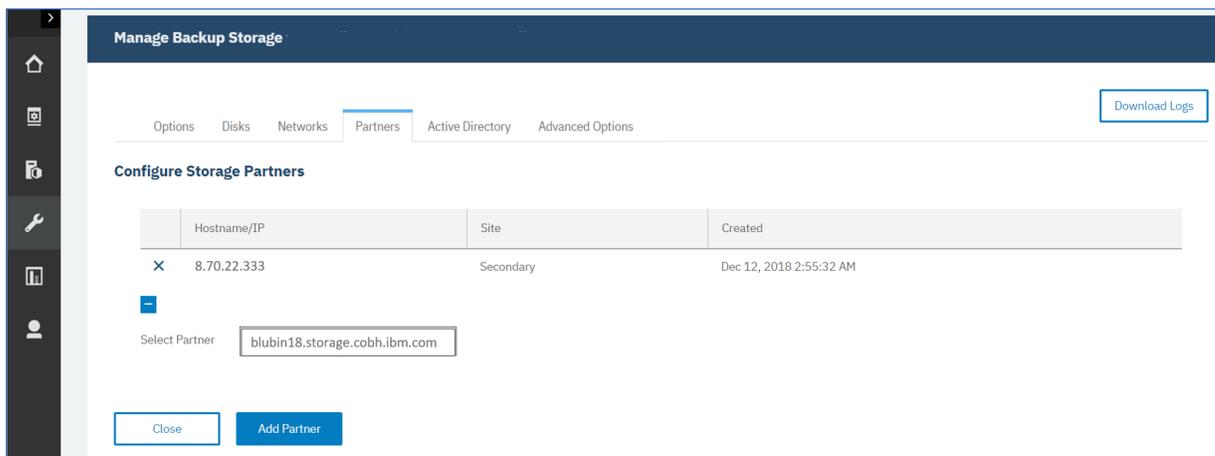
Before you begin

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

Procedure

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

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



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

Configuring network interface controllers

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

Before you begin

Versions of the vSnap server prior to V10.1.6 do not support this feature. To update a vSnap server, follow the instructions in [“Updating vSnap servers”](#) on page 14.

About this task

The network that is dedicated to sending management commands from IBM Spectrum Protect Plus to the vSnap server is indicated by the following icon in the **Network** page, .

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

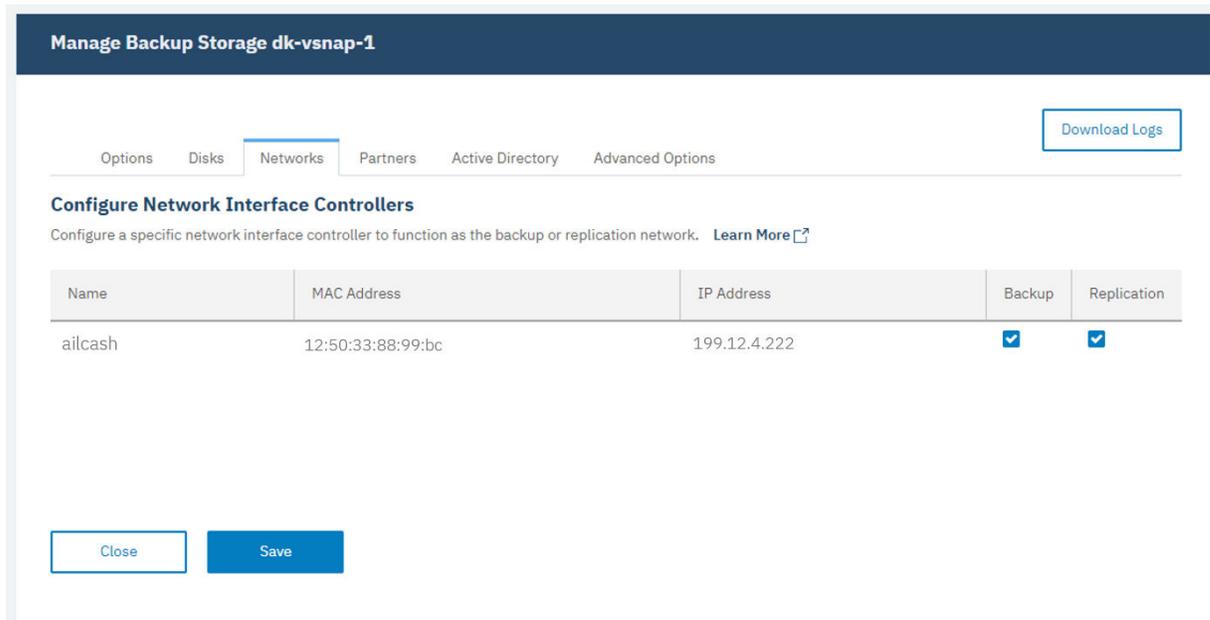
In the case that a second network interface card (NIC) is used, export `RMQ_SERVER_HOST` in the `/home/virgo/.bashrc` file on the IBM Spectrum Protect Plus appliance. The provided IP will be used for VADP to connect back to IBM Spectrum Protect Plus. The `<ip_address>` variable is the IP assigned to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

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

Procedure

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

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



The screenshot shows the 'Manage Backup Storage dk-vsnp-1' interface. At the top, there are tabs for 'Options', 'Disks', 'Networks', 'Partners', 'Active Directory', and 'Advanced Options'. A 'Download Logs' button is located in the top right corner. Below the tabs, the section is titled 'Configure Network Interface Controllers' with a 'Learn More' link. A table lists the network interface controllers with columns for Name, MAC Address, IP Address, Backup, and Replication. The table contains one entry for 'ailcash' with MAC Address '12:50:33:88:99:bc' and IP Address '199.12.4.222'. Both the Backup and Replication columns have checkmarks. At the bottom, there are 'Close' and 'Save' buttons.

Name	MAC Address	IP Address	Backup	Replication
ailcash	12:50:33:88:99:bc	199.12.4.222	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

3. Click **Save**.

Configuring an Active Directory

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

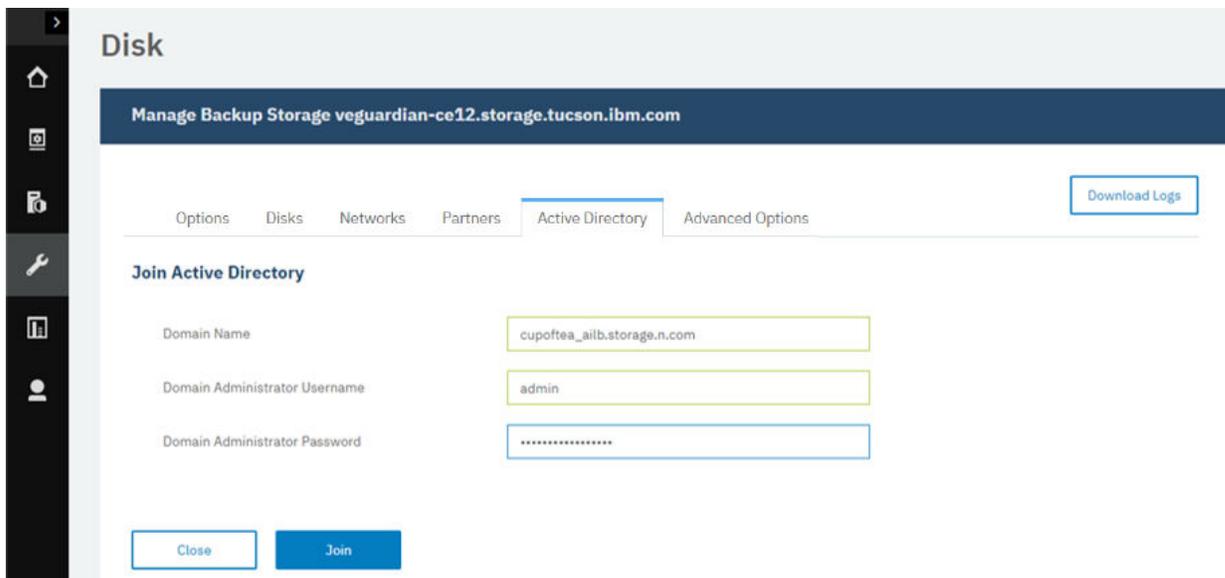
Before you begin

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

Procedure

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

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



The screenshot shows the 'Disk' configuration page in a web interface. The page title is 'Disk' and the subtitle is 'Manage Backup Storage veguardian-ce12.storage.tucson.ibm.com'. There are tabs for 'Options', 'Disks', 'Networks', 'Partners', 'Active Directory', and 'Advanced Options'. The 'Active Directory' tab is selected. Below the tabs is a 'Join Active Directory' form with three input fields: 'Domain Name' (containing 'cupoftea_aib.storage.n.com'), 'Domain Administrator Username' (containing 'admin'), and 'Domain Administrator Password' (containing a masked password). There are 'Close' and 'Join' buttons at the bottom of the form. A 'Download Logs' button is also visible in the top right corner.

4. Click **Join**.

Configuring advanced storage options

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

Procedure

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

1. In the navigation pane, click **System Configuration** , **Backup Storage** > **Disk**.
2. In the **Manage Backup Storage** pane, click the settings icon  that is associated with the host that you are managing.

3. On the **Advanced Options** tab, configure advanced options as shown in the following example:

The screenshot shows a web interface titled "Disk" with a sub-header "Manage Backup Storage". Below this is a navigation bar with tabs: "Options", "Disks", "Networks", "Partners", "Active Directory", and "Advanced Options" (which is selected). Underneath the navigation bar is a section titled "Set Advanced Options". This section contains four configuration rows, each with a label and a text input field:

Concurrent stream limit for copy to archive object storage	5
Concurrent stream limit for copy to standard object storage	5
Concurrent stream limit for replication	5
Interval in seconds between volume/snapshot deletions during space reclamation	300

At the bottom left of the configuration area is a "Close" button.

Figure 4. Manage backup storage advanced options.

- **Concurrent stream limit for copy to archive object storage:** This value defines the maximum number of concurrent streams that are used by this backup host when you are copying data to archive Object Storage.
- **Concurrent stream limit for copy to standard object storage:** This value defines the maximum number of concurrent streams that are used by this backup host when you are copying data to standard Object Storage.
- **Concurrent stream limit for replication:** This value defines the maximum number of concurrent streams that are used by this backup host when you are replicating data to other backup hosts.
- **Interval in seconds between volume/snapshot deletions during space reclamation:** This value defines the interval in seconds between successive deletions of volumes or snapshots on the vSnap server when space is reclaimed following a run of Maintenance jobs. Lowering the interval allows space to be reclaimed more aggressively, particularly when a large amount of data has expired in bulk.

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

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

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

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

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

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

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

4. Click **Close**.

Changing the throughput rate

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

Procedure

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

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

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

Site

Site Properties

Name:

Throughput Throttle

Enabled

Rate: MB/s

Schedule

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

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

Chapter 6. Managing secondary backup storage

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

Before you can access a secondary storage provider, you must add the provider to IBM Spectrum Protect Plus. During the process of adding the provider, you must supply an access key to ensure secure connection to the provider. A certificate might also be required.

You can add keys and certificates when you add a secondary backup storage provider to IBM Spectrum Protect Plus or you can add them in advance.

If you add keys and certificates in advance, select **Use existing key** or **Use existing certificate** and select the key or certificate when you add a secondary backup storage provider.

Related concepts

[“Copy snapshots to secondary backup storage” on page 3](#)

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

Managing cloud storage

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

Configuration for copying or archiving data to cloud

If you are planning to copy or archive IBM Spectrum Protect Plus data to cloud storage for long-term retention or for snapshot storage, you must configure secondary storage.

Tasks for configuring cloud storage

You must configure IBM Spectrum Protect Plus for backup and restore operations to cloud storage as shown in Table 1.

