IBM Spectrum Protect Plus
10.1.8

Installation and User's Guide



Note:

Before you use this information and the product it supports, read the information in "Notices" on page 685.

Seventh edition (1st October 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....xi
Who should read this publication.... xi
Publications xi

What's new in Version 10.1.8.... xiii

Getting involved in product development....xv
Sponsor user program.... xv
Beta program.... xv

Chapter 1. Product overview....1
Deployment storyboard.... 1
Product components....6
Overview of the serveradmin user account....9
Product dashboard.... 10
Alerts.... 11
Role-based access control.... 12
Incremental forever backup strategy....13
Replicate backup-storage data.... 14
Copy snapshots to secondary backup storage.... 15
IBM Spectrum Protect Plus on IBM Cloud.... 17
IBM Spectrum Protect Plus as a software offering on IBM Cloud.... 17
IBM Spectrum Protect Plus as a VMware service on IBM Cloud.... 18
IBM Spectrum Protect Plus on AWS....18
IBM Spectrum Protect Plus on Azure....19
Integration with IBM Spectrum Protect....20
Adding IBM Spectrum Protect Plus to the Operations Center....20
Entering the Operations Center URL....23
Accessing the Operations Center....24
Integration with IBM Cloud Pak® for Multicloud Management....25

Chapter 2. Installation overview.... 27
System requirements 27
Component requirements 27
Hypervisor and cloud instance requirements 57
File indexing and restore requirements....64
File system requirements....71
Container Backup Support requirements....75
Db2 requirements.... 83
Microsoft Exchange Server requirements.... 89
MongoDB requirements.... 96
Microsoft 365 requirements.... 101
Oracle requirements.... 106
Microsoft SQL Server requirements....113
Post installation tasks....121
Assigning a static IP address....121
Uploading the product key....121
Editing firewall ports.... 122

Chapter 3. Installing IBM Spectrum Protect Plus as a virtual appliance.... 125

Overview of virtual appliance deployment ... 125
Obtaining the IBM Spectrum Protect Plus installation package ... 125
Installing IBM Spectrum Protect Plus as a VMware virtual appliance ... 126
Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance ... 128

Chapter 4. Installing IBM Spectrum Protect Plus as a set of containers ... 131
Overview of container deployment ... 131
Support matrix ... 132
Storage requirements ... 132
Roles in IBM Spectrum Protect Plus ... 133
Installing on OpenShift Container Platform ... 133
Preparing to install the operator from IBM Entitled Registry ... 134
Installing the operator in an online environment ... 137
Installing the operator in an airgap environment ... 139
Installing the operator for IBM Cloud Pak for Multicloud Management at the command line ... 145
Creating an IBM Spectrum Protect Plus instance ... 147
Verifying the installation of the IBM Spectrum Protect Plus instance ... 149
Uninstalling IBM Spectrum Protect Plus containers ... 150
Uninstalling the instance ... 150
Uninstalling the operator ... 151

Chapter 5. Installing and managing vSnap servers ... 153
Installing a vSnap server ... 153
Installing a physical vSnap server ... 153
Installing a virtual vSnap server in a VMware environment ... 154
Installing a virtual vSnap server in a Hyper-V environment ... 155
Uninstalling a vSnap server ... 157
Managing vSnap servers ... 157
Registering a vSnap server ... 158
Initializing the vSnap server ... 169
Migrating onboard vSnap data to a stand-alone vSnap server ... 170
Expanding a vSnap storage pool ... 174
Changing the throughput rate ... 175
Replacing a failed vSnap server ... 176
Installing iSCSI initiator utilities ... 176
vSnap server administration reference ... 176
Troubleshooting vSnap servers ... 183

Chapter 6. Installing Container Backup Support ... 197
Prerequisites ... 197
Installing Helm 3 and renaming the binary file ... 197
Installing Velero to protect resources ... 198
For Kubernetes: Verifying whether Metrics Server is running ... 204
Defining the application and persistent volume claim relationship ... 205
Kubernetes or OpenShift: Installing from the command line ... 205
Setting up the installation variables ... 206
Installing in an airgap environment ... 214
Installing from IBM Helm Charts Repository ... 218
OpenShift: Installing from the OpenShift web console ... 220
Adding the IBM Entitled Registry to your Helm repository ... 221
Creating a project for Container Backup Support ... 222
Creating image pull secrets ... 222
Creating the credentials secret ... 223
Creating custom resource definitions ... 224
Installing Container Backup Support from the OpenShift web console ... 224
Updating Container Backup Support ... 227
Updating your credentials after installation ... 229

Uninstalling Container Backup Support....231
Configuration parameters....232

Chapter 7. Updating IBM Spectrum Protect Plus components.... 239
Updating IBM Spectrum Protect Plus in a virtual appliance environment.... 239
Managing updates....239
Updating the IBM Spectrum Protect Plus server.... 242
Updating IBM Spectrum Protect Plus in a container environment.... 244
Updating IBM Spectrum Protect Plus by using the OpenShift web console.... 244
Updating IBM Spectrum Protect Plus by using the user interface.... 246
Updating IBM Spectrum Protect Plus in a cloud environment....247
Updating vSnap servers....248
Updating the operating system for a physical vSnap server....248
Updating the operating system for a virtual vSnap server....249
Updating a vSnap server.... 250
Additional steps for updating virtual machines in Hyper-V Replica environments.... 251
Updating VADP proxies....251
Applying early availability updates.... 253

Chapter 8. Getting off to a quick start....255
Start IBM Spectrum Protect Plus....257
Manage sites....258
Create backup policies.... 259
Create a user account for the application administrator....262
Back up resources.... 264
Add resources to protect.... 264
Add resources to a job definition....266
Back up the IBM Spectrum Protect Plus catalog....268
Start a backup job....268
Run a report.... 269

Chapter 9. Configuring the system environment....271
Managing secondary backup storage....271
Managing cloud storage.... 271
Managing repository server storage.... 278
Managing sites.... 291
Adding a site....291
Editing a site.... 292
Deleting a site....293
Managing LDAP and SMTP servers....294
Adding an LDAP server....294
Adding an SMTP server.... 295
Editing settings for an LDAP or SMTP server....296
Deleting an LDAP or SMTP server....297
Managing keys and certificates for connection to IBM Spectrum Protect Plus components.... 297
Adding an access key....297
Deleting an access key....298
Adding a certificate.... 298
Deleting a certificate.... 299
Adding an SSH key.... 299
Deleting an SSH key.... 301
Managing certificates for connection to the IBM Spectrum Protect Plus user interface.... 301
Uploading a TLS certificate.... 302
Testing network connectivity.... 302
Running the Service Tool from a command line....303
Running the Service Tool remotely....304
Configuring global preferences.... 304

Configuring for virtual appliance installations ... 312
Logging on to the administrative console ... 312
Logging on to the virtual appliance ... 313
Setting the time zone ... 313
Adding virtual disks ... 314
Resetting the serveradmin password ... 318

Chapter 10. Managing SLA policies for backup operations ... 321
Protection Summary ... 321
Creating an SLA policy for hypervisors, databases, and file systems ... 324
Creating an SLA policy for Amazon EC2 instances ... 328
Creating an SLA policy for containers ... 329
Editing an SLA policy ... 333
Deleting an SLA policy ... 334

Chapter 11. Protecting virtualized systems ... 335
VMware ... 335
Adding a vCenter Server instance ... 335
Backing up VMware data ... 340
Managing VADP backup proxies ... 345
Restoring VMware data ... 350
Hyper-V ... 360
Adding a Hyper-V server ... 360
Backing up Hyper-V data ... 363
Restoring Hyper-V data ... 367
Amazon EC2 ... 373
Creating an AWS IAM user ... 373
Adding an Amazon EC2 account ... 375
Backing up Amazon EC2 data ... 376
Restoring Amazon EC2 data ... 378
Restoring files ... 380

Chapter 12. Protecting file systems ... 383
Windows file systems ... 383
Prerequisites for file systems ... 383
Adding a file system ... 384
Backing up file system data ... 388
Restoring file system data ... 395

Chapter 13. Protecting containers ... 401
Overview ... 401
Backup and restore types ... 403
SLA policies ... 405
User roles ... 406
Security features ... 407
Kubernetes ... 408
Registering a Kubernetes cluster ... 409
Backing up Kubernetes container data ... 415
Restoring Kubernetes container data ... 420
Restoring Kubernetes cluster-scoped and namespace-scoped resources ... 424
Expiring Kubernetes job sessions ... 429
OpenShift ... 430
Registering an OpenShift cluster ... 431
Backing up OpenShift container data ... 436
Restoring OpenShift container data ... 443
Expiring OpenShift job sessions ... 461
Viewing jobs and running reports ... 462

Viewing job logs....462
Creating reports for persistent volumes....462
Protecting containers by using commands....466
Container Backup Support requests....466
Backing up containers by using the command line....468
Restoring container data by using the command line....481
Restoring resources by using the command line....484
Managing container backup and restore jobs....486

Chapter 14. Protecting cloud management systems....491
Microsoft 365....491
Registering with Azure Active Directory....491
Registering the Microsoft 365 tenant with IBM Spectrum Protect Plus....492
Detailed process logs....494
Backing up Microsoft 365 data....495
Restoring Microsoft 365 data....495

Chapter 15. Protecting databases....497
Db2....497
Prerequisites for Db2....497
Adding a Db2 application server....500
Backing up Db2 data....504
Restoring Db2 data....511
Exchange Server....523
Prerequisites....523
Privileges....523
Adding an Exchange application server....525
Backing up Exchange databases....527
Incremental forever backup strategy....530
Restoring Exchange databases....530
Accessing Exchange database files with instant access mode....558
MongoDB....561
Prerequisites for MongoDB....561
Adding a MongoDB application server....564
Backing up MongoDB data....568
Restoring MongoDB data....573
Oracle....588
Adding an Oracle application server....588
Backing up Oracle data....590
Restoring Oracle data....593
SQL Server....600
Adding an SQL Server application server....601
Backing up SQL Server data....602
Restoring SQL Server data....607

Chapter 16. Protecting IBM Spectrum Protect Plus....615
Backing up the catalog....615
Restoring the catalog....615
Managing restore points....616
Expiring job sessions....616
Deleting resource metadata from the catalog....617

Chapter 17. Managing jobs and operations....619
Job types....619
Creating jobs and job schedules....620
Starting jobs on demand....621
Viewing jobs....622

Viewing backup job progress at the resource level....624
Viewing job logs....625
Viewing concurrent jobs....625
Pausing and resuming jobs....625
Editing jobs and job schedules....625
Canceling jobs....626
Deleting jobs....626
Rerunning partially completed backup jobs....627
Running an ad hoc backup job....627
Configuring scripts for backup and restore operations....628
Uploading a script....629
Adding a script to a server....629

Chapter 18. Managing reports and logs....631
Types of reports....631
Backup storage utilization reports....631
Protection reports....632
System reports....635
Running VM environment reports....636
Report actions....638
Running a report....638
Creating a custom report....639
Scheduling a report....640
Collecting and reviewing audit logs for actions....640

Chapter 19. Managing user access....643
Managing user resource groups....644
Creating a resource group....644
Editing a resource group....647
Deleting a resource group....648
Managing roles....648
Creating a role....650
Editing a role....652
Deleting a role....652
Managing user accounts....653
Creating a user account for an individual user....653
Creating a user account for an LDAP group....653
Editing user account credentials....654
Deleting a user account....654
Managing the superuser account....655
Managing identities....656
Adding an identity....656
Editing an identity....656
Deleting an identity....657

Chapter 20. Troubleshooting....659
Troubleshooting installation issues for IBM Spectrum Protect Plus as a set of containers....659
Collecting log files for troubleshooting....660
How do I tier data to tape or cloud storage?....660
How does SAN work with IBM Spectrum Protect Plus and a vSnap server?....660
Troubleshooting Container Backup Support....662
Troubleshooting installation issues....662
Installing a second instance of Velero....663
Collecting Container Backup Support log files....663
Setting the trace level of log files....664
Viewing trace logs for Container Backup Support....665
Quick reference....667

Troubleshooting backups and restores..... 672

Chapter 21. Product messages..... 679
Message prefixes..... 679

Appendix A. Search guidelines..... 681

Appendix B. Accessibility..... 683

Notices..... 685
Glossary..... 689

Index..... 691

About this publication

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

Who should read this publication

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

In this publication, it is assumed that you have an understanding of the applications that support IBM Spectrum Protect Plus as described in “System requirements ” on page 27.

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.

© Copyright IBM Corp. 2017, 2021

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.

© Copyright IBM Corp. 2017, 2021

Getting involved in product development

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

Sponsor user program

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

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

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

Are you ready to share your vision?

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

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

IBM Storage Sponsor User

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

Beta program

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

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

The program offers the following benefits:

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

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

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

Contact and enrollment information

You can enroll by completing the IBM Spectrum Protect Plus Beta Program Signup Form.

© Copyright IBM Corp. 2017, 2021

Chapter 1. IBM Spectrum Protect Plus overview

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

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

Deployment storyboard for IBM Spectrum Protect Plus

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

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

Before you start, review the system requirements for your environment. For more information, see technote 304861.

If you installed IBM Spectrum Protect Plus as a virtual appliance, see the stories in Table 1 and table Table 3.

Tip: If IBM Spectrum Protect Plus is installed as a virtual appliance, the steps in Table 1 rely on the information in the IBM Spectrum Protect Plus Blueprints and on the functioning of the *Sizer tool*. Video links are provided in Table 4 to help you with these tasks.

If you installed IBM Spectrum Protect Plus as a set of OpenShift® containers, see the stories in Table 2 and table Table 3.

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

Story	Procedure	Expected outcome
Prepare for sizing your capacity requirements by downloading the Blueprints and the Sizer Tool spreadsheet.	Download the IBM Spectrum Protect Plus Blueprints. The download includes the Sizer Tool. For sizing guidelines, see Chapters 1-3 of the Blueprints. For help with using the sizing spreadsheet, see the video links in Table 4.	You have the Sizer Tool spreadsheet and information you need to size your IBM Spectrum Protect Plus capacity requirements.

© Copyright IBM Corp. 2017, 2021

Table 1. IBM Spectrum Protect Plus installed as a virtual appliance (continued)		
Story	**Procedure**	**Expected outcome**
Size the capacity that is required for the primary storage in your environment.	Use the Sizer to size the primary storage. 1. Open the downloaded *Sizer Tool* spreadsheet and enable macros. Save a copy of the spreadsheet to your local drive for primary storage. 2. Complete the **Start Here** sheet by specifying your choices for global options for the primary storage. 3. Open the VMware tab and enter data for the vCenter capacity that includes daily rate change and annual growth. 4. Open the HyperV tab and enter data for your Hyper-V capacity. 5. For each application that you are planning to use, open an application tab and enter data for your capacity needs. 6. When all the data is entered, click the **Sizing Results** tab to review the calculated results. 7. Set the preferred vSnap server size. To automatically specify the value for the vSnap storage pool size, click `Automatic`. 8. Enter the percentage of vSnap server reserve that you require. This reserve is the percentage of the vSnap server storage that is reserved for usage, restore operations, and for any reuse. 9. Open IBM Spectrum Protect Plus, and navigate to **System Configuration** > **Global Preferences**. Input the global preferences percentages as shown in the *Sizer Tool*. Use these percentages to set the following options: • **Target free space error (percentage)** • **Target free space warning (percentage)** 10. Review the results of the Sizer for your primary storage. Save the Sizer, but leave it open for inputting settings that are required for secondary storage.	The Sizer Tool spreadsheet helps you to calculate the sizing information for primary storage. You saved a copy of the Sizer sizing spreadsheet. If capacity requirements change, you can update the spreadsheet accordingly. You also have details about the required number and size of the vSnap servers and, optionally, the number of required VMware vStorage API for Data Protection proxies. You have details about an eight-year view of growth based on your input into the spreadsheet. You set global preferences for triggering warning and errors from the vSnap when it reaches a specified threshold based on percentage usage.

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

Story	Procedure	Expected outcome
Size the capacity that is required for the secondary storage in your environment.	Use the Sizer to size the secondary storage by following these steps. Refer to Chapter 5 of the Blueprints. 1. Download the sizing spreadsheet from the Blueprints page and enable macros. Save a copy of the Sizer sheet to your local drive for secondary storage. 2. If there are any values, reset the *Sizer Tool* spreadsheet by clicking **Click to reset**. 3. Complete the **Start Here** sheet by specifying your choices for global options for the secondary storage. 4. Go to the **Results** tab of the primary storage *Sizer Tool* spreadsheet you previously saved. Copy the results that are listed in the Replication workload table and enter the values into the Optional Replication Input Workload table on the **Start Here** tab of the secondary storage Sizer Tool spreadsheet. 5. If you plan to protect application data, complete the application tabs. For example, you can specify options for copying data to object storage and replication policies. 6. Review the sizing results for your secondary storage. Save and close both Sizer Tool spreadsheets.	You have the sizing for the capacity for the secondary storage for your IBM Spectrum Protect Plus environment. You saved a copy of the Sizer for the secondary storage in your environment. If anything changes, you can alter the Sizer and make changes as required. You also have details about the vSnap server quantity for each year, the VADP proxy quantity, and the size of each vSnap server. You have details of an eight-year view of growth based on your inputs into the sizer. You set global preferences for triggering warning and errors from the vSnap when it reaches a percentage of usage.
Install or upgrade IBM Spectrum Protect Plus by using the ISO image for the version that you require. If you update the system environment, a new kernel is installed, and a restart is required.	Install IBM Spectrum Protect Plus, follow the instructions in “Installing IBM Spectrum Protect Plus as a VMware virtual appliance” on page 126 or “Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance” on page 128.	IBM Spectrum Protect Plus is installed.
Install or upgrade the vSnap server by using the ISO image for the version that you require. If you are using data deduplication, the vSnap server restart can take up to 15 minutes.	Install the vSnap server, follow the instructions in “Installing a physical vSnap server” on page 153. If you are installing a virtual vSnap server, follow the instructions in “Installing a virtual vSnap server in a Hyper-V environment” on page 155.	The vSnap server is installed. To verify that the vSnap server is installed, run the `vsnap show` command.

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

Story	Procedure	Expected outcome
Build the vSnap server with capacity that you derived from sizing by using the Blueprints and the *Sizing Tool*.	1. Create volumes and map vSnap devices. 2. Map volumes to VM cluster. 3. Refer to the steps for setting up a virtual or physical vSnap server in the Blueprints.	The vSnap server is built.
Add log space.	Create a Linux® Multiple Device driver with three partitions to store the vSnap server storage cache, cloud cache, and log files. For the cloud cache, the capacity is set at 128 GB by default. If you plan to copy data to the cloud, you must increase the capacity. For physical vSnap servers copy data to cloud storage, you must create the `/opt/vsnap-data` file system with the required capacity. For more information about this step, see *Configuring a physical vSnap server using storage software provided RAID*, and *Chapter 7 Configuring Cloud Object Storage* in the Blueprints.	You have set up log space for your virtual or physical vSnap servers.

Table 2. IBM Spectrum Protect Plus installed as a set of OpenShift containers

Story	Procedure	Expected outcome
Prepare for installation.	See the container requirements in IBM Spectrum Protect Plus as a set of containers requirements.	You are ready to install IBM Spectrum Protect Plus.
Install or upgrade IBM Spectrum Protect Plus by using one of the following methods: • Install the product from an online environment. • Download the product package and install the product in an airgap environment.	Install IBM Spectrum Protect Plus, follow the instructions in “Installing IBM Spectrum Protect Plus on OpenShift Container Platform” on page 133.	IBM Spectrum Protect Plus is installed.

Table 2. IBM Spectrum Protect Plus installed as a set of OpenShift containers (continued)		
Story	**Procedure**	**Expected outcome**
Install or upgrade the vSnap server by using the ISO image for the version that you require. If you are using data deduplication, the vSnap server restart can take up to 15 minutes.	Install the vSnap server, follow the instructions in "Installing a physical vSnap server" on page 153. If you are installing a virtual vSnap server, follow the instructions in "Installing a virtual vSnap server in a Hyper-V environment" on page 155.	The vSnap server is installed. To verify that the vSnap server is installed, run the `vsnap show` command.

Table 3. IBM Spectrum Protect Plus installed as a virtual appliance or as a set of OpenShift containers		
Story	**Procedure**	**Expected outcome**
Complete post installation tasks.	After you install IBM Spectrum Protect Plus, complete post-installation configuration tasks before you complete system management tasks. For more information and steps, see "Post installation tasks" on page 121.	You are ready to complete system management tasks to configure your IBM Spectrum Protect Plus environment.
Register the vSnap server.	Register the vSnap server. For more information and steps, see "Registering a vSnap server as a backup storage provider" on page 158.	The vSnap server is registered and added to IBM Spectrum Protect Plus.
Initialize the vSnap server.	After you install or upgrade IBM Spectrum Protect Plus, and added vSnap servers, initialize the vSnap servers. For information and steps, see "Completing a simple initialization" on page 170.	Depending on your choice, the vSnap server is initialized with or without encryption.
Configure the vSnap server.	Configure vSnap server storage options such as adding replication partners, see "Configuring backup storage options" on page 160.	If you configured the data replication feature, replication partners are set up.
(Optional) Configure the vSnap server as a VADP proxy.	If you are using a VADP proxy to optimize data movement to and from the vSnap server, you must register the vSnap server as a VADP proxy. For more instructions, see "Registering a VADP proxy on a vSnap server" on page 347.	The vSnap server is configured as a VADP proxy.
Set up the VMware environment that includes creating a vCenter, and registering a hypervisor.	To protect VMware data, you must first set up a vCenter Server. For instructions, see "Backing up and restoring VMware data" on page 335. Ensure that the required vCenter Server privileges are enabled. For more information about the required privileges, see "Virtual machine privileges " on page 336.	A vCenter is set up with the required permissions so that you can start to protect VMware data.
Add users.	Add the users who will be required to use IBM Spectrum Protect Plus. For more information, see "Creating a user account for an individual user" on page 653 by using the Add User form on the page.	The users are added and granted permissions to operate IBM Spectrum Protect Plus.

Table 3. IBM Spectrum Protect Plus installed as a virtual appliance or as a set of OpenShift containers (continued)		
Story	**Procedure**	**Expected outcome**
Create a service level agreement (SLA) policy.	Set up an SLA policy or policies for your IBM Spectrum Protect Plus workloads. For more information about SLA policies, see Chapter 10, "Managing SLA policies for backup operations," on page 321.	The SLA policies for your IBM Spectrum Protect Plus workloads are set up and you are ready to run backup jobs.
Update global preferences.	Administrators can edit the global preferences for all operations such as deduplication or encryption. For more information about global preferences, see "Configuring global preferences" on page 304.	If global preferences are set, they apply to the entire IBM Spectrum Protect Plus environment.