User scenario	Purpose	Steps
Store deduplicated data and non-deduplicated data in a cloud-container storage pool and restore the data as required.	Copy data to cloud storage. In the first copy operation, a full backup copy is created. Subsequent copies are incremental.	Choose one of the following providers: <ul style="list-style-type: none">• “Adding Amazon S3 Object Storage” on page 40• “Adding IBM Cloud Object Storage as a backup storage provider” on page 41• “Adding Microsoft Azure cloud storage as a backup storage provider” on page 42• “Adding S3 compatible object storage” on page 43

Adding Amazon S3 Object Storage

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

Before you begin

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

Ensure that cloud storage buckets are created for the IBM Spectrum Protect Plus data. For instructions about creating buckets, see [Amazon Simple Storage Service Documentation](#).

Procedure

To add Amazon S3 cloud storage as a backup Object Storage provider, complete the following steps:

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

Name

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

Region

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

Use existing key

Enable this option to select a previously entered key for the storage, and then select the key from the **Select a key** list.

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

Key name

Enter a meaningful name to help to identify the key.

Access key

Enter the AWS access key. Access keys are created in the AWS Management Console.

Secret key

Enter the AWS secret key. Secret keys are created in the AWS Management Console.

Enable Deep Archive

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

5. Click **Get Buckets** to connect IBM Spectrum Protect Plus to AWS to retrieve the list of available buckets.
6. Select the bucket that you plan to use as the copy target.
The **Standard object storage bucket** and **Archive object storage bucket** fields are displayed.
7. In the **Standard object storage bucket** field, select a bucket to serve as the copy target.
8. Optional: In the **Archive object storage bucket** field, select a cloud storage resource to serve as the archive target.
Archiving data creates a full data copy and can provide longer-term protection, cost, and security benefits.
9. Select **Deep Archive** to register Amazon S3 Glacier Deep Archive Buckets for long-term archiving.
10. Click **Register** to complete the operation.
The cloud storage is added to the cloud servers table.

Adding IBM Cloud Object Storage as a backup storage provider

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

Before you begin

When creating a bucket on IBM Cloud Object Storage (COS), ensure that both **Add Archive rule** and **Add Expiration rules** are not selected when creating buckets that are to be used for copy or archive. This can result in a failure with the “bucket has an unsupported lifecycle configuration” error when the job attempts to run in IBM Spectrum Protect Plus. The **Add Retention policy** option may be set for a bucket to be used for copy, but should not be set for a bucket that will be used for archiving.

The Cold Vault bucket of type should only be used when archiving, as it is the lowest-cost option and is described as ideal for long-term retention of data that will be minimally accessed.

When adding IBM Cloud Object Storage (COS), the method for obtaining the access and secret key will depend on the deployment model. If on-premise, keys can be obtained from the IBM COS Manager Console. For IBM COS IaaS, keys are created when a service account is created and can be obtained from the softlayer portal. If using IBM COS (COS as a Service), the access and secret key are not created by default; when a service account is created, check the **Include HMAC Credential** box, and add `{"HMAC": true}` to the **Add Inline Configuration Parameters** text area.

Procedure

To add IBM Cloud Object Storage as a backup storage provider, complete the following steps:

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

Name

Enter a meaningful name to help identify the cloud storage.

Endpoint

Select the endpoint of the cloud storage.

Use existing key

Enable to select a previously entered key for the storage, and then select the key from the **Select a key** list.

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

Key name

Enter a meaningful name to help to identify the key.

Access key

Enter the access key.

Secret key

Enter the secret key.

Certificate

Select a method of associating a certificate with the resource:

Upload

Select and click **Browse** to locate the certificate, then click **Upload**.

Copy and paste

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

Use existing

Select to use a previously uploaded certificate.

A certificate is not required if you are adding public IBM Cloud Object Storage.

5. Click **Get Buckets**, and then select a bucket to serve as the copy target.
After the buckets are generated, the **Standard object storage bucket** and **Archive object storage bucket** fields are displayed.
6. In the **Standard object storage bucket** field, select a bucket to serve as the copy target.
7. Optional: In the **Archive object storage bucket** field, select a cloud storage resource to serve as the archive target.
Archiving data creates a full data copy and can provide longer-term protection, cost, and security benefits. For more information about archiving data, see the information about copying data to cloud archive storage in [“Copy snapshots to secondary backup storage”](#) on page 3.
8. Click **Register**.
The cloud storage is added to the cloud servers table.

Adding Microsoft Azure cloud storage as a backup storage provider

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

Before you begin

Ensure that there are cloud storage buckets created for the IBM Spectrum Protect Plus data before you add the cloud storage in the following steps. For information how to create buckets, see Azure documentation.

Procedure

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

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

Name

Enter a meaningful name to help identify the cloud storage.

Endpoint

Select the endpoint of the cloud storage.

Use existing key

Enable to select a previously entered key for the storage, and then select the key from the **Select a key** list.

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

Key name

Enter a meaningful name to help identify the key.

Storage Account Name

Enter the Microsoft Azure access storage account name. This is from the Azure Management Portal.

Storage Account Shared Key

Enter the Microsoft Azure key from any one of the key fields in the Azure Management Portal, either key1 or key2.

5. Click **Get Buckets**, and then select a bucket to serve as the copy target.
After the buckets are generated, the **Standard object storage bucket** and **Archive object storage bucket** fields are displayed.
6. In the **Standard object storage bucket** field, select a bucket to serve as the copy target.
7. Optional: In the **Archive object storage bucket** field, select a cloud storage resource to serve as the archive target.

Archiving data creates a full data copy and can provide longer-term protection, cost, and security benefits. For more information about archiving data, see the information about copying data to cloud archive storage in [“Copy snapshots to secondary backup storage”](#) on page 3.

8. Click **Register**.

The cloud storage is added to the cloud servers table.

Adding S3 compatible object storage

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

Before you begin

Tip:

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

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

Ensure that cloud storage buckets are available. For more information about cloud storage buckets, see the documentation for the S3 compatible storage provider.

Procedure

To add S3 compatible cloud storage as a backup target, complete the following steps:

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

Name

Enter a meaningful name to help identify the cloud storage.

Endpoint

Enter the endpoint of the cloud storage.

Use existing access key

Enable this option to select a previously entered key for the storage, and then select the key from the **Select a key** list.

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

Key name

Enter a meaningful name to identify the key.

Access key

Enter the S3 compatible access key. For instructions about obtaining access keys, see the documentation for the S3 compatible storage provider.

Secret key

Enter the S3 compatible secret key. For instructions about obtaining access keys, see the documentation for the S3 compatible storage provider.

Certificate

Select the appropriate option to add a certificate for the S3 compatible storage:

Upload

To upload a certificate, click **Browse** to locate and select the certificate. Click **Upload**.

Copy and paste

Enter a name for the certificate and paste the certificate into the text area. Click **Create**.

Use existing

If a certificate exists, select the certificate from the **Select a certificate** list.

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

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

6. In the **Standard object storage bucket** field, select a bucket to serve as the backup target.
7. Optional: In the **Archive object storage bucket** field, select a cloud storage resource to serve as the archive target.

Archiving data creates a full data copy and can provide longer-term protection, cost, and security benefits. For more information about archiving data, see the information about copying data to cloud archive storage in [“Copy snapshots to secondary backup storage”](#) on page 3.

8. Click **Register**.

The cloud storage is added to the cloud servers table.

Editing settings for cloud storage

Edit the settings for a cloud storage provider to reflect changes in your cloud environment.

Procedure

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

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

Deleting cloud storage

Delete a cloud storage provider to reflect changes in your cloud environment. Ensure that the provider is not associated with any SLA policies before deleting the provider.

Procedure

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

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

Managing repository server storage

You can copy data to a repository server for longer-term data protection. For the current release of IBM Spectrum Protect Plus, the repository server must be an IBM Spectrum Protect server Version 8.1.7 or later. To copy data to tape, IBM Spectrum Protect server Version 8.1.8 or later is required.

You can choose to replicate the IBM Spectrum Protect Plus data that is copied to the IBM Spectrum Protect server to a target server. However, IBM Spectrum Protect Plus is not aware of subsequent IBM Spectrum Protect server replication operations and you cannot restore the replicated data from the target IBM Spectrum Protect server to IBM Spectrum Protect Plus.

Configuration for copying or archiving data to IBM Spectrum Protect

If you are planning to copy or archive IBM Spectrum Protect Plus data to an IBM Spectrum Protect server, there are three possible configurations. Choosing which one to configure depends on which scenario applies to your data protection needs. For each scenario, there are steps that are required in both the IBM Spectrum Protect Plus and IBM Spectrum Protect server environments to complete the setup.

Tasks for configuring IBM Spectrum Protect

You must configure the IBM Spectrum Protect server to communicate with the IBM Spectrum Protect Plus server, and to enable process requests for backup and restore operations. The Amazon Simple Storage Service (S3) protocol enables communication between the two servers.

User scenario	Purpose	Steps
Copying to standard object storage when you are running daily or less frequent copies to standard object storage.	Copy data to standard object storage. In the first copy operation, a full backup copy is created. Subsequent copies are incremental. Copying data to standard object storage is useful if you want relatively fast backup and recovery times and do not require the longer-term protection, cost, and security benefits that are provided by tape storage.	To copy data to standard object storage to the IBM Spectrum Protect server, you must create a cloud-container or directory-container storage pool, and set up the object agent component of IBM Spectrum Protect. Adding the object agent is a mandatory step. In addition to setting up the required storage pool, follow steps 2-4 listed, here .
Copying to tape when you are creating a weekly or less frequent full-copy of your data to tape storage. Important: Archiving data to tape cannot be run more frequently than once a week. Recovery time objectives (RTO) should be considered when recovering data from archive copies in your disaster recovery action plan. Therefore, for disaster recovery, recovering from archive data should only be used as a last resort.	When you copy data to tape, a full copy of the data is created at the time of the copy process. Copying data to tape provides extra security benefits. By storing tape volumes at a secure, offsite location that is not connected to the internet, you can help to protect your data from online threats such as malware and hackers. However, because copying to these storage types requires a full data copy, the time that is required to copy data increases. In addition, the recovery time can be unpredictable and the data might take longer to process before it is usable.	To copy data to tape, you must create a cloud-container or directory-container storage pool for tape, and a cold-data-cache storage pool on the IBM Spectrum Protect server. Adding the object agent is a mandatory step. Follow steps 1-4 listed, here .
Mixture of both standard object storage and long-term copying to tape	Secure your data in incremental backups on the IBM Spectrum Protect server, as well as retaining data on tape for longer term security.	This is a combination of the previous cases: data is stored to tape and data is stored on standard object storage at the IBM Spectrum Protect server. As well as setting up the required data storage pools for both scenarios, the creation of an object agent is mandatory.

The four steps required to set up and configure the data transfer communication between IBM Spectrum Protect Plus and the IBM Spectrum Protect server are as follows:

1. If you are setting up storage pools for copying data to tape follow Step1. Create storage pools on the IBM Spectrum Protect server by using the IBM Spectrum Protect Operations Center. For instructions, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape” on page 46](#). This step is required only if you are setting IBM Spectrum Protect for archiving with copies run once a week or less frequently.

2. Create a policy domain that points to the storage pool or pools. The policy domain defines the rules that control the backup services for IBM Spectrum Protect Plus. For instructions, see [“Step 2: Configuring an object policy domain”](#) on page 48.
3. If you are copying data to a standard storage pool or to tape, you must add standard object storage on the IBM Spectrum Protect server. For instructions, see [“Step 3: Setting up standard object storage”](#) on page 49.
4. Add an object agent on the IBM Spectrum Protect server. The object agent provides a gateway between the IBM Spectrum Protect Plus server and the IBM Spectrum Protect server. For instructions, see [“Step 4: Adding an object agent for copying data”](#) on page 52.
5. To complete the setup, you must add an object client on the IBM Spectrum Protect server. The object client identifies the IBM Spectrum Protect Plus server and enables it to store objects at the IBM Spectrum Protect server. The same credentials as those that you used for IBM Spectrum Protect Plus are used for the object client, which is the object client that is associated with the policy domain as set up in Step 2. For instructions to set up an object client, see [“Step 5: Adding and configuring an object client for copying data”](#) on page 53.