Resources and video library

The blueprints must be used for sizing your IBM Spectrum Protect Plus environment. The videos that are listed in the following table can help you with that process.

Table 4. Blueprints and sizing	
Task or topic	**Video link**
Introduction to the Sizer tool	IBM Spectrum Protect Plus Sizer and Blueprints: 1. Sizer introduction - Demo
Sizer worksheet overview	IBM Spectrum Protect Plus Sizer & Blueprints: 2. Sizer Worksheet Overview – Demo
Sizer Global values	IBM Spectrum Protect Plus Sizer & Blueprints: 3. Sizer Global Values – Demo
Adding a hypervisor	IBM Spectrum Protect Plus Sizer & Blueprints: 4. Adding a Hypervisor workload to the sizer – Demo
Adding an application	IBM Spectrum Protect Plus Sizer & Blueprints: 5. Adding Application workload to the sizer– Demo
Evaluating the results	IBM Spectrum Protect Plus Sizer & Blueprints: 6. Evaluating the sizer's results – Demo
Adding secondary storage	IBM Spectrum Protect Plus Sizer & Blueprints: 7. Adding a secondary site to sizer – Demo
What if scenarios	IBM Spectrum Protect Plus Sizer & Blueprints: 8. What if sizing scenarios – Demo
What's new in the blueprints	IBM Spectrum Protect Plus Sizer & Blueprints: 9. What's new in 10.1.5 sizer – Presentation
Using the Sizer results for deployment	IBM Spectrum Protect Plus Sizer & Blueprint: 10. Tying the blueprints, sizer and install together - Demo

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.

User interfaces

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

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

A key component of the interface is the dashboard, which provides summary information about the health of your environment. For more information about the dashboard, see “Product dashboard” on page 10.

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

Item	Description
IBM Spectrum Protect icon 	This icon opens IBM Spectrum Protect Operations Center to provide expanded data protection. This icon is active only when the URL is entered in the **IBM Spectrum Protect Operations Center URL** preference field on the **Global Preferences** page. For information about this preference, see “Configuring global preferences” on page 304.
Alerts icon 	This icon opens the **Alerts** window. For more information about alerts, see “Alerts” on page 11.

Item	Description
Help icon	This icon opens the online help system.
User menu	This menu shows the name of the user who is logged on. The menu provides access to product information and what's new, quick start, and API documentation. You can also use this menu to complete tasks such as accessing logs and testing connections between IBM Spectrum Protect Plus and nodes. If you are logged on to IBM Spectrum Protect Plus as the superuser, you also use this menu to manage Transport Layer Security (TLS) certificates and the product license. If IBM Spectrum Protect Plus is installed as a set of OpenShift containers, you can use this menu to update IBM Spectrum Protect Plus. The IBM Spectrum Protect Plus superuser is the user who is assigned the SUPERUSER role. There is only one IBM Spectrum Protect Plus superuser. For more information about the superuser, see “Managing the superuser account” on page 655.

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

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

vSnap command-line interface

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

Administrative console

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

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

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

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

Example 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.

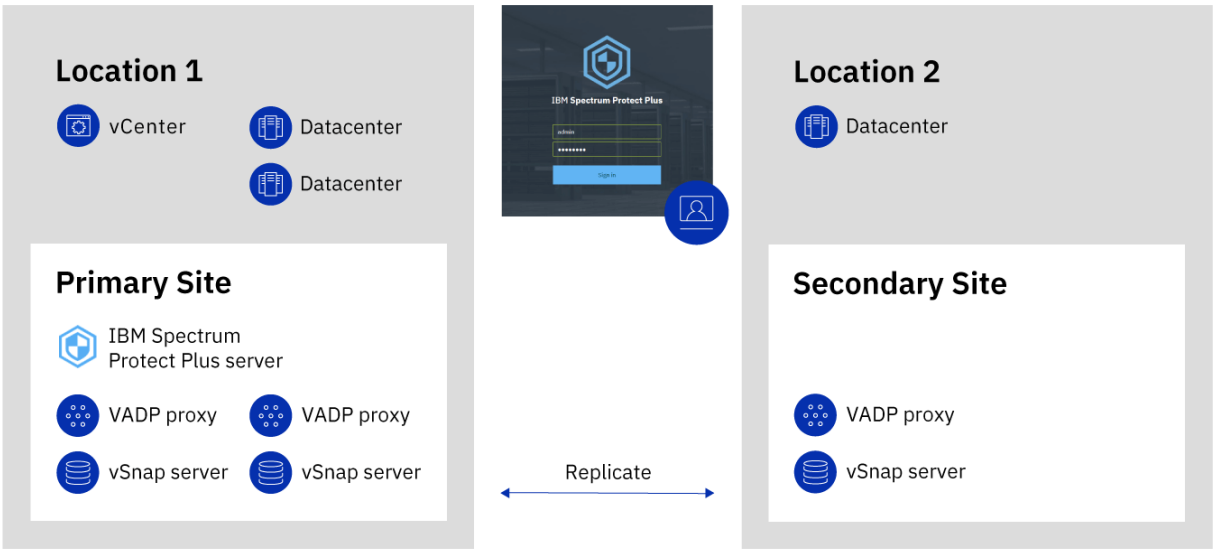


Figure 1. IBM Spectrum Protect Plus deployment across two geographical locations

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 "Resetting the serveradmin password" on page 318.

Product dashboard

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

Jobs and Operations

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

Currently Running

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

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

History

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

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

100 x Successful Jobs / Total Jobs = Success Rate

Completed jobs are shown by job status:

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

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

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

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

Destinations

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

Capacity Summary

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

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

Device Status

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

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

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

Data Reduction

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

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

Coverage

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

Source Protection

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

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

Policies

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

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

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

Alerts

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

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

The alert list can include the following alert types:

Alert types

Job failed
Is displayed when a job fails.

Job partially succeeded
Is displayed when a job partially succeeds.

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

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

System memory low
Is displayed when memory usage exceeds 95%.

System CPU usage high
Is displayed when processor usage exceeds 95%.

Hypervisor VM not found
Is displayed when the VM is not found.

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

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

SQL log backup failure
Is displayed when log backup fails for a database.

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

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

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

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

Role-based access control

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

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

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

To set up and manage user access, see Chapter 19, “Managing user access,” on page 643.

Creating a superuser account with the SUPERUSER role

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

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

For the steps required to set the username and password for the superuser account, see “Start IBM Spectrum Protect Plus” on page 257.

To manage the superuser account after it is created, see “Managing the superuser account” on page 655.

Incremental forever backup strategy

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

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

The incremental forever backup solution provides the following advantages:

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

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

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

VMware vCenter server instances
Back up resources such as virtual machines (VMs), datastores, folders, vApps, and datacenters. For more information about VMware backup operations, see “Backing up VMware data” on page 340.

Microsoft Hyper-V server instances
Back up resources such as VMs, virtual hard disks (VHDX), and VHDX files. For more information about Hyper-V backup operations, see “Backing up Hyper-V data” on page 363.

File systems
You can back up the directories and files that are associated with physical and virtualized Windows file system servers. For more information about Windows file system backup operations, see “Windows file systems ” on page 383

Cloud management systems
You can back up resource that are associated with Microsoft 365 accounts. These resources include mailboxes, calendars, contacts, and OneDrive. For more information about Microsoft 365 backup operations, see “Backing up Microsoft 365 data” on page 495.

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

Db2®
For more information about Db2 backup operations, see “Backing up Db2 data” on page 504.

Oracle
For more information about Oracle backup operations, see “Backing up Oracle data” on page 590.

MongoDB
For more information about MongoDB operations, see “MongoDB ” on page 561.

Microsoft Exchange Server
For more information about Exchange Server backup operations, see “Backing up Exchange databases” on page 527.

Microsoft SQL Server
For more information about SQL Server backup operations, see “Backing up SQL Server data” on page 602.

Backup for Amazon EC2 and containers

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

Amazon EC2 instances
For information about the backup configuration for Amazon EC2 instances, see "Backing up and restoring Amazon EC2 data" on page 373.

Kubernetes and Red Hat® OpenShift containers
For information about the backup configuration options for containers, see the "Backup and restore types" on page 403 topic for containers.

Replicate backup-storage data

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

Enabling replication of backup-storage data

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

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

   Ensure that the pool on the partner server is sufficiently large enough to hold replicated data from the primary server's pool.
2. Enable replication of backup-storage data. The replication feature is enabled by using backup policies, which are also referred to as service level agreement (SLA) policies.

   These policies define parameters that are applied to backup jobs, including the frequency of backup operations and the retention policy for the backups. For more information about SLA policies, see Chapter 10, "Managing SLA policies for backup operations," on page 321.

   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 660. 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 278.

For information about how snapshot data is copied to standard object storage and archive object storage for each cloud storage system, see "Cloud storage requirements" on page 54.

Adding secondary backup storage and creating backup policies

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

Action	How to
To copy snapshots to a repository server • Set up IBM Spectrum Protect Plus as an object client in the IBM Spectrum Protect server environment. • Add the storage to IBM Spectrum Protect Plus.	See “Configuration for copying or archiving data to IBM Spectrum Protect” on page 278 and “Registering a repository server as a backup storage provider” on page 289.
To copy snapshots to cloud storage, add the storage to IBM Spectrum Protect Plus.	Follow the instructions for your selected storage type: • “Adding Amazon S3 Object Storage” on page 272 • “Adding IBM Cloud Object Storage as a backup storage provider” on page 273 • “Adding Microsoft Azure cloud storage as a backup storage provider” on page 275 • “Registering a repository server as a backup storage provider” on page 289
Create a backup policy that includes the storage.	See “Create backup policies” on page 259.

Example deployments

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

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

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

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

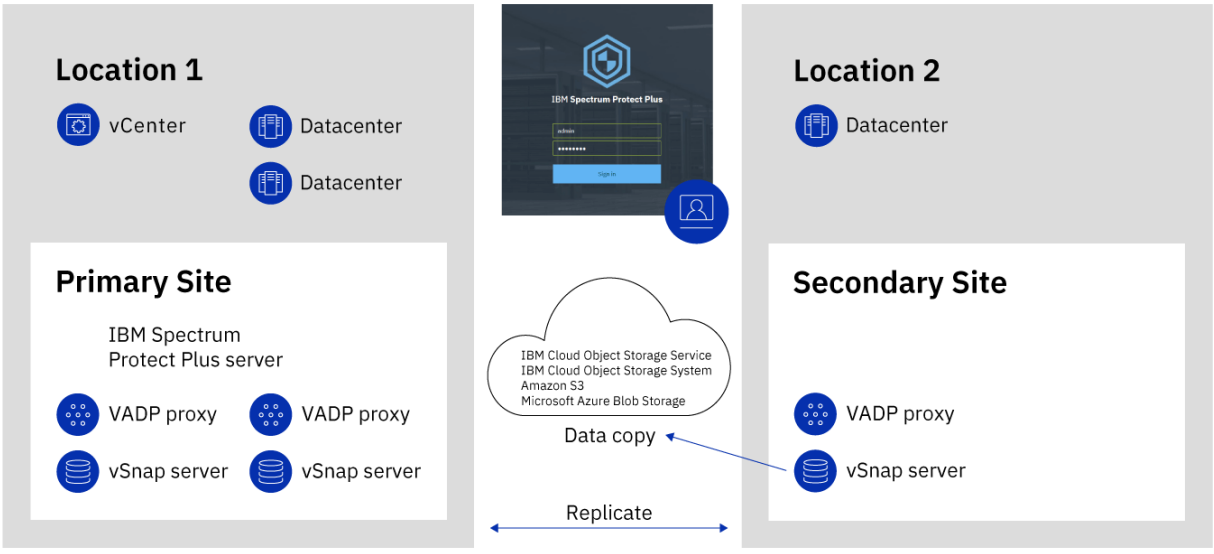


Figure 2. IBM Spectrum Protect Plus deployment across two geographical locations with copy to cloud storage