Tip: Alternatively, enter the **DEFINE STGPOOL** command to create a storage pool as described in the following topics:

What to do next

1. After you complete the tasks required for IBM Spectrum Protect storage, you must add the IBM Spectrum Protect server to IBM Spectrum Protect Plus. For information about how to do this, follow the instructions in [“Registering a repository server as a backup storage provider”](#) on page 55.
2. When that is done, you can create an SLA policy that defines the IBM Spectrum Protect server as the backup storage target. For more information to help you choose which type of policy you need, see [Managing SLA policies for backup operations](#).

Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape

Before you can copy data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server for archiving purposes, you must configure an object agent service. For long-term archiving of data, you must configure a cold data storage pool. If you are not planning to archive data to tape on the IBM Spectrum Protect server, you can skip this step.

About this task

Before you start, ensure that you have sized your cold cache storage needs by using the sizing tool and the Blueprints. For information about how to do this, see the [Blueprints](#).

Object client data that is specified with an S3 Glacier storage class is not frequently accessed. To enable the copying of this data, which is often called *cold data*, to tape storage, the data is written temporarily to a storage pool that meets the requirements for handling object data. The data is then moved to the tape device or VTL. This storage pool, called a *cold-data-cache storage pool*, is assigned to a policy domain for object clients. Only data from object clients can be written to or restored from a cold-data-cache storage pool.

Procedure

If you are not using the Operations Center, you can use the **define stgpool** command. The command can be defined as follows:

```
define stgpool NAME
stgtype=colddatacache
```

Note: To configure standard pools for object storage, follow these steps but when you define the type of storage pool, select Standard.

To configure the IBM Spectrum Protect server to copy data from an object client to physical tape media or a VTL, complete the following configuration steps:

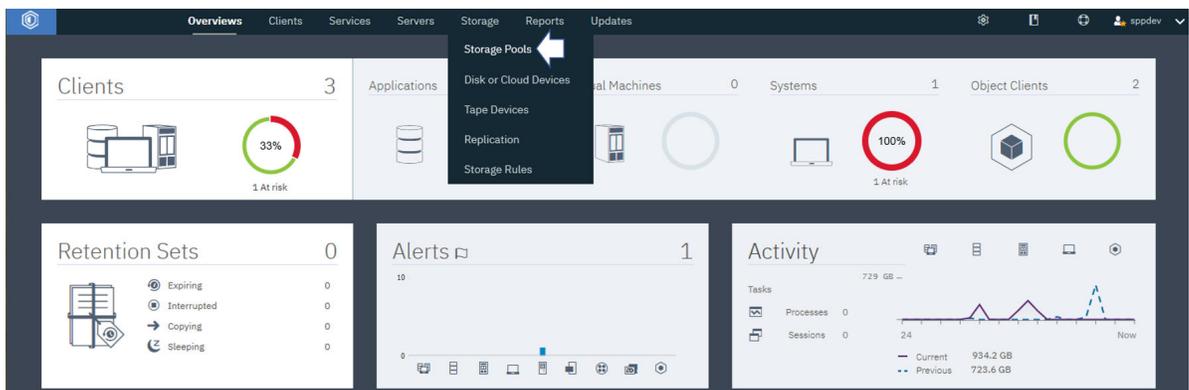
1. On the IBM Spectrum Protect server, configure a primary storage pool that represents a tape device or VTL. This primary storage pool is the destination for the object data that you want to copy.

Later, when you define the cold-data-cache storage pool, you must specify this tape pool as the next storage pool for the cold-data-cache pool.

Restrictions: The following restrictions apply to the tape storage pool:

- You cannot replicate object client data to or from the tape storage pool.
- The tape storage pool cannot be deduplicated.
- A next storage pool cannot be specified for the tape storage pool.

- a) On the Operations Center menu bar, click **Storage > Storage Pools**.



- b) On the **Storage Pools** page, click **Storage Pool** .

- c) In the **Add Storage Pool** wizard, select **Object Client** to enable object clients to copy data to tape.

2. Step through the wizard steps to configure a cold-data-cache storage pool.

A cold-data-cache storage pool consists of one or more file system directories on disk. It is an intermediary storage pool between the object client and a tape device or VTL and is linked to the primary sequential access storage pool that represents the tape device or VTL. Identify one or more existing file system directories for temporary disk storage and the primary sequential access storage pool that represents the tape device or VTL.

3. On the **Cold Data Cache** page, specify one or more existing file system directories for disk storage. Enter a fully qualified path name that conforms to the syntax that is used by the server operating system.

For example, enter `c:\temp\dir1\` for Microsoft Windows, or `/tmp/dir1/` for UNIX.

The object data is stored in sequential volumes in the file system directories. An object client can copy infrequently accessed data, or cold data, to physical tape media or to a VTL. When an object client copies cold data, the data is first stored in the cold data cache. The data is then migrated, without a migration delay, to the primary tape storage pool that represents the physical tape media or VTL. After the data is migrated to tape, it is deleted from the cold data cache. The cold data cache is used as a staging area for restoring cold data to the object client. During restore operations, the data is copied to the cold data cache. The data remains in the cold data cache for a period that is specified by the object client. Data is restored to the object client from the cold data cache, and not directly from the tape or VTL.

If you specify multiple directories for performance enhancement, ensure that the directories correspond to separate physical volumes. Although the cold data cache is used for temporary storage, it must be large enough to hold the data that is copied from the object client before the data is migrated to tape. It must also be large enough to hold data during restore operations for the period that is specified by the object client.

What to do next

When you complete the configuration of the cold data cache storage pool, create the object domain. For instructions about how to do that, see [“Step 2: Configuring an object policy domain” on page 48](#).

Step 2: Configuring an object policy domain

Before you copy data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, you must create and configure an object policy domain. The policy domain defines the rules that control the backup services for IBM Spectrum Protect Plus. You must add a standard storage pool which is with a directory or cloud container based storage for copies, and a cold pool if you are copying data to tape or archiving data.

Procedure

1. Verify the settings for the policy domain that you plan to use for copying data. Object clients that are defined or updated in the IBM Spectrum Protect server V8.1.8 or later must be assigned to policy domains that are created with the **DEFINE OBJECTDOMAIN** command. An object client node is associated with this policy domain when the node is registered or updated with the **REGISTER NODE** or **UPDATE NODE** command.

Restriction: Beginning with IBM Spectrum Protect server V8.1.8, all new object client nodes must be assigned to object policy domains.

For object client nodes that were assigned to non-object policy domains before V8.1.8, you do not have to update the assignment after you upgrade the server to IBM Spectrum Protect server V8.1.8. However, if any update to the object client node's domain is required, the node must be assigned to an object policy domain.

2. Review the following considerations for specifying policy domains for copy operations.
 - For IBM Spectrum Protect server, a policy domain can specify management classes for standard storage pools (cloud-container or directory-container storage pools), cold-data-cache storage pools, or both standard and cold-data-cache storage pools.

However, to copy data from IBM Spectrum Protect Plus, you must specify the following management classes depending on whether you are copying data to a cloud-container or directory-container storage pool or are copying data to a cold-data-cache storage pool for storage on physical tape media or in a virtual tape library (VTL):

- To copy data to a cloud-container or directory-container storage pool, use the **STANDARDPOOL** parameter to define the storage pool for the policy domain as shown in the following example:

```
define objectdomain mydomain standardpool=hotpool
```

- To copy data to a cold-data-cache storage pool, you must specify both a standard pool and a cold pool for the policy domain. A standard pool is required to store metadata that is used for restore and other IBM Spectrum Protect Plus operations. To define a cold-data-cache storage pool for the policy domain, use the **COLDPOOL** parameter, as shown in the following example:

```
define objectdomain mydomain standardpool=hotpool coldpool=coldpool
```

- All objects are uniquely named. There are no inactive versions of objects. When you define a policy domain, the following Storage Management policies are specified automatically:
 - The `Versions Data Exists` field is set to 1.
 - The `Retain Extra Versions` and the `Retain Only Version` fields are set to 0.
- The IBM Spectrum Protect Plus server controls the time when objects are deleted.

Example: Display detailed information about a policy domain for an IBM Spectrum Protect Plus copy operation

When the policy domain was created, it was assigned management classes and copy groups. You can use the **QUERY COPYGROUP** command to view information about the destination storage pools for the policy

domain. In the following example, the policy domain name is XYZ. The destination storage pools are HOTPOOL and COLDPOOL.

```
query copygroup xyz standard f=d
```

```
Policy Domain Name: XYZ
Policy Set Name: STANDARD
Mgmt Class Name: COLD
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 1
Versions Data Deleted: 1
Retain Extra Versions: 0
Retain Only Version: 0
Copy Mode: Modified
Copy Serialization: Shared Static
Copy Frequency: 0
Copy Destination: COLDPOOL
Table of Contents (TOC) Destination:
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 05/22/20 17:03:46
Managing profile:
Changes Pending: No

Policy Domain Name: XYZ
Policy Set Name: STANDARD
Mgmt Class Name: STANDARD
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 1
Versions Data Deleted: 1
Retain Extra Versions: 0
Retain Only Version: 0
Copy Mode: Modified
Copy Serialization: Shared Static
Copy Frequency: 0
Copy Destination: HOTPOOL
Table of Contents (TOC) Destination:
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 03/05/20 22:15:18
Managing profile:
Changes Pending: No
```

What to do next

After you create the object domain, proceed to the next step [“Step 3: Setting up standard object storage”](#) on page 49.

Step 3: Setting up standard object storage

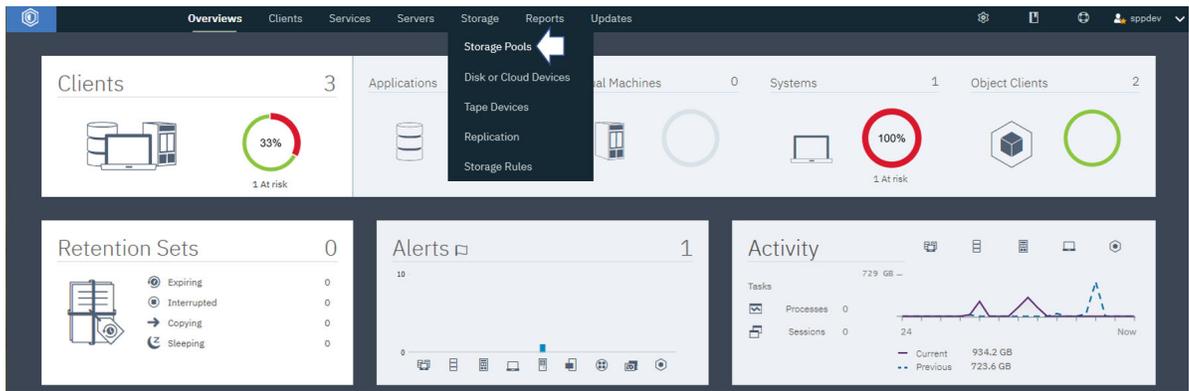
To set up standard object storage for copying data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, log in to the Operations Center and follow the procedure to set up storage pools. Complete the process by following the steps to create an object agent service by using the Operations Center wizard.

Before you begin

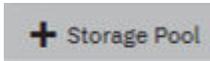
Before you start you must set up storage pools for standard storage or for copying to tape. If you are copying to tape, you must set up the cold data cache storage pool, and for standard object storage you must create and configure storage pools as required. For instructions about how to set up the cold data cache storage pool, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape”](#) on page 46.

Procedure

1. Create a directory-container storage pool by completing the following steps:
 - a) On the Operations Center menu bar, click **Storage > Storage Pools**.



b) On the **Storage Pools** page, click **Storage Pool**



c) Complete the steps in the **Add Storage Pool** wizard.