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

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

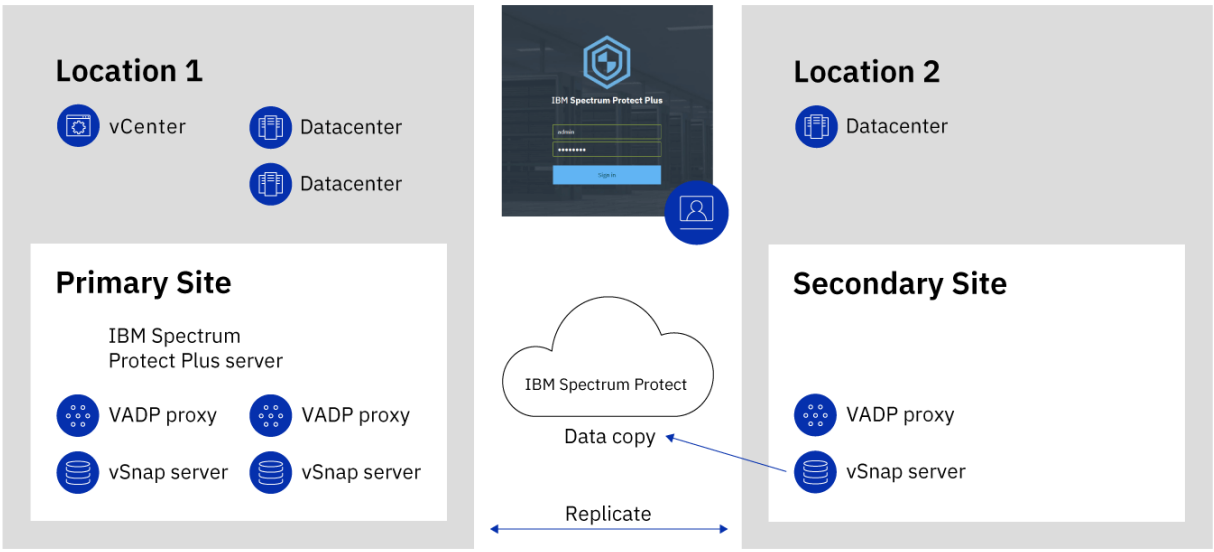


Figure 3. IBM Spectrum Protect Plus deployment across two geographical locations with copy to IBM Spectrum Protect

IBM Spectrum Protect Plus on IBM Cloud

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

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

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

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

If you want to use the current version of IBM Spectrum Protect Plus on IBM Cloud, follow the instructions in Chapter 7, “Updating IBM Spectrum Protect Plus components,” on page 239 to complete an upgrade.

IBM Spectrum Protect Plus as a software offering on IBM Cloud

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

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

All-on-cloud environment

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

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

Hybrid environment

In this configuration, only the vSnap server is deployed in IBM Cloud on an existing VMware environment. The IBM Spectrum Protect Plus server is installed and maintained on-premises. This option might benefit existing IBM Spectrum Protect Plus users who want to continue protecting

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

Deploying IBM Spectrum Protect Plus on IBM Cloud

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

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

For videos and detailed information about IBM Cloud, see the IBM Cloud documentation.

IBM Spectrum Protect Plus as a VMware service on IBM Cloud

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

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

IBM Cloud for VMware Solutions provides the following major benefits:

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

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

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

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

Expansion and contraction flexibility
Expand and contract your VMware workloads according to your business requirements.

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

For more information

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

IBM Spectrum Protect Plus on the AWS cloud platform

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

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

All-on-cloud environment

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

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

Hybrid environment

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

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

Deploying IBM Spectrum Protect Plus to AWS

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

The IBM Spectrum Protect Plus on AWS deployment includes IBM Spectrum Protect Plus version 10.1.6. If you want to use the current version of IBM Spectrum Protect Plus, follow the instructions in Chapter 7, “Updating IBM Spectrum Protect Plus components,” on page 239 to complete an upgrade.

IBM Spectrum Protect Plus on the Microsoft Azure cloud platform

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

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

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

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

All-on-cloud environment

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

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

Hybrid environment

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

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

Deploying IBM Spectrum Protect Plus on Microsoft Azure

Follow the instructions in the IBM Spectrum Protect Plus on Microsoft Azure Deployment Guide to deploy IBM Spectrum Protect Plus on Azure.

Upgrading IBM Spectrum Protect Plus to a later version

The IBM Spectrum Protect Plus on Azure deployment includes IBM Spectrum Protect Plus version 10.1.7. If you want to use the current version of IBM Spectrum Protect Plus, follow the instructions in "Updating IBM Spectrum Protect Plus in a cloud environment" on page 247 to complete an upgrade.

Integration with IBM Spectrum Protect

You can monitor your IBM Spectrum Protect Plus environment from IBM Spectrum Protect Operations Center. For convenience, you can also access the Operations Center directly from IBM Spectrum Protect Plus.

Monitor IBM Spectrum Protect Plus from the Operations Center

The Operations Center includes a dashboard for IBM Spectrum Protect Plus that provides the following information:

- A summary of job activities for a selected time period. You can view the percentages of backup, restore, and other jobs that succeeded and failed. From this summary information, you can go to more detailed information for each job type.
- A summary of the capacity and availability of vSnap servers. You can view the total disk capacity that is available to the IBM Spectrum Protect Plus server through all vSnap servers. You can also view the available capacity for each vSnap server.
- A summary of service level agreement (SLA) policies that are defined on the IBM Spectrum Protect Plus server. You can view the number of policies that have associated backup jobs. You can also view the percentage of resources that are protected by backup jobs, and the number of resources that are not protected. From this summary information, you can go to more detailed policy information.

To enable this feature, a system administrator must add the IBM Spectrum Protect Plus server to the Operations Center.

Access the Operations Center from the IBM Spectrum Protect Plus GUI

To access the Operations Center from IBM Spectrum Protect Plus, a system administrator must add the Operations Center URL on the **Global Preferences** page of the IBM Spectrum Protect Plus GUI.

You can then access the Operations Center from the IBM Spectrum Protect icon on the menu bar.

Adding IBM Spectrum Protect Plus to the Operations Center

When you add an IBM Spectrum Protect Plus server to the Operations Center, you establish a connection between the server and the Operations Center. After this connection is established, you can use the Operations Center to monitor the IBM Spectrum Protect Plus environment.

Before you begin

Ensure that you have the URL for the Operations Center and user credentials to log on.

Procedure

To add an IBM Spectrum Protect Plus server to the Operations Center, complete the following steps:

1. On the Operations Center menu bar, click **Overviews** > **Protect Plus** and take one of the following actions to open the **Add Server** wizard:

Current configuration	Action
No IBM Spectrum Protect Plus servers are connected to the Operations Center.	A message indicates that no IBM Spectrum Protect Plus servers are configured. Click **+Add Server**.
One or more IBM Spectrum Protect Plus servers are connected to the Operations Center.	The IBM Spectrum Protect Plus dashboard is displayed. From the list of servers on the monitoring dashboard, select **+Add Server**  .

2. To add the IBM Spectrum Protect Plus server, follow the directions in the wizard.

   On the **Authorization** page of the wizard, you are prompted to specify user credentials to access and monitor the IBM Spectrum Protect Plus server. If you have an IBM Spectrum Protect Plus account whose credentials match the Operations Center credentials, you can use that account. If you don't have matching credentials, you must create an account.

   Use Operations Center credentials
   Select this option to use an existing IBM Spectrum Protect Plus user account that matches the user name and password of the administrator account that you used to log on to the Operations Center.

   Create a monitoring user account
   Select this option to have the wizard create an IBM Spectrum Protect Plus user account.

   To enable the Operations Center to access IBM Spectrum Protect Plus and create the account, provide credentials for an IBM Spectrum Protect Plus user account that is assigned to the SYSADMIN role. Enter the credentials in the **User name** and **Password** fields as shown in the following figure.

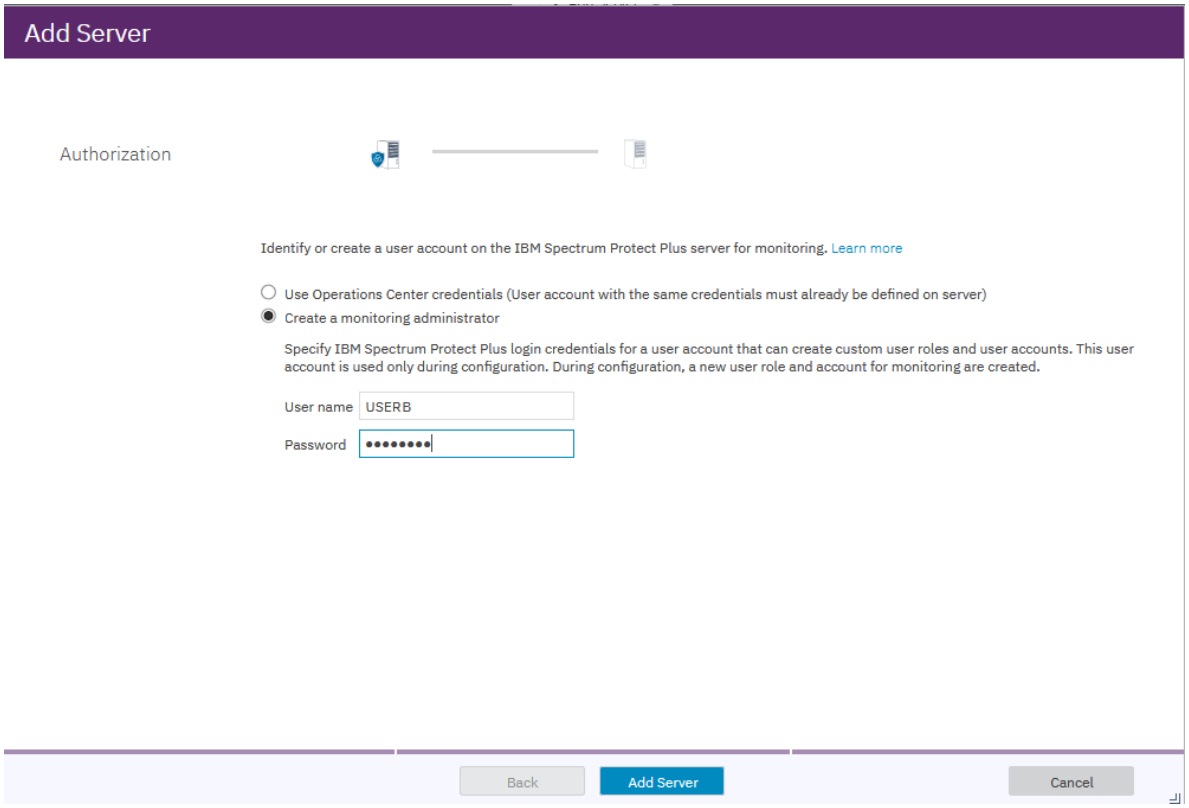


Figure 4. Entering IBM Spectrum Protect Plus credentials

The credentials that are entered here are not saved. The Operations Center logs on to the IBM Spectrum Protect Plus server by using these account credentials and creates the user account OC_MONITOR_*number*, where *number* is a random number for identification. The Operations Center will connect to the IBM Spectrum Protect Plus environment by using the new account.

3. Click **Add Server**.

   If the operation is successful, results are displayed as shown in the following figure:

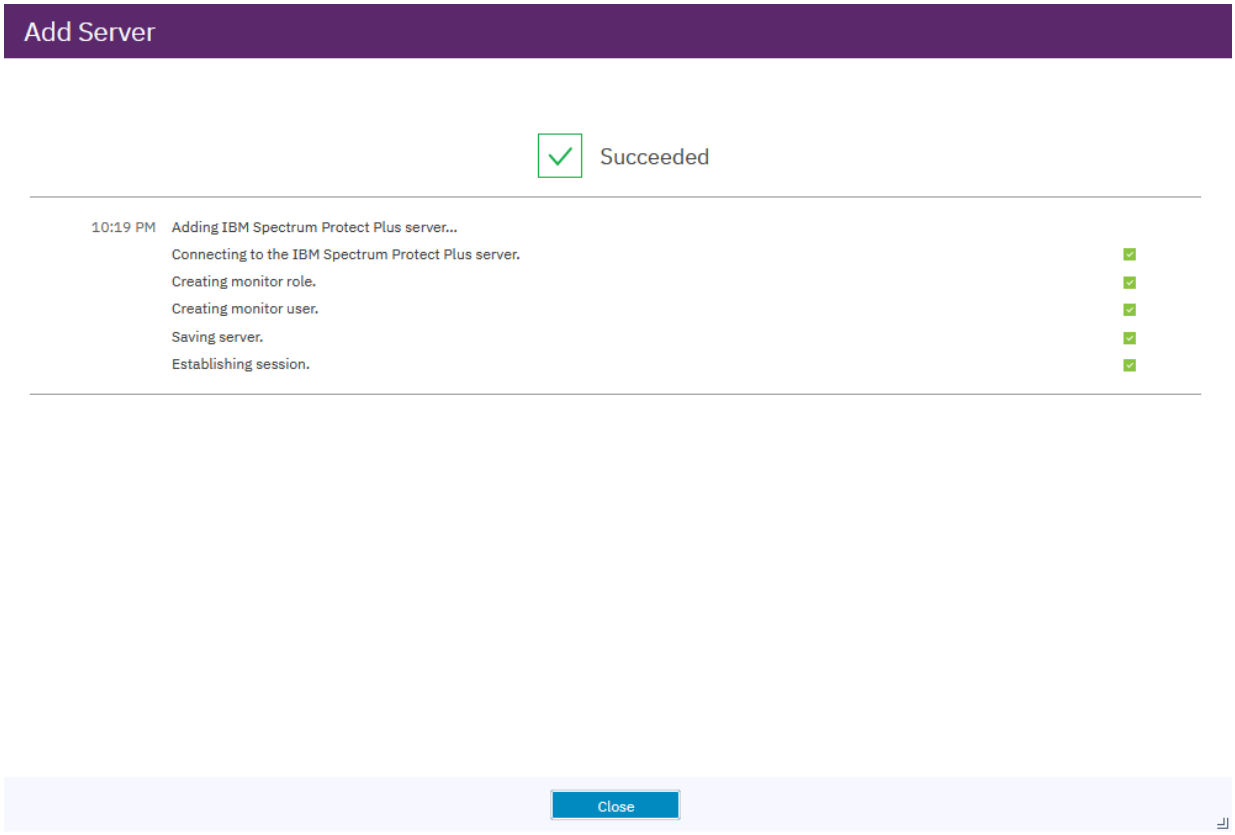


Figure 5. IBM Spectrum Protect Plus added successfully

Entering the Operations Center URL

To access the Operations Center from IBM Spectrum Protect Plus, enter the URL for the Operations Center in the IBM Spectrum Protect Plus global preferences.

About this task

You must have IBM Spectrum Protect Plus administrator credentials to configure global preferences.

When this preference is entered, the IBM Spectrum Protect icon is active on the IBM Spectrum Protect Plus menu bar.

Procedure

To enter the URL for the Operations Center, complete the following steps:

1. In the navigation panel, click **System Configuration** > **Global Preferences**.
2. Enter the URL for the Operations Center in the **IBM Spectrum Protect Operations Center URL** field.

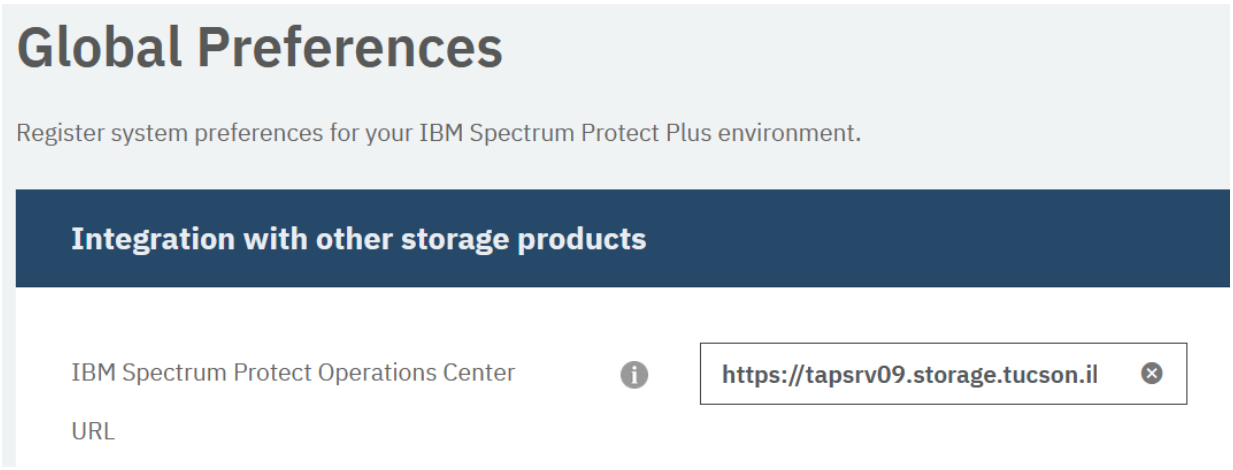


Figure 6. Entering the Operations Center URL

3. To activate the IBM Spectrum Protect icon on the IBM Spectrum Protect Plus menu bar, log off IBM Spectrum Protect Plus and log back on again.

Accessing the Operations Center

Start the Operations Center to monitor your IBM Spectrum Protect Plus environment.

Before you begin

Ensure that you completed the following tasks:

- "Adding IBM Spectrum Protect Plus to the Operations Center" on page 20
- "Entering the Operations Center URL" on page 23

Procedure

To access the Operations Center and monitor your IBM Spectrum Protect Plus environment, complete the following steps:

1. On the IBM Spectrum Protect Plus menu bar, click the IBM Spectrum Protect icon .
2. Log on to the Operations Center.
3. On the Operations Center menu bar, click **Overviews** > **Protect Plus**.

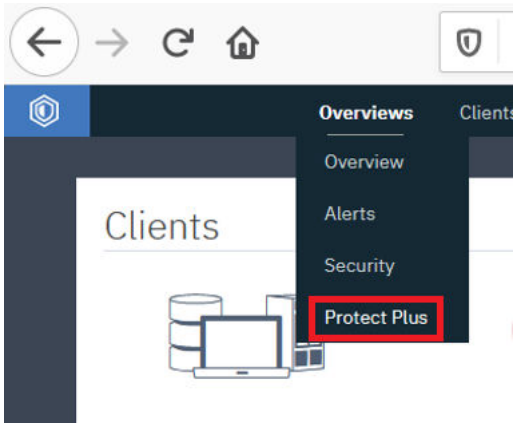


Figure 7. Selecting IBM Spectrum Protect Plus in the Operations Center

4. View the status of your IBM Spectrum Protect Plus environment on the IBM Spectrum Protect Plus monitoring dashboard as shown in the following example figure:

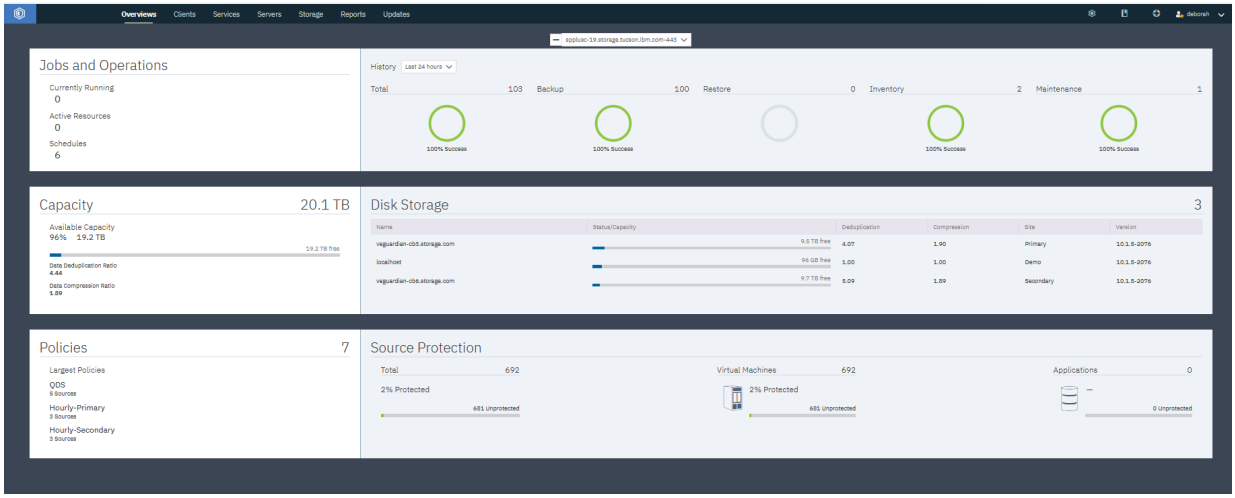


Figure 8. Viewing the IBM Spectrum Protect Plus dashboard

Integration with IBM Cloud Pak for Multicloud Management

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

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

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

For more information about IBM Cloud Pak for Multicloud Management, see the product information and demo on the IBM Cloud Pak for Multicloud Management product page.

Architecture

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

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

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

For instructions about installing the IBM Cloud Pak for Multicloud Management operator for IBM Spectrum Protect Plus, see Chapter 4, “Installing IBM Spectrum Protect Plus in a container environment,” on page 131.

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

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

Chapter 2. Installation overview

You can install IBM Spectrum Protect Plus as a VMware or Microsoft Hyper-V virtual appliance or as a set of Red Hat OpenShift containers.

Installing as a virtual appliance

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

For more information about installing IBM Spectrum Protect Plus as as a virtual appliance, see "Overview of IBM Spectrum Protect Plus virtual appliance deployment" on page 125.

Installing as a set of containers

IBM Spectrum Protect Plus can be installed on any supported cloud computing system or virtual environment where an OpenShift cluster is installed. The installation process uses the IBM Spectrum Protect Plus operator, which deploys and manages all the IBM Spectrum Protect Plus components on Kubernetes.

For more information about installing IBM Spectrum Protect Plus as set of containers, see "Overview of IBM Spectrum Protect Plus container deployment" on page 131.

Installation prerequisites

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

- IBM Spectrum Protect Plus Blueprints
- "Deployment storyboard for IBM Spectrum Protect Plus" on page 1
- "System requirements " on page 27

System requirements

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

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861.

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

Component requirements

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

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861

IBM Spectrum Protect Plus support for third-party platforms, applications, services, and hardware depend on the third-party vendors. When a third-party vendor product or version enters extended

© Copyright IBM Corp. 2017, 2021

support, self-serve support, or end of life, IBM Spectrum Protect Plus supports the product or version at the same level as the vendor.