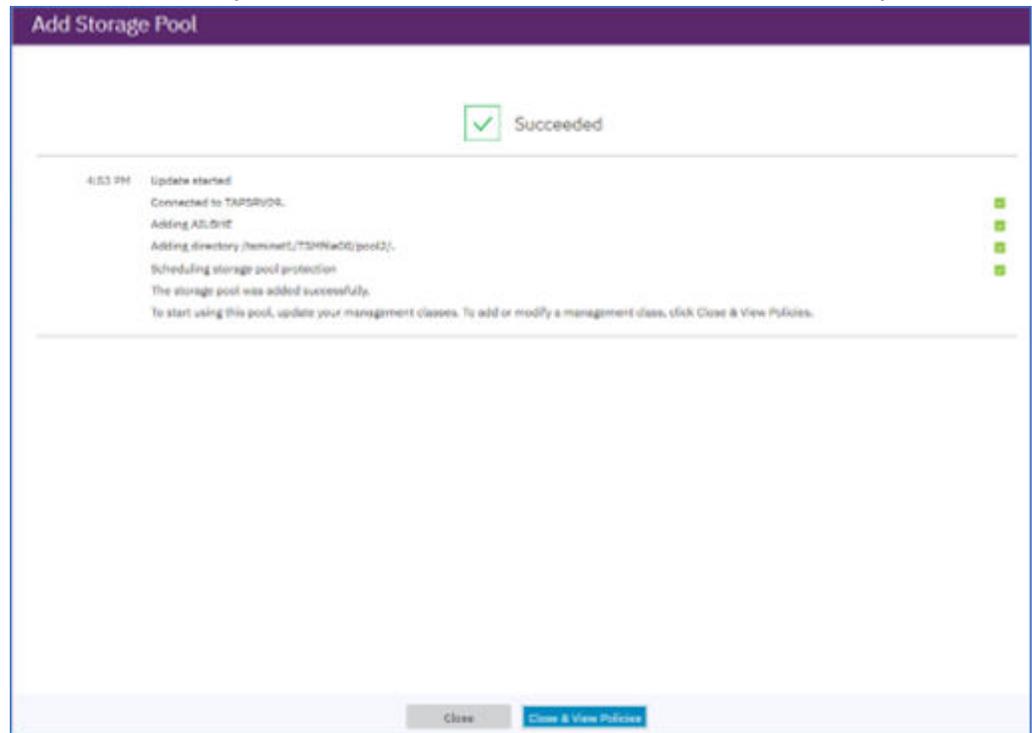
Tip: Select **Directory** for the type of container-based storage, and add directories with the + icon. Click **Next** to continue.

d) Review the **Protect Pool** summary, and click **Next**.

e) Specify an overflow pool if that is required.

f) Click **Add Storage Pool** to complete the creation of the storage pool.

If the operation was successful, you will see an icon to indicate success with a summary of the



storage pool.

2. In the **Services > Policies** page, select a policy, and click **Details**.

Policy Domain	Server	Clients	Mgmt Classes	Option Sets	Schedules	Default Mgmt Class	Backup Destination	Archive Destination	Migration
IBM_DEPLOY_CLI...	P9B-AIX1	0	1	0	0	IBM_DEPLOY_CLIENT		DEDUPPOOL	
JASON	P9B-AIX1	0	2	0	0	STANDARD	DEDUPPOOL		
P9B-AIX1_DATABA...	P9B-AIX1	0	4	0	1	BACKUP_DISK_KEE...	DEDUPPOOL		
P9B-AIX1_DB2	P9B-AIX1	0	1	0	0	BACK_ARCH_DISK	DEDUPPOOL	DEDUPPOOL	

- You can edit an existing domain policy by following these steps:
 - Update one or more management classes to use the new pool by editing the **Backup Destination** field of the table.
 - Click **Save**.
 - Or, you can create a new domain by running the **define objectdomain** command. For more information, see the previous step “Step 2: Configuring an object policy domain” on page 48.
3. On the **Details** page, click **Policy Sets**. Click the **Configure** toggle to make the policy sets editable.

JASON P9B-AIX1

Active policy set: STANDARD Activated Apr 1, 2020, 8:25 PM

Default management class: STANDARD

Active STANDARD Configure

Management Class	Default	Backup Destination
COLD		(None)
STANDARD	✓	DEDUPPOOL

Cancel Save

4. Change the Backup Destination to the newly created storage pool, or add a new management class,



to point to the new storage pool.

- Click **Activate**.
Changing the active policy set might result in data loss. A summary of the differences between the active policy set and the new policy set is displayed before the change is made.
- Review the differences between corresponding management classes in the two policy sets, and consider the consequences on client files. Client files that are bound to management classes in the currently active policy set are, after activation, bound to the management classes with the same names in the new policy set.
- Identify management classes in the currently active policy set that do not have counterparts in the new policy set, and consider the consequences on client files. Client files that are bound to these management classes are, after activation, managed by the default management class in the new policy set.
- If the changes implemented by the policy set are acceptable, select the **I understand that these updates can cause data loss** checkbox and click **Activate**.

What to do next

Create and configure an object client for the storage pool or pools you created. For more information, see [“Step 5: Adding and configuring an object client for copying data” on page 53](#)

Step 4: Adding an object agent for copying data

Before you can copy data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, you must add and configure the object agent. This step is the fourth step in setting up IBM Spectrum Protect Plus with the IBM Spectrum Protect server for archiving data or copying data to object storage.

Before you begin

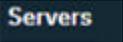
Ensure that the following steps are complete before you start to create the object client.

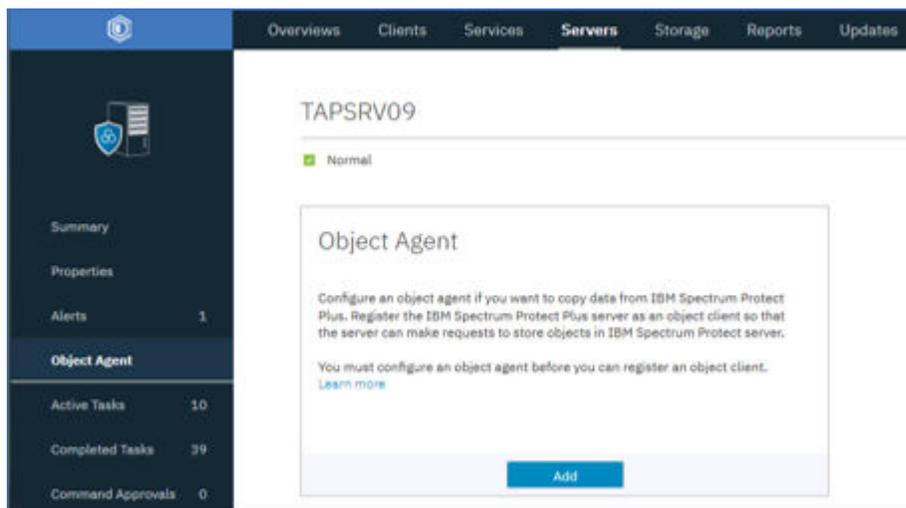
1. Ensure that you are logged in to the IBM Spectrum Protect server with an instance user ID.
2. Ensure that you have set up storage pools either for standard storage or for copying to tape. For instructions, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape” on page 46](#) or [“Step 3: Setting up standard object storage” on page 49](#).
3. Ensure that you have created an object domain.

About this task

This procedure is based on an environment where the IBM Spectrum Protect server is installed on an IBM AIX® operating system AIX Version 7.2 TL 1 and SP 4 or later, running on an IBM POWER8® or later server. (LINK TO a previous version)

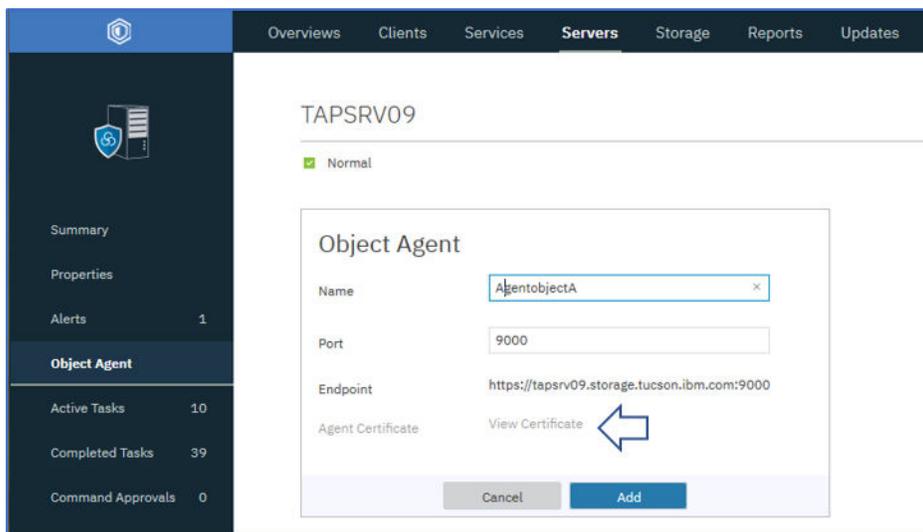
Procedure

1. On the Operations Center menu bar, click **Servers** .
2. Select a server and click **Details**.
3. From the navigation pane, click **Object Agent**; click **Add** to add an object agent.



Tip: If you are using the command line, run the **DEFINE SERVER** command to create an object agent. Specify **OBJECTAGENT=YES**. Follow the instructions in the command output. When these actions are completed, the object agent service automatically starts on the system that is hosting the IBM Spectrum Protect server.

4. To authenticate to the object agent, use the certificate that is generated.



5. Install the object agent service by running the command that can be copied from the wizard like in the following examples:

```
[root@servername-os: /]# /opt/tivoli/tsm/server/bin/spObjectAgent service install
/home/tsminst1/tsminst1/SPPOBJAGENT/spObjectAgent_SPPOBJAGENT_1500.config
2020-03-31 15:50:07.631021 I | Installed and started system service as
nameportnumberobjectagentname
```

Here is an example

```
[root@p9b-aix1: /]# /opt/tivoli/tsm/server/bin/spObjectAgent service install
/home/tsminst1/tsminst1/SPPOBJAGENT/spObjectAgent_SPPOBJAGENT_1500.config
2020-03-31 15:50:07.631021 I | Installed and started system service as spoa9000SPPOBJAGENT
```

6. Complete the configuration by starting an object agent service by running the **startObjectAgent** command. Here is an example for **AGENTOBJECTA** object agent.

```
"/opt/tivoli/tsm/server/bin/spObjectAgent" service install
"/home/tsminst1/tsminst1/AGENTOBJECTA/spObjectAgent_AGENTOBJECTA_1500.config"
```

7. Set up the object agent service to start automatically on startup by running a command similar to the following command for AIX:

```
spobj:2:once:/usr/bin/startsrc -s nameportnumberobjectagentname
```

Here is an example:

```
spobj:2:once:/usr/bin/startsrc -s spoa9000SPPOBJAGENT
```

Step 5: Adding and configuring an object client for copying data

Before you can copy data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, you must configure the object client. This step is the last step in setting up the IBM Spectrum Protect server for archiving and copying of data with the Operations Center.

Before you begin

Ensure that the following steps are complete before you start to create the object client.

1. Ensure that you are logged in to the IBM Spectrum Protect server with an instance user ID.
2. Ensure that the storage pools for either standard storage or for copying to tape are set up and ready. For instructions, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape”](#) on page 46 or [“Step 3: Setting up standard object storage”](#) on page 49.
3. Ensure that an object domain and an object agent are created before you start.

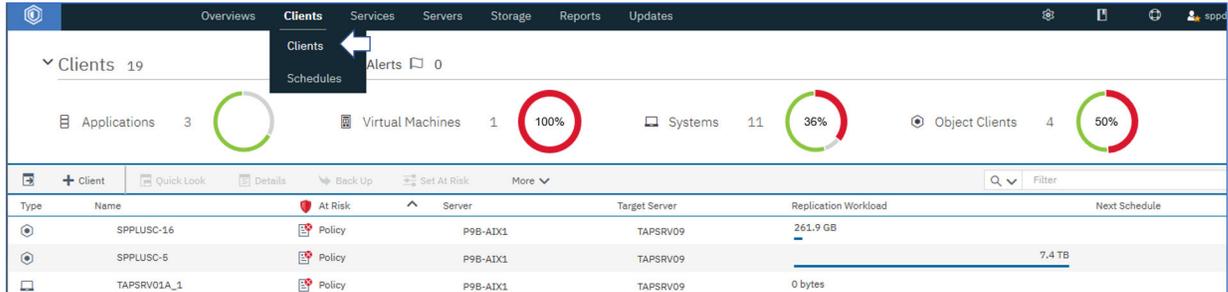
Tip: If you create an object client before you create the corresponding object agent, the **Add Client** wizard forces the creation of the object agent.

About this task

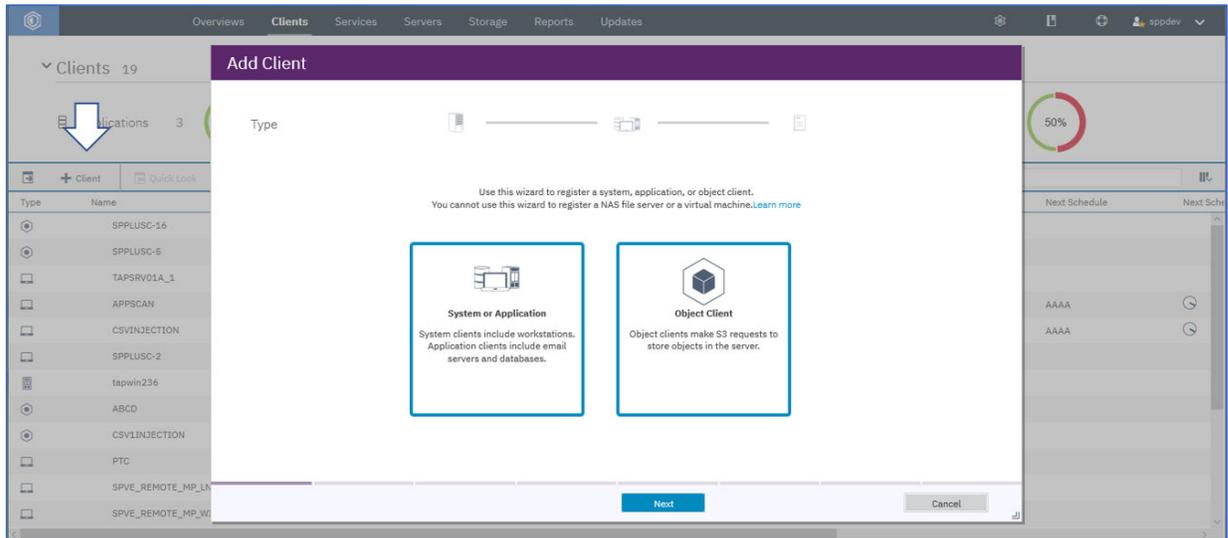
This procedure is based on an environment where the IBM Spectrum Protect server is installed on an IBM AIX operating system AIX Version 7.2 TL 1 and SP 4 or later, running on an IBM POWER8 or later server.

Procedure

1. On the Operations Center menu bar, click **Clients**.



2. Click **Client** to add a client as shown.



3. Select **Object Client** and click **Next** to start the **Add Client** wizard.

In the wizard screens, you are asked for to make the following choices and definitions for the client you are setting up.

- You can also choose to enable replication for this client.
- You must assign a client name and contact name, and an email address for reporting which you define in the final step of the wizard.
- You must assign a policy domain, which you set up in step 2, “[Step 2: Configuring an object policy domain](#)” on page 48.
- You can define at risk reporting for the client, such as a once-a-day report to the email address that you specified.

4. Click **Add Client**.

Note:

After the process finishes, you are provided with the endpoint for communicating with the object agent on the server, the access key ID, the secret access key, and the certificate for connecting securely.

When IBM Spectrum Protect Plus is an object client, it directs requests to the endpoint, and uses this information in the form of the access key ID, the secret access key, and the secure certificate.

Important: Ensure that a copy of each credential is saved to a secure location.

Tip: If you are using the command line, run the **REGISTER NODE** command to create an object client. Specify TYPE=OBJECTCLIENT. The script runs under the instance user ID.

What to do next

As a next step, you must register the IBM Spectrum Protect server as a repository server. For information about how to do this, see [“Registering a repository server as a backup storage provider”](#) on page 55. Once that is completed, you can create SLA policy jobs to copy data to the IBM Spectrum Protect server for standard storage or for archive to tape.

Registering a repository server as a backup storage provider

Add and register a repository server to enable IBM Spectrum Protect Plus to copy data to the server.

Before you begin

Configure the key and certificate that are required for the repository server. For instructions, see [Adding an access key](#) and [Adding a certificate](#).

For the current release of IBM Spectrum Protect Plus, the repository server must be an IBM Spectrum Protect server.

Configure IBM Spectrum Protect Plus as an object client to the IBM Spectrum Protect server. The object client node transfers and stores copied data. After you complete the setup procedure, the wizard provides you with the endpoint for communicating with the object agent on the server, and the access ID, secret key, and certificate for connecting securely.

Certificates can be obtained from the IBM Spectrum Protect server Operations Center by navigating to the following pane: **Server > Object Agent > Agent Certificate**. Alternatively, the certificate can be obtained from the IBM Spectrum Protect Plus appliance by running the following command: `openssl s_client -showcerts -connect <ip-address>:9000 </dev/null 2>/dev/null | openssl x509`

Copy retention settings are fully controlled through associated SLA policies in IBM Spectrum Protect Plus. IBM Spectrum Protect server copygroup retention settings are not used for copy operations.

Procedure

To add and register an IBM Spectrum Protect server as a backup storage provider, complete the following steps:

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

Name

Enter a meaningful name to help identify the repository server.

Hostname

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

Port

Enter the communications port of the repository server.

Use existing key

Enable to select a previously entered key for the repository, and then select the key from the **Select a key** list.

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

Key name

Enter a meaningful name to help to identify the key.

Access key

Enter the access key.

Secret key

Enter the secret key.

Certificate

Select a method of associating a certificate with the resource. If copying the certificate, the BEGIN and END lines of text must be included.

Upload

Select and click **Browse** to locate the certificate, then click **Upload**.

Copy and paste

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

Use existing

Select to use a previously uploaded certificate.

4. Click Register.

The IBM Spectrum Protect server is added to the repository servers table.

Related concepts

[“Configuration for copying or archiving data to IBM Spectrum Protect” on page 5](#)

If you are planning to copy or archive IBM Spectrum Protect Plus data to an IBM Spectrum Protect server, there are three possible configurations. Choosing which one to configure depends on which scenario applies to your data protection needs. For each scenario, there are steps that are required in both the IBM Spectrum Protect Plus and IBM Spectrum Protect server environments to complete the setup.

Editing settings for a repository server

Edit the settings for a repository server provider to reflect changes in your cloud environment.

Procedure

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

1. In the 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.

Chapter 7. vSnap server administration reference

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

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

Managing vSnap by using the command line interface

The vSnap server can be managed through the command-line interface and is the primary means of administering a vSnap server. Run the **vsnap** command from the vSnap server's interface after connecting through SSH using the user ID `serveradmin` or any other operating system user who has been assigned vSnap admin privileges. The initial `serveradmin` password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus” on page 12](#).

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

Managing vSnap by using the IBM Spectrum Protect Plus user interface

Some of the most common operations can also be completed from the IBM Spectrum Protect Plus user interface. Log in to the user interface and click **System Configuration > Backup Storage > Disk** in the navigation pane. Click the manage icon  for a vSnap server to edit its settings.

Related tasks

[“Managing vSnap servers” on page 19](#)

To enable backup and restore jobs, IBM Spectrum Protect Plus requires at least one vSnap server. The vSnap server is its own appliance, either deployed virtually or installed physically on a system that meets the minimum requirements. Each vSnap server in the environment must be registered in IBM Spectrum Protect Plus so that it is recognized.

[“Configuring advanced storage options” on page 35](#)

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

Storage management

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

Managing disks

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



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

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

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

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

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

Enabling encryption

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

vSnap encryption utilizes the following algorithm:

Cipher name

Advanced Encryption Standard (AES)

Cipher mode

xts-plain64

Key

256 bits

Linux Unified Key Setup (LUKS) header hashing

sha256

Managing encryption keys

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

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

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

2. Copy the key files to the temporary location:

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

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

Detecting disks

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

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

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

Showing disks

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

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

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

Showing storage pool information

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

Creating a storage pool

If you completed the simple initialization procedure that is described in [“Completing a simple initialization” on page 17](#), 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 59](#). For information about available options, use the `--help` option for any command or subcommand.

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

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

raid0

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

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

raid5

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

raid6

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

Expanding a storage pool

Before you expand a pool, ensure that one or more unused disks are available as described in [“Showing disks”](#) on page 59.

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

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

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

Managing the cache and log for storage pools

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

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

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



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

- `vsnap pool addcache`
- `vsnap pool addlog`
- `vsnap pool removecache`
- `vsnap pool removelog`

Installing kernel headers and tools

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

About this task

When vSnap is installed or updated, Linux kernel Version 4.19 is installed by default. If you opt out of the kernel upgrade to V4.19 and remain on the V3.10, a kernel V3.10 that is compatible with the vSnap server is installed and used. In both cases, kernel headers and tools associated with the kernel are not installed. If you plan to compile or use custom drivers, modules or other software, you must install the kernel packages. The Red Hat Package Manager (RPM) installers for the kernel headers and tools are available in the vSnap installation directory.

Procedure

1. Log on to the vSnap server as the `serveradmin` user. The initial password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in [“Start IBM Spectrum Protect Plus”](#) on page 12.
2. To determine the Linux kernel version, open a command line and issue the following command:

```
$ uname -r
```

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

```
$ 4.19.xxxx
```

3. Navigate to this directory:

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

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

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

Results

The kernel header or tool is installed.

User management

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

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

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

```
$ vsnap user create
```

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

--username <username>

Enter the username of the user.

--password <password>

Enter the password of the user.

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

```
$ vsnap user grant
```

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

--username <username>

Enter the username of the user.

--password <password>

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

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

```
$ vsnap user revoke
```

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

--username <username>

Enter the username of the user.

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

```
$ vsnap user show
```

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

```
$ vsnap user update
```

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

--username <username>

Enter the username of the user.

--password <old_password>

Enter the old password of the user.

--new_password <new_password>

Enter the new password of the user.

Chapter 8. Resetting the serveradmin password

The `serveradmin` account is a system user account that is preconfigured on IBM Spectrum Protect Plus and vSnap server. It is used to manage both of the virtual appliances. If you forget the `serveradmin` password, you must use the `root` account to reset the password. Because the password for the `root` account is not provided for deployments of IBM Spectrum Protect Plus or vSnap server, you must reset the `root` password and then reset the `serveradmin` through the virtual console.

About this task

Changing the `root` password will require that the IBM Spectrum Protect Plus virtual appliance be rebooted.

For IBM Spectrum Protect Plus, the preferred method is to avoid using the `root` account and instead use the `serveradmin` account for administration purposes.

Important: Pick a strong password that is a minimum of 15 characters in length and must contain at least one character from each of the classes (numbers, uppercase letters, lowercase letters, and other).

Procedure

1. Prepare to restart the IBM Spectrum Protect Plus virtual appliance by pausing all scheduled jobs. Then, wait for any running jobs to be completed.

Important: By ensuring that no jobs are running when you shut down the virtual appliance, you help to prevent possible issues.