General

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

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

IBM Spectrum Protect Plus server requirements

IBM Spectrum Protect Plus as a virtual appliance requirements

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

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

VIRTUAL APPLIANCE CONFIGURATION

Before you deploy IBM Spectrum Protect Plus to the host, ensure that one of the following virtualization products is installed on the host:

- VMware vSphere 6.0, including all updates and patch levels
- VMware vSphere 6.5, including all updates and patch levels
- vSphere 6.7, including all updates and patch levels (beginning with IBM Spectrum Protect Plus V10.1.2)
- vSphere 7.0, including all updates and patch levels (beginning with IBM Spectrum Protect Plus V10.1.6)
- Microsoft® Hyper-V 2016
- Microsoft Hyper-V 2019 (beginning with IBM Spectrum Protect Plus V10.1.3)

VIRTUAL APPLIANCE HARDWARE

The listed requirements are the minimum requirements for installation. Depending on the capacity and configuration of the storage pool, additional resources might be required. For more information about how to size and build an IBM Spectrum Protect Plus solution, see the IBM Spectrum Protect Plus Blueprints

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

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

IBM Spectrum Protect Plus as a set of containers requirements

IBM Spectrum Protect Plus can be installed on a Red Hat OpenShift cluster environment. The installation process uses the IBM Spectrum Protect Plus operator, which deploys and manages all the IBM Spectrum Protect Plus components on Red Hat OpenShift.

The IBM Spectrum Protect Plus operator is a Docker image that uses Ansible® Operator technology. The image contains the Kubernetes configuration files that are required to deploy and upgrade IBM Spectrum Protect Plus.

If you plan to install IBM Spectrum Protect Plus in an environment that has IBM Cloud Pak for Multicloud Management 2.2 installed, you must use the IBM Spectrum Protect Plus operator for IBM Cloud Pak for

Multicloud Management. This operator also works on environments that do not have the IBM Cloud Pak for Multicloud Management installed.

CONTAINER CONFIGURATION

Before you deploy IBM Spectrum Protect Plus to a Red Hat OpenShift cluster, ensure that the following requirements are met:

- **Supported container platform:** Red Hat OpenShift Container Platform Version 4.5 and later maintenance and modification levels
- **Supported cloud management platform:** IBM Cloud Pak for Multicloud Management Version 2.2 and later maintenance and modification levels
- **Supported cloud:** On premises (private cloud)

You can install the operator in an online environment or in an air-gapped environment. Before you can install an instance of the IBM Spectrum Protect Plus server, ensure that the following tools are installed or updated to the required version:

- OpenShift command-line tool (oc) is which delivered with the supported container platform version.
- Kubernetes command-line tool (kubectl) v1.18.0 or later.

  Note: Kubernetes levels supported in earlier IBM Spectrum Protect Plus reached end of life, see Kubernetes Patch Releases
- IBM Cloud Private command-line interface (cloudctl) v3.5.0 or later
- Skopeo v1.2.2 or later.

  Note: The skopeo utility is used by the load_files.sh script to copy the installation images from IBM® Entitled Registry to your private registry. For information about downloading the skopeo v1.2.2 utility, see https://github.com/containers/skopeo

You must run all commands on the Linux operating system.

IBM Spectrum Protect Plus containers are deployed on OpenShift Container Platform. IBM Spectrum Protect Plus consists of 10 core components that run as separate containers. The following IBM Spectrum Protect Plus containers are deployed in an OpenShift cluster:

- virgo
- vadp
- UI
- Node.js
- kc
- postgres
- MongoDB (three containers)
- redis
- awsebs
- awsec2

In addition to these core components, the IBM Spectrum Protect Plus operator also deploys the following containers:

- proxy: Used for internal communications between the virgo container and other containers
- manager: Used to update the IBM Spectrum Protect Plus instance from the IBM Spectrum Protect Plus user interface

For an example of system configuration, go to IBM Knowledge Center and see Figure.1

CONTAINER HARDWARE

Persistent storage:

In order for IBM Spectrum Protect Plus to run on an OpenShift cluster, persistent storage is required. The IBM Spectrum Protect Plus operator submits requests for storage by using persistent volume claims (PVCs). The OpenShift cluster completes these requests by using an existing storage driver. A storage class must be configured to allow IBM Spectrum Protect Plus to create persistent volumes dynamically. For the storage volume access mode, all PVC access modes are set to RWO (`readwriteonce`).

The following table list the minimum storage capacity for the persistent volumes (PVs):

Table 5. Persistent storage size

Persistent volume	Size	Mount Path	Permissions	Containers that access the PVC
Virgo logs	10GB	/data/log	drwxrwsr-x	virgo
plugin logs	10GB	/data/platform/log	drwxrwsr-x	awsec2 awsebs
MongoDB	50GB	/var/lib/mongodb/data	drwxrwsr-x	mongodb
MongoDB catalog	100GB	/var/lib/mongodb/data	drwxrwsr-x	mongodb2
Postgres	2GB	/var/lib/pgsql/data	drwxrwsr-x	postgres
Apache Lucene	150GB	/data/lucene	drwxrwsr-x	virgo
Nodejs logs	2GB	/data/log/node-cdm-service	drwxrwsr-x	nodejs
VMware vStorage API for Data Protection proxy (VADP proxy) logs	10GB	/data/log/vmdkbackupproxy	drwxrwsr-x	vadp

Networking:

An ingress controller on OpenShift handles external communications for IBM Spectrum Protect Plus. The IBM Spectrum Protect Plus operator deploys the ingress controller, which decrypts the encrypted traffic and directs it to the proxy container. The proxy container then routes the request internally to the proper service. Each IBM Spectrum Protect Plus container uses a corresponding Kubernetes service to communicate internally with other containers.

Timeouts:

By default, the ingress timeout is set to 900 seconds. This value can be updated by using the `haproxy.router.openshift.io/timeout` annotation of the ingress resource definition. Proxy timeouts can also be updated from the spp-proxy-config ConfigMap. The default value for the proxy timeout is set to 600 seconds.

On any external load balancers that are being used, also set the timeout values to at least 900 seconds. For example, for an OpenShift cluster on Amazon Web Services (AWS), change the default value for the `idle timeout` setting of the Elastic Load Balancing (ELB) service from 60 seconds to 900 seconds.

CPU and memory resources:

The following table lists the minimum CPU and memory resources that are required for each IBM Spectrum Protect Plus container:

Table 6. IBM Spectrum Protect Plus container CPU and memory requirements

Container	CPU (request)	CPU (limit)	Memory (request)	Memory (limit)
virgo	1000m	2000m	4Gi	8Gi
vadp	100m	250m	300Mi	500Mi

Table 6. IBM Spectrum Protect Plus container CPU and memory requirements (continued)				
Container	**CPU (request)**	**CPU (limit)**	**Memory (request)**	**Memory (limit)**
ui	50m	100m	100Mi	250Mi
nodejs	50m	100m	50Mi	150Mi
kc	50m	100m	300Mi	500Mi
postgres	50m	100m	50Mi	150Mi
mongodb (x3)	50m	150m	250Mi	2Gi
redis	100m	250m	100Mi	500Mi
awsebs	50m	250m	500Mi	2Gi
awsec2	50m	250m	500Mi	2Gi
The CPU resource is measured in Kubernetes *cpu* units. Memory is specified in units of bytes. For more information about CPU units and memory, see Managing Resources for Containers				

IBM Spectrum Protect Plus server additional requirements

The “Connectivity requirements” on page 54 must be met.

Only default Active Directory non-nested security groups are supported.

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

IBM Spectrum Protect Plus server browser support

IBM Spectrum Protect Plus was tested and validated with the following web browsers:

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

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

IBM Spectrum Protect Plus server ports

IBM Spectrum Protect Plus server and associated services use the following ports.

Table 7. Communication ports when the target is an IBM Spectrum Protect Plus server				
Port	**Protocol**	**Initiator**	**Target**	**Description**
22	Transmission Control Protocol (TCP)	vSnap server	IBM Spectrum Protect Plus server	Provides access for troubleshooting and maintenance tasks on the IBM Spectrum Protect Plus server by using the Secure Shell (SSH) protocol.

Table 7. Communication ports when the target is an IBM Spectrum Protect Plus server (continued)

Port	Protocol	Initiator	Target	Description
443	TCP	IBM Spectrum Protect Plus user interface	IBM Spectrum Protect Plus server	Provides web access by using the Hypertext Transfer Protocol Secure (HTTPS) protocol. This port is the main entry point for client connections that use the Transport Layer Security (TLS) protocol. This port is also used for Representational State Transfer application programming interface (REST API) queries.
443	TCP	VADP proxy host	IBM Spectrum Protect Plus	Provides web access by using the HTTPS protocol. This port is the main entry point for client connections that use the TLS protocol. This port is also used for REST API queries.
443	TCP	All agents(except IBM Db2	IBM Spectrum Protect Plus	Used for IBM Spectrum Protect Plus issued commands to run backup, restore, inventory, and other operations

Table 7. Communication ports when the target is an IBM Spectrum Protect Plus server (continued)

Port	Protocol	Initiator	Target	Description
8090	TCP	Kubernetes plug-in	IBM Spectrum Protect Plus server	Provides access to the built-in Kubernetes (K8s) Kube proxy in support of Amazon Elastic Compute Cloud (EC2). Not every port in this range is used at the same time. Rather, a small subset of ephemeral ports will be opened as required by NodePort services.

Table 8. Communication ports when the initiator is an IBM Spectrum Protect Plus server

Port	Protocol	Initiator	Target	Description
22	TCP	IBM Spectrum Protect Plus server	vSnap server	Provides access for troubleshooting and maintenance tasks on vSnap servers by using the SSH protocol.
22	TCP	IBM Spectrum Protect Plus	VADP proxy host	Provides access for troubleshooting and maintenance tasks on the VADP proxy hosts by using the SSH protocol.
22	TCP	IBM Spectrum Protect Plus	Agents that use the Network File System (NFS) client. Except Kubernetes and OpenShift agents	Provides access to troubleshoot and maintain remote proxy host servers running guest application components by using the SSH protocol for Kubernetes and OpenShift agents the port will assigned by NodePort service in Kubernetes

Table 8. Communication ports when the initiator is an IBM Spectrum Protect Plus server (continued)

Port	Protocol	Initiator	Target	Description
Assigned by the NodePort service in Kubernetes	TCP	IBM Spectrum Protect Plus	Kubernetes or OpenShift agent	Used by IBM Spectrum Protect Plus to connect to the data mover container to deploy and run agents By default, port 30001 is used for SSH connections to the Kubernetes or OpenShift agent containers. This port is configurable and is port-forwarded to port 22. This port is used only when a containerized IBM Spectrum Protect Plus server opens an SSH connection to the Kubernetes or OpenShift agent container. SSH connections are never used within the Container Backup Support containers.
25	TCP	IBM Spectrum Protect Plus server	Email server that can be accessed by using the Simple Mail Transfer Protocol (SMTP)	Provides access to an email service.
389	TCP	IBM Spectrum Protect Plus server	Lightweight Directory Access Protocol (LDAP) server	Provides access to Active Directory Services.
443	TCP	IBM Spectrum Protect Plus server	Hypervisor: VMware Elastic Sky X Integrated (ESXi) host and vCenter	Provides access to ESXi and vCenter for managing operations.
443	TCP	IBM Spectrum Protect Plus	Hypervisor: EC2	Provides access to Amazon Web Services (AWS) for managing operations.
636	TCP	IBM Spectrum Protect Plus server	LDAP server	Provides access to Active Directory Services by using the TLS protocol.

Table 8. Communication ports when the initiator is an IBM Spectrum Protect Plus server (continued)

Port	Protocol	Initiator	Target	Description
902	TCP	IBM Spectrum Protect Plus server	Hypervisor: VMware ESXi host	Used for the Network File Copy (NFC) protocol, which provides a file-type-aware File Transfer Protocol (FTP) service for vSphere components.By default, ESXi uses NFC for operations such as copying and moving data between datastores.
5985	TCP	IBM Spectrum Protect Plus server	Hypervisor: Hyper-V	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.
5985	TCP	IBM Spectrum Protect Plus server	Agents that use the iSCSI initiator	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.
5986	TCP	IBM Spectrum Protect Plus	Hypervisor: Hyper-V	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.
5986	TCP	IBM Spectrum Protect Plus	Agents that use the iSCSI initiator	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.
8098	TCP	IBM Spectrum Protect Plus server	VADP proxy host	Supports REST API communications between the IBM Spectrum Protect Plus server and the VADP proxy by using the TLS protocol.

Table 8. Communication ports when the initiator is an IBM Spectrum Protect Plus server (continued)				
Port	**Protocol**	**Initiator**	**Target**	**Description**
8900	TCP	IBM Spectrum Protect Plus server	vSnap server	Supports REST API communications between the IBM Spectrum Protect Plus server and the vSnap server by using the TLS protocol.

vSnap server requirements

A vSnap server is the primary backup destination for IBM Spectrum Protect Plus.

vSnap server configuration

- **vSnap server VM installation**

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

  - vSphere 6.0, including all updates and patch levels
  - vSphere 6.5, including all updates and patch levels
  - vSphere 6.7, including all updates and patch levels (beginning with IBM Spectrum Protect Plus V10.1.2)
  - vSphere 7.0, including all updates and patch levels (beginning with IBM Spectrum Protect Plus V10.1.6)
  - Microsoft Hyper-V 2016
  - Microsoft Hyper-V 2019 (beginning with IBM Spectrum Protect Plus V10.1.3)

- **vSnap server physical installation**

  When installing or upgrading vSnap on CentOS or RedHat Enterprise Linux, the installer contains additional operating system packages that the vSnap software depends on. The dependencies included in the installer are sufficient to satisfy requirements on a 'Minimal' configuration of CentOS or RedHat Enterprise Linux. If the system was created with a different OS configuration, additional dependencies may be needed. If the system is configured with access to online repositories, the vSnap installer will attempt to download the required dependencies automatically. If online repositories are not accessible, you may need to manually install or update additional packages. Refer to the on-screen messages shown by the vSnap installer to determine which additional packages need to be installed or updated.

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

- The following Linux operating systems are supported for IBM Spectrum Protect Plus V10.1.8 physical vSnap server installations:

  - CentOS 7.1804 (7.5) (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.2)
  - CentOS 7.1810 (7.6) (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.3 patch 1)
  - CentOS 7.1908 (7.7) (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.5 patch 1)
  - CentOS 7.2003 (7.8) (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.7)
  - CentOS 7.2009 (7.9) (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.8)
  - CentOS 8.1911 (8.1) (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.8)
  - CentOS 8.2004 (8.2) (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.8)
  - CentOS 8.2011 (8.3) (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.8)

- RHEL 7.5 (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.2)
- RHEL 7.6 (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.3 patch1)
- RHEL 7.7 (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.5 patch1)
- RHEL 7.8 (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.7)
- RHEL 7.9 (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.8)
- RHEL 8.1 (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.8)
- RHEL 8.2 (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.8)
- RHEL 8.3 (x86_64) (beginning with IBM Spectrum Protect Plus V10.1.8)

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

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

vSnap server hardware

The listed requirements are the minimum requirements for installation. Depending on the capacity and configuration of the storage pool, additional resources might be required. For more information about how to size and build an IBM Spectrum Protect Plus solution, see the IBM Spectrum Protect Plus Blueprints

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

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

Restrictions:

- UEFI Secure Boot must be disabled
- SELinux mode must be Permissive in the /etc/selinux/config file druring physical vSnap installation

vSnap server additional requirements

The “Connectivity requirements” on page 54 must be met.

vSnap server ports

The following ports are used by vSnap servers.

Table 9. Communication ports when the target is a vSnap server

Port	Protocol	Initiator	Target	Description
22	TCP	IBM Spectrum Protect Plus server	vSnap server	Provides access for troubleshooting and maintenance tasks on vSnap servers by using SSH protocol.

Table 9. Communication ports when the target is a vSnap server (continued)

Port	Protocol	Initiator	Target	Description
22	TCP	Hypervisors	vSnap server	Provides access for troubleshooting and maintenance tasks on vSnap servers by using SSH protocol.
22	TCP	Agents that use the NFS client	vSnap server	Provides access for troubleshooting and maintenance tasks on vSnap servers by using SSH protocol.
111	TCP and User Datagram Protocol (UDP)	Hypervisor: VMware ESXi host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
111	TCP and UDP	VADP proxy host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
445	TCP	Application agents that use the SMB or the CIFS protocol	vSnap server	Used for SMB or CIFS file sharing by the vSnap server during backup and restore operations.
2049	TCP and UDP	Hypervisor: VMware ESXi host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
2049	TCP and UDP	VADP proxy host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.

Table 9. Communication ports when the target is a vSnap server (continued)

Port	Protocol	Initiator	Target	Description
2049	TCP and UDP	Agents that use the NFS client	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
3260	TCP	Hypervisor: Microsoft Hyper-V	vSnap server	Used for Microsoft Internet Small Computer System Interface (iSCSI) data transfer to and from logical unit number's (LUN)s mounted from vSnap servers during backup and restore operations.
3260	TCP	Agents that use the iSCSI client	vSnap server	Used for iSCSI data transfer to and from LUNs mounted from vSnap servers during backup and restore operations.
8900	TCP	IBM Spectrum Protect Plus server	vSnap server	Supports REST API communications between the IBM Spectrum Protect Plus server and the vSnap server by using the TLS protocol.
8900	TCP	vSnap server	vSnap server	Supports REST API communications between two vSnap servers during replication by using the TLS protocol.
20048	TCP and UDP	Hypervisor: VMware ESXi host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.

Table 9. Communication ports when the target is a vSnap server (continued)

Port	Protocol	Initiator	Target	Description
20048	TCP and UDP	VADP proxy host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
20048	TCP and UDP	Agents that use the NFS client	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.

Important security information: Process requests to vSnap data ports (NFS, SMB, and iSCSI) only when the request comes from a node in the internal network. Requests that come from external (non-private) network nodes must be blocked. To ensure that proper security practices are followed, work with your network security administrator.

Table 10. Communication ports when the initiator is a vSnap server

Port	Protocol	Initiator	Target	Description
22	TCP	vSnap server	IBM Spectrum Protect Plus server	Provides access for troubleshooting and maintenance tasks on the IBM Spectrum Protect Plus server by using SSH protocol.
443	TCP	vSnap server	Cloud server endpoints	Allows the vSnap server to communicate with Amazon Simple Storage Service (S3), Microsoft Azure, or IBM Cloud Object Storage endpoints.
8900	TCP	vSnap server	vSnap server	Supports REST API communications between two vSnap servers during replication by using the TLS protocol.

Table 10. Communication ports when the initiator is a vSnap server (continued)				
Port	**Protocol**	**Initiator**	**Target**	**Description**
9000	TCP	vSnap server	Repository server endpoints	Allows the vSnap server to communicate with IBM Spectrum Protect (repository server) endpoints.

VADP proxy requirements

In IBM Spectrum Protect Plus, running VM backup jobs through VADP requires significant system resources. By creating VADP backup job proxies, you enable load sharing and load balancing for IBM Spectrum Protect Plus backup jobs. If proxies exist, the entire processing load is shifted from the IBM Spectrum Protect Plus server onto the proxies.

VADP proxy configuration

This feature is supported only in 64-bit quad core or higher configurations with a minimum kernel version of v2.6.32 in the following Linux environments:

- CentOS 6.5 and later maintenance and modification levels (beginning with IBM Spectrum Protect Plus V10.1.1 patch 1)
- CentOS 7.0 and later maintenance and modification levels (beginning with IBM Spectrum Protect Plus V10.1.1 patch 1)
- CentOS 8.0 and later maintenance and modification levels (beginning with IBM Spectrum Protect Plus V10.1.8)
- RHEL 6.4 and later maintenance and modification levels (beginning with IBM Spectrum Protect Plus V10.1.1)
- RHEL 7 and later maintenance and modification levels (beginning with IBM Spectrum Protect Plus V10.1.1)
- RHEL 8 and later maintenance and modification levels (beginning with IBM Spectrum Protect Plus V10.1.8)
- SUSE Linux Enterprise Server (SLES) 12 and later maintenance and modification levels (beginning with IBM Spectrum Protect Plus V10.1.1)
- SLES 15 and later maintenance and modification levels (beginning with IBM Spectrum Protect Plus V10.1.7)

For more information about how to build an IBM Spectrum Protect Plus solution, see the IBM Spectrum Protect Plus Blueprints

VADP proxy hardware

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

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

Notes:

- Because of increased processor usage and concurrency on the VADP proxy server, the memory that is allocated on the proxy server must be increased.
- Disk space is consumed in /opt/IBM/SPP for installation and the job log file. Approximately 100 MB is used for installation and the job log file size will vary based on job execution. Additionally, the /tmp directory is used for temporary files where the amount of disk space used depends on job execution.

VADP proxy additional requirements

The “Connectivity requirements” on page 54 must be met.

To create VADP proxies, you must have a user ID with the SYSADMIN role assigned. For more information about roles, see “Managing roles” on page 648.

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

VADP proxy ports

The following ports are used by VADP proxies.

Table 11. Communication ports when the target is a VADP proxy host

Port	Protocol	Initiator	Target	Description
22	TCP	IBM Spectrum Protect Plus server	VADP proxy host	Provides access for troubleshooting and maintenance tasks on VADP proxy hosts by using the SSH protocol.
8098	TCP	IBM Spectrum Protect Plus server	VADP proxy host	Supports REST API communications between the IBM Spectrum Protect Plus server and the VADP proxy by using the TLS protocol.

Notes®:

- VADP proxies can be pushed and installed to Linux-based servers over SSH port 22.
- Port 8098 on the VADP proxy server must be open when the proxy server firewall is enabled.

Table 12. Communication ports when the initiator is a VADP proxy host

Port	Protocol	Initiator	Target	Description
111	TCP and UDP	VADP proxy host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
443	TCP	VADP proxy host	Hypervisor: VMware ESXi host and vCenter	Provides access to ESXi and vCenter for managing operations.

Table 12. Communication ports when the initiator is a VADP proxy host (continued)

Port	Protocol	Initiator	Target	Description
443	TCP	VADP proxy host	IBM Spectrum Protect Plus server	Provides web access by using the HTTPS protocol. This port is the main entry point for client connections that use TLS protocol. This port is also used for REST API queries.
902	TCP	VADP proxy host	Hypervisor: VMware ESXi host	Used for the Network File Copy (NFC) protocol, which provides a file-type-aware File Transfer Protocol (FTP) service for vSphere components. By default, ESXi uses NFC for operations such as copying and moving data between datastores.
2049	TCP and UDP	VADP proxy host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
20048	TCP and UDP	VADP proxy host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.