2. Log in to the vSphere Client.
3. Restart the IBM Spectrum Protect Plus virtual appliance from the Actions menu by selecting Restart Guest OS.
4. Launch the web console. During the restart, the boot loader will appear. There will be at least two CentOS versions displayed.
5. Select the first CentOS version from the list and press the **e** key.
6. Locate the line that begins with `linux16 /vmlinuz`. In that line, locate `ro`.
7. Replace `ro` with the following string:

```
rw init=/sysroot/bin/sh
```

8. Press **Ctrl + X** to start in single user mode. A prompt appears.
9. Enter the `chroot` command:

```
:/# chroot /sysroot
```

10. Initiate the password change for the `root` account:

```
:/# passwd root
```

11. Enter the new password for the `root` account. You will be prompted to enter the password a second time.
12. Update the Security-Enhanced Linux parameters:

```
:/# touch /.autorelabel
```

13. Exit the current context with the `exit` command:

```
:/# exit
```

14. Restart the IBM Spectrum Protect Plus virtual appliance:

```
:/# reboot
```

15. Using secure shell (SSH) or the console, log in to the IBM Spectrum Protect Plus virtual appliance with the username `root` and the password that was created in a previous step.
16. Change the `serveradmin` password with the `passwd` command at the prompt:

```
:/# passwd serveradmin
```

17. Enter the new password for the `serveradmin` account. You will be prompted to enter the password a second time.
18. Close the SSH session to the IBM Spectrum Protect Plus virtual appliance.
19. Log in to IBM Spectrum Protect Plus using the `serveradmin` account with the newly created password to verify that the new password is set.

Results

The `root` and `serveradmin` account passwords for the IBM Spectrum Protect Plus virtual appliance will be reset to the specified password.

Chapter 9. Troubleshooting vSnap servers

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

Preventing job failures by synchronizing vSnap and CIFS passwords

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

About this task

For information about how to synchronize passwords, see [“User management” on page 61](#).

How do I tier data to tape or cloud storage?

You cannot tier data from IBM Spectrum Protect Plus to tape storage. You can tier data from IBM Spectrum Protect Plus to cloud storage, but only to cloud storage classes that support the rapid recall of data. When you are copying data to tape from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, it is not a good idea to use the IBM Spectrum Protect tiering function. If you are archiving data to tape, you must use a cold cache storage pool.

Review the guidelines about tape and cloud storage:

- Although you cannot tier data from IBM Spectrum Protect Plus to tape, you can archive or copy IBM Spectrum Protect Plus data to tape. To do this, define a cold-data-cache storage pool, as described in [Step 1: Creating a tape storage pool and cold-data-cache storage pool for copying data to tape](#).
- You can tier data from IBM Spectrum Protect Plus to cloud-container storage pools, but only to cloud storage classes that support the rapid recall of data. If you are using Amazon Web Services (AWS) with the Simple Storage Service (S3) protocol to move data to cloud container pools, do not move the data to Amazon S3 Glacier. For scenarios and instructions about copying or archiving data to cloud storage, see [Configuration for copying or archiving data](#). For instructions about tiering data to the cloud, see [Tiering data to cloud or tape storage](#) in the IBM Spectrum Protect product documentation.

You cannot tier data from IBM Spectrum Protect Plus to tape. To store IBM Spectrum Protect Plus data on tape, copy the data to an IBM Spectrum Protect server for storage on physical tape media or in a virtual tape library. For different scenarios and more information about how to set up storage, see [“Configuration for copying or archiving data to IBM Spectrum Protect” on page 5](#) and [“Configuration for copying or archiving data to cloud” on page 39](#). You

To set up a cold cache storage pool for archiving or copying data to tape, see [“Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape” on page 46](#).

Why is the vSnap server still offline?

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

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

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

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

VMware production or clone restore operations can use VMware SAN transport mode, which transports data in a storage area network (SAN) environment. To run a SAN-based restore operation, you can use the advanced setting **Enable Streaming (VADP) restore**, which was introduced in IBM Spectrum Protect Plus V10.1.5. This restore operation option is set by default. Coupled with this option, you can specify SAN transport mode in the VADP proxy options for a particular site.

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

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

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

Communication

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

Specifying SAN as a data transport mode

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

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

VADP Proxy

VADP Proxies

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

Total: 4

Auto Refresh

Proxy Details for spicemouse12.storage.clifden.golf.ie

- Server Address: spicemouse12.storage.clifden.golf.ie
- Site: Secondary
- Number of Cores: 8
- Available Memory: 8.8 GB

View Tasks

Actions: Suspend, Proxy Options, Uninstall, Unregister, Edit

4. From the **Transport Modes** list, select SAN.
5. Click **Save**.

How do I repair a failed source vSnap in an IBM Spectrum Protect Plus environment?

The vSnap servers in an IBM Spectrum Protect Plus environment provide disk storage for protecting data through backup and replication processes. You can repair and replace a failed vSnap server that is configured in your IBM Spectrum Protect Plus environment to act as the *source* for backup and replication services. The source vSnap server must be repaired so that backup and replication services can resume.

Before you begin

Important: It is assumed that all vSnap servers in the environment are protected by replication. If a vSnap server is not replicated and it fails, it cannot be recovered to a state that would allow it to continue as a disk storage source or target. In the absence of replication processes, you must create a new vSnap server and set up service level agreement (SLA) policies. When you run the policies, a new full backup process runs to the new vSnap server.

To determine which type of repair process is applicable to your vSnap server, see [technote 1103847](#).

About this task

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

This procedure establishes a new source vSnap server in your IBM Spectrum Protect Plus environment to replace the failed source vSnap server. The new source vSnap server will contain only the most recent recovery points.

Note: The version of the new vSnap server must match the version of the deployed IBM Spectrum Protect Plus appliance.

Procedure

1. Log in to the target vSnap server console with the ID `serveradmin` by using Secure Shell (SSH) protocol.

Enter the following command: `$ ssh serveradmin@MGMT_ADDRESS`

For example, `$ ssh serveradmin@10.10.10.2`

2. Obtain the ID of the failed source vSnap server by opening a command prompt and entering the following command:

```
$ vsnap partner show
```

The output is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
MGMT ADDRESS: 10.10.10.1
API PORT: 8900
SSH PORT: 22
```

3. Verify that the MGMT ADDRESS is the address of the failed source vSnap server. Take note of the failed source vSnap server's ID number.
4. In the environment with the source vSnap server, install a new vSnap server of the same type and version, and with the same storage allocation, as the failed source vSnap server.

For instructions about installing a vSnap server, see [Installing a physical vSnap server](#).

Important: Do not register the new vSnap server with IBM Spectrum Protect Plus. Do not use the Add Disk Storage wizard.

- a) You will first need to initialize the vSnap server with the following command:

```
$ vsnap system init --skip_pool --id partner_id
```

For example: `$ vsnap system init --skip_pool --id 12345678901234567890123456789012` using the failed source vSnap partner ID. A message indicates when the initialization is completed.

Note: This command is different to the vSnap initialization command listed in the IBM Knowledge Center and in the Blueprints.

5. Complete the vSnap server and pool creation process as outlined in *Chapter 5: vSnap Server Installation and Setup* in the [Blueprints](#).
6. Place the new source vSnap server into maintenance mode by entering the following command:

```
$ vsnap system maintenance begin
```

Placing the vSnap server into maintenance mode suspends operations such as snapshot creation, data restore jobs, and replication operations.

7. Initialize the new source vSnap server with the failed source vSnap server's partner ID. Enter the following command:

```
$ vsnap system init --id partner_id
```

The following command is an example: `$ vsnap system init --id 12345678901234567890123456789012`

8. On the new source vSnap server, add the partner vSnap servers. Each partner must be added separately. To add a partner, enter the following command:

```
$ vsnap partner add --remote_addr remote_ip_address --local_addr local_ip_address
```

where, *remote_ip_address* specifies the IP address of the source vSnap server, and *local_ip_address* specifies the IP address of the new source vSnap server.

The following command is an example:

```
$ vsnap partner add --remote_addr 10.10.10.2 --local_addr 10.10.10.1
```

9. When prompted, enter the user ID and password for the target vSnap server.

Informational messages indicate when the partners are created and updated successfully.

10. Create a repair task on the new source vSnap server by entering the following command:

```
$ vsnap repair create --async
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: N/A
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: PENDING
MESSAGE: The repair has been scheduled
```

11. Monitor the number of volumes that are involved in the repair operation by entering the following command:

```
$ vsnap repair show
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: 3
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: ACTIVE
MESSAGE: Created 0 volumes. There are 3 primary volumes that have recoverable snapshots,
the latest snapshot of each will be restored. Restoring 3 snapshots: 3 active, 0 pending, 0
completed, and 0 failed
```

The number of volumes that are involved in the repair operation is indicated in the TOTAL VOLUMES field.

12. Monitor the status of the repair task by viewing the repair.log file on the new source vSnap server, in the following directory /opt/vsnap/log/repair.log. Alternatively, you can enter the following command:

```
$ vsnap repair show
```

The output of this command is similar to the previous example. The following status messages can be displayed during the repair process:

- STATUS: PENDING indicates that the repair job is about to run.
- STATUS: ACTIVE indicates that the repair job is active.
- STATUS: COMPLETED indicates that the repair job is completed.
- STATUS: FAILED indicates that the repair job failed and must be resubmitted.

13. During the repair operation, run the vsnap repair show command to verify when the status is COMPLETED.

```
$ vsnap repair session show
```

The output of this command is similar to the following example:

```
ID: 1 RELATIONSHIP: 72b19f6a9116a46aae6c642566906b31
PARTNER TYPE: vsnap
LOCAL SNAP: 1313
REMOTE SNAP: 311
STATUS: ACTIVE
SENT: 102.15GB
STARTED: 2019-11-01 15:51:18 UTC
ENDED: N/A
Created 0 volumes.
There are 3 replica volumes whose snapshots will be restored on next replication.
```

A session for each volume involved in the repair operation is displayed.

Periodically issue the `$ vsnap repair session show` command to ensure that the amount of data being sent for each volume is increasing in increments. As the sessions finish you will see the status change to `COMPLETED`. When all the sessions finish, issue the `$ vsnap repair session show` command to verify that the overall status is `COMPLETED`. A final message indicating the number of volumes for which snapshots were restored is displayed. The message output is similar to the following example:

```
Created 0 volumes.  
There are 3 primary volumes that have recoverable snapshots, the latest snapshot of each  
will be restored.  
Restored 3 snapshots.
```

14. For any snapshots that are not restored and that indicate a `FAILED` status, resubmit the repair process by entering the following command:

```
$ vsnap repair create --async --retry
```

15. When the repair process reports a `COMPLETED` status, you can resume normal operations for the vSnap server by moving it out of maintenance mode. To resume normal processing, enter the following command:

```
$ vsnap system maintenance complete
```

16. Remove saved SSH host keys from the repaired source vSnap server and the target vSnap servers.

Run the following commands on both the source and target vSnap servers:

```
$ sudo rm -f /home/vsnap/.ssh/known_hosts
```

```
$ sudo rm -f /root/.ssh/known_hosts
```

Removing the SSH keys ensures that subsequent replication transfers do not produce errors that result from the changed host key of the repaired vSnap server.

17. Restart the vSnap service on the replaced server by entering the following command:

```
$ sudo systemctl restart vsnap
```

18. Click **System Configuration** > **Backup Storage** > **Disk** to verify that the new vSnap server is correctly registered, as follows:

- If the new vSnap server is using the same host name or IP address for registration, no change is required.
- If the new vSnap server is using a different host name or IP address for registration, you must update the registration by selecting the pencil icon.

19. To remove recovery points that are no longer available on the source vSnap server, start a maintenance job from the IBM Spectrum Protect Plus user interface.

For instructions, see [Creating jobs and job schedules](#).

Tip: You might see informational messages that are similar to the following example:

```
CTGGA1843 storage snapshot spp_1004_2102_2_16de41fcbc3 not found on live Storage2101  
Snapshot Type vsnap
```

20. To resume jobs that failed after the vSnap server became unavailable, run a storage server inventory job. For instructions, see [Creating jobs and job schedules](#)

Results

The source vSnap server has been repaired with only the most recent recovery points. The next backup job that runs as part of an SLA will back up data incrementally. If you create a restore job, only the most recent recovery point will be available in the backup repository. All other recovery points will be available in the replication repositories, and in the object storage and archive storage repositories if applicable to your environment.

How do I repair a failed target vSnap in an IBM Spectrum Protect Plus environment?

The vSnap servers in an IBM Spectrum Protect Plus environment provide disk storage for protecting data through backup and replication processes. You can repair and replace a failed vSnap server that is configured in your IBM Spectrum Protect Plus environment to act as the *target* for backup and replication services. The source vSnap server must be repaired so that backup and replication services can resume.

Before you begin

Important: It is assumed that all vSnap servers in the environment are protected by replication. If a vSnap server is not replicated and it fails, it cannot be recovered to a state that would allow it to continue as a disk storage source or target. In the absence of replication processes, you must create a new vSnap server and set up service level agreement (SLA) policies. When you run the policies, a new full backup process runs to the new vSnap server.

About this task

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

This procedure establishes a new target vSnap server in your IBM Spectrum Protect Plus environment to replace the failed target vSnap server. The new target vSnap server will not contain any data but will be populated with the most recent recovery points during the next scheduled replication operation.

Note: The version of the new vSnap server must match the version of the deployed IBM Spectrum Protect Plus appliance.

To determine which type of repair process is applicable to your vSnap server, see [technote 1103847](#).

Procedure

1. Log in to the functioning vSnap server console with the ID `serveradmin` by using Secure Shell (SSH) protocol.

Enter the following command: `$ ssh serveradmin@MGMT_ADDRESS`

For example, `$ ssh serveradmin@10.10.10.1`

2. Obtain the ID of the failed vSnap server by opening a command prompt and entering the following command:

```
$ vsnap partner show
```

The output is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
MGMT ADDRESS: 10.10.10.2
API PORT: 8900
SSH PORT: 22
```

3. Verify that the MGMT ADDRESS is the address of the failed vSnap server. Take note of the failed vSnap server's ID number.
4. In the environment with the target vSnap server, install a new vSnap server of the same type and version, and with the same storage allocation, as the failed target vSnap server.

For instructions about installing a vSnap server, see [Installing a physical vSnap server](#).

Important: Do not register the new vSnap server with IBM Spectrum Protect Plus. Do not use the Add Disk Storage wizard.

- a) You will first need to initialize the vSnap server with the following command:

```
$ vsnap system init --skip_pool --id <partner_id>
```

For example: `$ vsnap system init --skip_pool --id 12345678901234567890123456789012` using the failed source vSnap partner ID. A message indicates when the initialization is completed.

Note: This command is different to the vSnap initialization command listed in the IBM Knowledge Center and in the Blueprints.

5. Complete the vSnap server and pool creation process as outlined in *Chapter 5: vSnap Server Installation and Setup* in the [Blueprints](#).
6. Place the new vSnap server into maintenance mode by entering the following command:

```
$ vsnap system maintenance begin
```

Placing the vSnap server into maintenance mode suspends operations such as snapshot creation, data restore jobs, and replication operations.

7. Initialize the new target vSnap server with the failed target vSnap server's partner ID. Enter the following command:

```
$ vsnap system init --id <partner_id>
```

The following command is an example:

```
$ vsnap system init --id 12345678901234567890123456789012
```

8. On the new target vSnap server, add the partner vSnap servers. Each partner must be added separately. To add a partner, enter the following command:

```
$ vsnap partner add --remote_addr <remote_ip_address> --local_addr <local_ip_address>
```

where, `<remote_ip_address>` specifies the IP address of the source vSnap server, and `<local_ip_address>` specifies the IP address of the new target vSnap server.

The following command is an example:

```
$ vsnap partner add --remote_addr 10.10.10.1 --local_addr 10.10.10.2
```

9. When prompted, enter the user ID and password for the source vSnap server. Informational messages indicate when the partners are created and updated successfully.
10. Create a repair task on the new source vSnap server by entering the following command:

```
$ vsnap repair create --async
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: N/A
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: PENDING
MESSAGE: The repair has been scheduled
```

11. Monitor the number of volumes that are involved in the repair operation by entering the following command:

```
$ vsnap repair show
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: 3
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: ACTIVE
MESSAGE: Creating 3 volumes for partner 670d61a10f78456bb895b87c45e20999
```

The number of volumes that are involved in the repair operation is indicated in the TOTAL VOLUMES field.

12. Monitor the status of the repair task by viewing the repair.log file on the new source vSnap server, in the following directory /opt/vsnap/log/repair.log. Alternatively, you can enter the following command:

```
$ vsnap repair show
```

The output of this command is similar to the previous example. The following status messages can be displayed during the repair process:

- STATUS: PENDING indicates that the repair job is about to run.
 - STATUS: ACTIVE indicates that the repair job is active.
 - STATUS: COMPLETED indicates that the repair job is completed.
 - STATUS: FAILED indicates that the repair job failed and must be resubmitted.
13. During the repair operation, run the vSnap repair show command to verify when the status is COMPLETED.

```
$ vsnap repair session show
```

The final message indicates the number of volumes whose snapshots will be restored on the next replication, as follows:

```
Created 0 volumes.
There are 3 replica volumes whose snapshots will be restored on next replication.
```

14. For any snapshots that are not restored and indicate a FAILED status, resubmit the repair process by entering the following command:

```
$ vsnap repair create --async --retry
```

15. When the repair process reports a COMPLETED status, you can resume normal operations for the vSnap server by moving it out of maintenance mode. To resume normal processing, enter the following command:

```
$ vsnap system maintenance complete
```

16. Remove saved SSH host keys from the repaired source vSnap server and the target vSnap servers.

Run the following commands on both the source and target vSnap servers:

```
$ sudo rm -f /home/vsnap/.ssh/<known_hosts>
```

```
$ sudo rm -f /root/.ssh/<known_hosts>
```

Removing the SSH keys ensures that subsequent replication transfers do not produce errors that result from the changed host key of the repaired vSnap server.

17. Restart the vSnap service on the replaced server by entering the following command.

```
$ sudo systemctl restart vsnap
```

18. Click **System Configuration > Backup Storage > Disk** to verify that the new vSnap is correctly registered, as follows:
 - If the new vSnap server is using the same hostname or IP address for registration, no change is required.
 - If the new vSnap server is using a different hostname or IP address for registration, you must update the registration by selecting the pencil icon.
19. To remove recovery points that are no longer available on the source vSnap server, start a maintenance job from the IBM Spectrum Protect Plus user interface.

Tip: You might see informational messages that are similar to the following example:

```
CTGGA1843 storage snapshot spp_1004_2102_2_16de41fcbc3 not found on live Storage2101
Snapshot Type vsnap
```

20. To resume jobs that failed after the vSnap server became unavailable, run a storage server inventory job.

Results

The target vSnap server has been repaired. A new backup job must be run on the source vSnap server before any additional action is taken on the new target vSnap server.

If a replication job is attempted on the new target vSnap server, a message is displayed as follows:

```
CTGGA0289 - Skipping volume <volume_id> because there are no new snapshots since last backup
```

After a new backup job is run on the source vSnap server, the next scheduled replication job replicates the recovery points that are created by the backup job. At this point, if you create a restore job, only the most recent recovery point will be available in the replication repository. If the target vSnap server was also acting as a copy source to object or archive storage, the replication job must first run on the target vSnap server before any additional copy operations can complete successfully. The first copy of data to object storage will be a full copy.

How do I repair a failed dual-role vSnap in an IBM Spectrum Protect Plus environment?

You can repair and replace a failed vSnap server that is configured in your IBM Spectrum Protect Plus environment to act as both the *source* and *target* for backup and replication services.

About this task

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

This procedure establishes a new vSnap server in your IBM Spectrum Protect Plus environment to replace the failed vSnap server. After the repair process is completed, the new vSnap server is recovered to a point where backup jobs can continue to back up incremental changes (no full backup required) and replication jobs can continue.

To determine which type of repair process is applicable to your vSnap server, see [technote 1103847](#).

Note: The version of the new vSnap server must match the version of the deployed IBM Spectrum Protect Plus appliance.

Procedure

1. Log in to the functioning vSnap server in your environment console with the ID `serveradmin` by using Secure Shell (SSH) protocol.
Enter the following command: `$ ssh serveradmin@MGMT_ADDRESS`

For example, `$ ssh serveradmin@10.10.10.2`

2. Obtain the ID of the failed vSnap server by opening a command prompt and entering the following command:

```
$ vsnap partner show
```

The output is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
MGMT ADDRESS: 10.10.10.1
API PORT: 8900
SSH PORT: 22
```

3. Verify that the MGMT ADDRESS is the address of the failed vSnap server. Take note of the failed vSnap server's ID number.
4. On the target vSnap server, install a new vSnap server of the same type and version, and with the same storage allocation, as the failed source vSnap server.

For instructions about installing a vSnap server, see [Installing a physical vSnap server](#).

Important: Do not register the new vSnap server with IBM Spectrum Protect Plus. Do not use the Add Disk Storage wizard.

- a) You will first need to initialize the vSnap server with the following command:

```
$ vsnap system init --skip_pool --id partner_id
```

For example: `$ vsnap system init --skip_pool --id 12345678901234567890123456789012` using the failed source vSnap partner ID. A message indicates when the initialization is completed.

Note: This command is different to the vSnap initialization command listed in the IBM Knowledge Center and in the Blueprints.

5. Complete the vSnap server and pool creation process as outlined in *Chapter 5: vSnap Server Installation and Setup* in the [Blueprints](#).
6. Place the new vSnap server into maintenance mode by entering the following command:

```
$ vsnap system maintenance begin
```

Placing the vSnap server into maintenance mode suspends operations such as snapshot creation, data restore jobs, and replication operations.

7. Initialize the new target vSnap server with the failed target vSnap server's partner ID. Enter the following command to initialize the vSnap:

```
$ vsnap system init --id partner_id
```

The following command is an example: `$ vsnap system init --id 12345678901234567890123456789012`

8. On the new target vSnap server, add the partner vSnap servers. If there is more than one partner server, each partner must be added separately. To add a partner, enter the following command:

```
$ vsnap partner add --remote_addr remote_ip_address --local_addr local_ip_address
```

where, `remote_ip_address` specifies the IP address of the source vSnap server, and `local_ip_address` specifies the IP address of the new target vSnap server.

The following command is an example:

```
$ vsnap partner add --remote_addr 10.10.10.1 --local_addr 10.10.10.2
```

9. When prompted, enter the user ID and password for the source vSnap server.

Informational messages indicate when the partners are created and updated successfully.

10. Create a repair task on the new source vSnap server by entering the following command:

```
$ vsnap repair create --async
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: N/A
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: PENDING
MESSAGE: The repair has been scheduled
```

11. Monitor the number of volumes that are involved in the repair operation by entering the following command:

```
$ vsnap repair show
```

The output of this command is similar to the following example:

```
ID: 12345678901234567890123456789012
PARTNER TYPE: vsnap
PARTNER ID: abcdef7890abcdef7890abcdef7890ab
TOTAL VOLUMES: 6
SNAPSHOTS RESTORED: N/A
RETRY: No
CREATED: 2019-11-01 15:49:31 UTC
UPDATED: 2019-11-01 15:49:31 UTC
ENDED: N/A
STATUS: ACTIVE
MESSAGE: Created 0 volumes
There are 3 replica volumes whose snapshots will be restored on next replication.
There are 3 primary volumes that have recoverable snapshots, the latest snapshot of each
will be restored.
The number of volumes that are involved in the repair operation are indicated in the TOTAL
VOLUMES field
```

12. Monitor the status of the repair task by viewing the repair.log file on the new source vSnap server, in the following directory /opt/vsnap/log/repair.log. Alternatively, you can enter the following command:

```
$ vsnap repair show
```

13. When the status of the repair operation is in the ACTIVE state, you can view the status of individual repair sessions by entering the following command:

```
$ vsnap repair session show
```

The output is similar to this example:

```
ID: 1
RELATIONSHIP: 72b19f6a9116a46aae6c642566906b31
PARTNER TYPE: vsnap
LOCAL SNAP: 1313
REMOTE SNAP: 311
STATUS: ACTIVE
SENT: 102.15GB
STARTED: 2019-11-01 15:51:18 UTC
ENDED: N/A
```

View a session for each of the source volumes in the repair operation. The amount of data that is sent for each volume shows increasing incremental values until the process completes. The final message indicates the number of volumes whose snapshots will be restored by the next replication operation, as shown in this example:

```
Created 0 volumes. There are 3 replica volumes whose snapshots will be restored on next
replication.
```

14. For any snapshots that are not restored and indicate a FAILED status, resubmit the repair process by entering the following command:

```
$ vsnap repair create --async --retry
```

15. When the repair process reports a COMPLETED status, you can resume normal operations for the vSnap server by moving it out of maintenance mode. To resume normal processing, enter the following command:

```
$ vsnap system maintenance complete
```

16. Optional: To view the total volumes and number of snapshots that were restored during the repair operation, run the show command for the vSnap server.

The output includes the following information:

- **Total volumes** lists the total number of volumes that were inspected during the repair operation. This list includes the source volumes (primary volumes) where the latest recovery point backup was restored, and target volumes (replica volumes) that are repopulated during upcoming replication operations as scheduled in SLAs.
- **SNAPSHOTS RESTORED** lists the number of source volumes that were restored.

17. Remove saved SSH host keys from the repaired source vSnap server and the target vSnap servers.

Run the following commands on both the source and target vSnap servers:

```
$ sudo rm -f /home/vsnap/.ssh/known_hosts
```

```
$ sudo rm -f /root/.ssh/known_hosts
```

Removing the SSH keys ensures that subsequent replication transfers do not produce errors that result from the changed host key of the repaired vSnap server.

18. Restart the vSnap service on the replaced server by entering the following command:

```
$ sudo systemctl restart vsnap
```

19. Click **System Configuration > Backup Storage > Disk** to verify that the new vSnap server is correctly registered, as follows:

- If the new vSnap server is using the same hostname or IP address for registration, no change is required.
- If the new vSnap server is using a different hostname or IP address for registration, you must update the registration by selecting the pencil icon.

20. To remove recovery points that are no longer available on the source vSnap server, start a maintenance job from the IBM Spectrum Protect Plus user interface.

Follow the instructions here to do this, [Creating jobs and job schedules](#).

Tip: You might see informational messages that are similar to the following example:

```
CTGGA1843 storage snapshot spp_1005_2102_2_16de41fcbc3 not found on live Storage2101  
Snapshot Type vsnap
```

21. To resume jobs that failed after the vSnap server became unavailable, run a storage server inventory job. For instructions, see [Creating jobs and job schedules](#).

Results

For primary backup data that is stored on the repaired vSnap server, the latest recovery point for primary backup data is now available. Subsequent backups to the repaired vSnap server continue to send only incremental changes since the last backup. For replicated data stored on the repaired vSnap server, no replicated data is available immediately after the repair. Subsequent replication jobs from the partner vSnap server will repopulate any backups that are created on the partner vSnap server after the repair process was completed. If a replication job is attempted on the partner vSnap server before a backup is completed on the partner vSnap server, a warning message is displayed indicating that there are no new snapshots since the last backup:

CTGGA0289 - Skipping volume <volume_id> because there are no new snapshots since last backup

If the repaired vSnap server was acting as a copy source to object or archive storage, a backup job must first be run on the repaired vSnap server before any additional copy operations will be successful. The first copy of data to object storage will be a full copy.

How do I delete and recreate a vSnap storage pool?

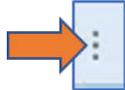
When a scenario arises that results in the requirement to delete a vSnap storage pool due to corruption or any other reason, you can follow the steps to delete and recreate the storage pool. This procedure is a destructive operation that discards all data in an existing vSnap storage pool. All backup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

Procedure

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

For more information about unregistering the vSnap server, see [“Unregistering a vSnap server” on page 28](#).

2. Run a maintenance job on the vSnap server by opening **Job and Operations > Schedule**. Find the

Maintenance job in the list. Click the actions icon,  and click **Start**.

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

For more information about maintenance jobs, see [Job types](#).

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

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

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

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

```
$ vsnap pool show
```

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

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

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

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

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

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

```
$ sudo reboot -n
```

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

7. When the restart finishes, run the status command:

```
$ vsnap_status
```

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

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

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

```
$ vsnap disk show
```

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

Sample output

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

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

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

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

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

10. Create the new pool with the following command:

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

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

The following command is an example of the create command:

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

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

Tip:

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

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

```
$ vsnap pool show
```

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

12. Register the vSnap server in IBM Spectrum Protect Plus under a chosen site to finalize the setup. For more information about how to register a vSnap server, see [“Registering a vSnap server as a backup storage provider”](#) on page 19.

Chapter 10. Product messages

IBM Spectrum Protect Plus components send messages with prefixes that help to identify which component they come from. Use the search option to find a particular message by using its unique identifier.

Messages consist of the following elements:

- A five-letter prefix.
- A number to identify the message.
- Message text that is displayed on screen and written to message logs.

Tip: Use your browser's search capability by using Ctrl+F to find the message code you are looking for.

The following example contains the Db2® agent prefix. When you click More, extra details that explain the reason for the message are shown.

```
Warning
Apr 16, 2019
9:14:37 AM
GTGGH0098
[myserver1.myplace.irl.ibm.com]
Database AC7 will not be backed up as it is ineligible for the backup operation. More
```

IBM Spectrum Protect Plus message prefixes

Messages have different prefixes to help you to identify the component that issues the message.

The following table identifies the prefix that is associated with each component.

Prefix	Component
CTGGA	IBM Spectrum Protect Plus
CTGGE	IBM Spectrum Protect Plus for Microsoft SQL Server
CTGGF	IBM Spectrum Protect Plus for Oracle
CTGGG	IBM Spectrum Protect Plus for Microsoft Exchange Server
CTGGH	IBM Spectrum Protect Plus for IBM Db2
CTGGI	IBM Spectrum Protect Plus for MongoDB
CTGGK	IBM Spectrum Protect Plus for Containers
CTGGL	IBM Spectrum Protect Plus for Amazon EC2
CTGGR	IBM Spectrum Protect Plus for Microsoft Office 365
CTGGT	IBM Spectrum Protect Plus for file systems

For a list of all messages, see IBM Knowledge Center [here](#).

Appendix A. Search guidelines

Use filters to search for an entity such as a file or a restore point.

You can enter a character string to find objects with a name that exactly matches the character string. For example, searching for the term `string.txt` returns the exact match, `string.txt`.

Regular expression search entries are also supported. For more information, see [Search Text with Regular Expressions](#).

You can also include the following special characters in the search. You must use a backslash (`\`) escape character before any of the special characters:

```
+ - & | ! ( ) { } [ ] ^ " ~ * ? : \
```

For example, to search for the file `string[2].txt`, enter the `string\[2\].txt`.

Searching with wildcards

You can position wildcards at the beginning, middle, or end of a string, and combine them within a string.

Match a character string with an asterisk

The following examples show search text with an asterisk:

- `string*` searches for terms like `string`, `strings`, or `stringency`
- `str*ing` searches for terms like `string`, `straying`, or `straightening`
- `*string` searches for terms like `string` or `shoestring`

You can use multiple asterisk wildcards in a single text string, but multiple wildcards might considerably slow down a large search.

Match a single character with a question mark

The following examples show search text with a question mark:

- `string?` searches for terms like `strings`, `stringy`, or `string1`
- `st??ring` searches for terms like `starring` or `steering`
- `???string` searches for terms like `hamstring` or `bowstring`

Appendix B. Accessibility features for the IBM Spectrum Protect product family

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

Overview

The IBM Spectrum Protect family of products includes the following major accessibility features:

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

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

The product documentation in IBM Knowledge Center is enabled for accessibility.

Keyboard navigation

This product uses standard navigation keys.

Interface information

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

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

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

Vendor software

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

Related accessibility information

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

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

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

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

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

Trademarks

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

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

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

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

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

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

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

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

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

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

Terms and conditions for product documentation

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

Applicability

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

Personal use

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

Commercial use

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

Rights

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

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

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

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

Privacy policy considerations

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

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

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

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

Glossary

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

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

Index

A

- accessibility features [85](#)
- adding
 - vSnap servers [19](#)
- Advanced backup options [35](#)

B

- backup storage
 - advanced options, managing [35](#)
 - storage options, managing disks [31](#)
 - storage options, managing partners [33](#)
- backup storage server
 - storage options, managing [33](#), [35](#)

C

- cloud provider
 - deleting [44](#)
 - editing [44](#)
- cloud server
 - adding a Microsoft azure cloud resource [42](#)
 - adding an Amazon S3 [40](#)
 - adding an IBM Cloud Object Storage resource [41](#)
 - adding an s3 compatible cloud resource [43](#)
- cold-data-cache storage pool [46](#)
- Configuring backup storage
 - storage options, adding disks [32](#)
- copying data to tape [46](#)

D

- data copy to tape
 - configuring [46](#)
- data protection [52](#), [53](#)
- DEFINE STGPOOL command [46](#)
- deleting
 - demo [26](#)
- demo
 - site [26](#)
 - SLA [26](#)
 - vSnap [26](#)
- disability [85](#)

F

- files
 - searching for [83](#)

I

- IBM Knowledge Center [v](#)
- IBM spectrum protect server
 - adding a repository server [55](#)
 - registering a repository server [55](#)

- installing
 - vSnap servers
 - Hyper-V environment [11](#)
 - physical environment [9](#)
 - VMware environment [10](#)

K

- keyboard [85](#)
- Knowledge Center [v](#)

L

- localhost
 - vSnap [26](#)

M

- message
 - prefixes [81](#)
- messages [81](#)

N

- Network configuration [33](#)
- New in IBM Spectrum Protect Plus Version Version 10.1.8 [vii](#)
- NICs [33](#)

O

- object client [52](#), [53](#)
- Object Storage
 - Amazon S3 [40](#)

P

- publications [v](#)

R

- registering
 - vSnap servers [19](#)
- removing
 - demo [26](#)
- Replication partners [33](#)
- repository server provider
 - deleting [56](#)
 - editing [56](#)
- Rescan
 - After expanding storage [32](#)
- Rescan vSnap [32](#)

S

- starting

starting (*continued*)
IBM Spectrum Protect Plus [12](#)

U

Updating
vSnap server [14](#)

V

virtual environments [52](#), [53](#)

vSnap

updating [15](#)

vSnap server

administering

kernel headers

kernel tools [60](#)

storage administration [57](#)

user administration [61](#)

change throughput [37](#)

editing [25](#)

initializing

advanced [17](#)

simple [17](#)

replication partnership, establishing [21](#)

storage pools, expanding [20](#)

Unregistering [28](#)

vSnap servers

adding [19](#)

installing

Hyper-V environment [11](#)

physical environment [9](#)

VMware environment [10](#)

registering [19](#)

uninstalling [16](#)

SPP Reuse URLs

[System requirements](#)

[System requirements](#)

[Creating catalog backup job](#)

[Permissions](#)

[Backing up the IBM Spectrum Protect Plus application](#)

[Adding a key](#)

[Adding a certificate](#)

[Backing up and restoring VMware data](#)

[Virtual machine privileges](#)

[Creating a user account for an individual user](#)

[SLA policies](#)

[Configuring global preferences](#)

[Managing SLA policies](#)

[Assigning a static IP address](#)

[VADP system requirements](#)

[Adding an access key](#)

[Installing a vSnap server](#)

[Tiering](#)

[Installing a physical vSnap](#)

[Creating jobs and job schedules](#)

[Job types](#)

[Product messages](#)



Product Number: 5737-F11