If the firewall command script is not available on your system, edit the firewall manually to open or close the necessary ports, and restart the firewall. For instructions about editing firewall ports, see "Editing firewall ports" on page 122.

VADP proxy on vSnap server

VADP proxies can be installed on the vSnap servers in your IBM Spectrum Protect Plus environment. A combination VADP proxy and vSnap server must meet the minimum requirements of both devices. Consider the system requirements of both devices and add the core and RAM requirements together to identify the minimum requirements of the combination VADP proxy and vSnap server.

For a VADP proxy installed on a virtual vSnap server, the following requirements must be met:

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

All required VADP proxy ports and vSnap server ports must be open on the combination VADP proxy and vSnap server.

IBM Spectrum Protect Plus Ports and Diagram

IBM Spectrum Protect Plus communication ports

The following table lists the all ports between the IBM Spectrum Protect Plus components.

Table 13. Communication ports

Port	Protocol	Initiator	Target	Description
111	TCP and User Datagram Protocol (UDP)	Hypervisor: VMware ESXi host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
111	TCP and UDP	VADP proxy host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
111	TCP and UDP	Agents that use the NFS client	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
20048	TCP and UDP	Hypervisor: VMware ESXi host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
20048	TCP and UDP	VADP proxy host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.

Table 13. Communication ports (continued)

Port	Protocol	Initiator	Target	Description
20048	TCP and UDP	Agents that use the NFS client	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
2049	TCP and UDP	Hypervisor: VMware ESXi host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
2049	TCP and UDP	VADP proxy host	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
2049	TCP and UDP	Agents that use the NFS client	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations
22	Transmission Control Protocol (TCP)	vSnap server	IBM Spectrum Protect Plus server	Provides access for troubleshooting and maintenance tasks on the IBM Spectrum Protect Plus server by using Secure Shell (SSH) protocol.
22	TCP	IBM Spectrum Protect Plus server	vSnap server	Provides access for troubleshooting and maintenance tasks on vSnap servers by using the SSH protocol.
22	TCP	IBM Spectrum Protect Plus server	VADP proxy host	Provides access for troubleshooting and maintenance tasks on VADP proxy hosts by using the SSH protocol.

Table 13. Communication ports (continued)

Port	Protocol	Initiator	Target	Description
22	TCP	IBM Spectrum Protect Plus server	Agents that use the NFS client. Except Kubernetes and OpenShift agents	Provides access to troubleshoot and maintain remote proxy host servers running guest application components by using the SSH protocol **Note:** for Kubernetes and OpenShift agents the port will assigned by NodePort service in Kubernetes
22	TCP	Hypervisors	vSnap server	Provides access for troubleshooting and maintenance tasks on vSnap servers by using SSH protocol.
22	TCP	Agents that use the NFS client	vSnap server	Provides access for troubleshooting and maintenance tasks on vSnap servers by using SSH protocol.
25	TCP	IBM Spectrum Protect Plus server	Email server that can be accessed by using the Simple Mail Transfer Protocol (SMTP)	Provides access to an email service.
30000 - 32767	TCP	Kubernetes plug-in	IBM Spectrum Protect Plus server	Provides access to the built-in Kubernetes (K8s) Kube proxy in support of EC2. **Note:** Not every port in this range is used at the same time. Rather, a small subset of ephemeral ports will be opened as required by NodePort services.

Table 13. Communication ports (continued)

Port	Protocol	Initiator	Target	Description
3260	TCP	Hypervisor: Microsoft Hyper-V	vSnap server	Used for Microsoft Internet Small Computer System Interface (iSCSI) data transfer to and from logical unit number's (LUN)s s mounted from vSnap servers during backup and restore operations.
3260	TCP	Agents that use the iSCSI client	vSnap server	Used for iSCSI data transfer to and from LUNs mounted from vSnap servers during backup and restore operations.
389	TCP	IBM Spectrum Protect Plus server	Lightweight Directory Access Protocol (LDAP) server	Provides access to Active Directory Services.
443	TCP	IBM Spectrum Protect Plus user interface	IBM Spectrum Protect Plus server	Provides web access by using the Hypertext Transfer Protocol Secure (HTTPS) protocol. This port is the main entry point for client connections that use the Transport Layer Security (TLS) protocol. This port is also used for Representational State Transfer application programming interface (REST API) queries.

Table 13. Communication ports (continued)

Port	Protocol	Initiator	Target	Description
443	TCP	VADP proxy host	IBM Spectrum Protect Plus server	Provides web access by using the HTTPS protocol. This port is the main entry point for client connections that use TLS protocol. This port is also used for REST API queries.
443	TCP	IBM Spectrum Protect Plus server	Hypervisor: VMware Elastic Sky X Integrated (ESXi) host and vCenter	Provides access to ESXi and vCenter for managing operations.
443	TCP	IBM Spectrum Protect Plus server	Hypervisor: Amazon EC2	Provides access to Amazon Web Services (AWS) for managing operations.
443	TCP	VADP proxy host	Hypervisor: VMware ESXi host and vCenter	Provides access to ESXi and vCenter for managing operations.
443	TCP	All agents(except IBM Db2)	IBM Spectrum Protect Plus server	Port that allows the agents to communicate with IBM Spectrum Protect Plus for making representational state transfer application programming interface (REST API) calls to run backup, restore, inventory, and other operations.
443	TCP	vSnap server	Cloud server endpoints	Allows the vSnap server to communicate with Amazon Simple Storage Service (S3), Microsoft Azure, or IBM Cloud Object Storage endpoints.

Table 13. Communication ports (continued)

Port	Protocol	Initiator	Target	Description
445	TCP	Application agents that use the SMB or the CIFS protocol	vSnap server	Used for SMB or CIFS file sharing by the vSnap server during backup and restore operations.
5985	TCP	IBM Spectrum Protect Plus server	Hypervisor: Hyper-V	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.
5985	TCP	IBM Spectrum Protect Plus server	Agents that use the iSCSI initiator	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.
5986	TCP	IBM Spectrum Protect Plus server	Hypervisor: Hyper-V	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.
5986	TCP	IBM Spectrum Protect Plus server	Agents that use the iSCSI initiator	Provides access to the Microsoft Windows Remote Management (WinRM) service for Windows-based servers.
636	TCP	IBM Spectrum Protect Plus server	LDAP server	Provides access to Active Directory Services by using the Transport Layer Security (TLS) protocol.
8090	TCP	IBM Spectrum Protect Plus administrative console	IBM Spectrum Protect Plus server	Provides access for system administration. This extensible framework supports plug-ins that run operations such as system and network updates.

Table 13. Communication ports (continued)

Port	Protocol	Initiator	Target	Description
8098	TCP	IBM Spectrum Protect Plus server	VADP proxy host	Supports REST API communications between the IBM Spectrum Protect Plus server and the VADP proxy by using the Transport Layer Security (TLS) protocol.
8900	TCP	IBM Spectrum Protect Plus server	vSnap server	Supports REST API communications between the IBM Spectrum Protect Plus server and the vSnap server by using the TLS protocol.
8900	TCP	vSnap server	vSnap server	Supports REST API communications between two vSnap servers during replication by using the TLS protocol.
9000	TCP	vSnap server	Repository server endpoints	Allows the vSnap server to communicate with IBM Spectrum Protect (repository server) endpoints.
902	TCP	IBM Spectrum Protect Plus server	Hypervisor: VMware ESXi host	Used for the Network File Copy (NFC) protocol, which provides a file-type-aware File Transfer Protocol (FTP) service for vSphere components. By default, ESXi uses NFC for operations such as copying and moving data between datastores.

Table 13. Communication ports (continued)

Port	Protocol	Initiator	Target	Description
902	TCP	VADP proxy host	Hypervisor: VMware ESXi host	Used for the Network File Copy (NFC) protocol, which provides a file-type-aware File Transfer Protocol (FTP) service for vSphere components. By default, ESXi uses NFC for operations such as copying and moving data between datastores.
Assigned by the NodePort service in Kubernetes	TCP	IBM Spectrum Protect Plus server	Kubernetes or OpenShift agent	Used by IBM Spectrum Protect Plus to connect to the data mover container to deploy and run agents **Note:** By default, port 30001 is used for SSH connections to the Kubernetes or OpenShift agent containers. This port is configurable and is port-forwarded to port 22. This port is used only when a containerized IBM Spectrum Protect Plus server opens an SSH connection to the Kubernetes or OpenShift agent container. SSH connections are never used within the Container Backup Support containers.

Port updates:

- Ports 111, 2029, and 20048: In earlier versions, these ports were used for catalog backup operations to vSnap server via the Network File System (NFS) client. Beginning with IBM Spectrum Protect Plus V10.1.7, the IBM Spectrum Protect Plus server uses the Secure File Transfer protocol (SFTP) to back up catalogs to vSnap servers. For that reason, ports 111, 2029, and 20048 are no longer required.

- Ports 137, 138, and 139: In earlier versions, ports 137, 138, and 139 on the vSnap server were used by application agents that use SMBv1. Beginning with IBM Spectrum Protect Plus V10.1.6, the SMBv1 protocol is not used. All agents use SMBv2 or later, which does not require ports 137, 138, or 139.
- Port 3260: In earlier versions, this port was used for Internet Small Computer System Interface (iSCSI) data transfer by the vSnap server. Beginning with IBM Spectrum Protect Plus V10.1.7, the IBM Spectrum Protect Plus server does not include an onboard vSnap server. For that reason, the port is no longer required.
- Port 9090: In earlier versions, this port was used for online help. Starting with V10.1.4, this port is no longer required for online help. No further action is required.
- Port 8761: In earlier versions, this port was used to automatically discover VADP proxies and for IBM Spectrum Protect Plus VM backup operations. Beginning with IBM Spectrum Protect Plus V10.1.6, the VADP proxy architecture is modified and port 8761 is no longer required to be open. When IBM Spectrum Protect Plus is updated to V10.1.6 or later, the associated VADP proxies in the environment are also upgraded.
- Port 5671: In earlier versions, this port was used for internal and external message and log management. Beginning with IBM Spectrum Protect Plus V10.1.7, the VADP proxy architecture is modified and port 5671 is no longer required to be open.

  Note: If you upgrade to IBM Spectrum Protect Plus version 10.1.7 from a previous version, you can close TCP port 5671 since it is no longer used in V10.1.7 and later. Log in to IBM Spectrum Protect Plus as the serveradmin user and issue the following commands to close the port:

```
$ sudo firewall-cmd --zone=public --permanent --remove-port=5671/tcp
$ sudo firewall-cmd --reload
```

IBM Spectrum Protect Plus communication paths diagram

The following diagram is an overview of the communication paths that are managed by IBM Spectrum Protect Plus. This diagram is intended to provide a high level representation of components and their associated ports. This diagram can provide assistance for troubleshooting and network configuration for deployment scenarios. The diagram may not contain all communication paths. For a list of initiator and receiver ports, see the Communication ports table.

- The labeled resources on the gray background represent the core services of IBM Spectrum Protect Plus.
- The colors of the various modules represent different types of services as defined by the key.
- The area that is labeled **Firewall** represents the network firewall.
- Services that appear in the **Firewall** area are indicative of the ports that are open on the firewall.
- Dashed arrows represent communication among resources and services.
- Arrows flow toward the listening port (target).
- The port numbers that must be open are indicated by the listening port.

  For example:

  – The vSnap service is represented as being external to the IBM Spectrum Protect Plus server. The vSnap service is listening on port 8900 and other ports.

  – A component in the server establishes a communication path with a connection to the vSnap service at port 8900.

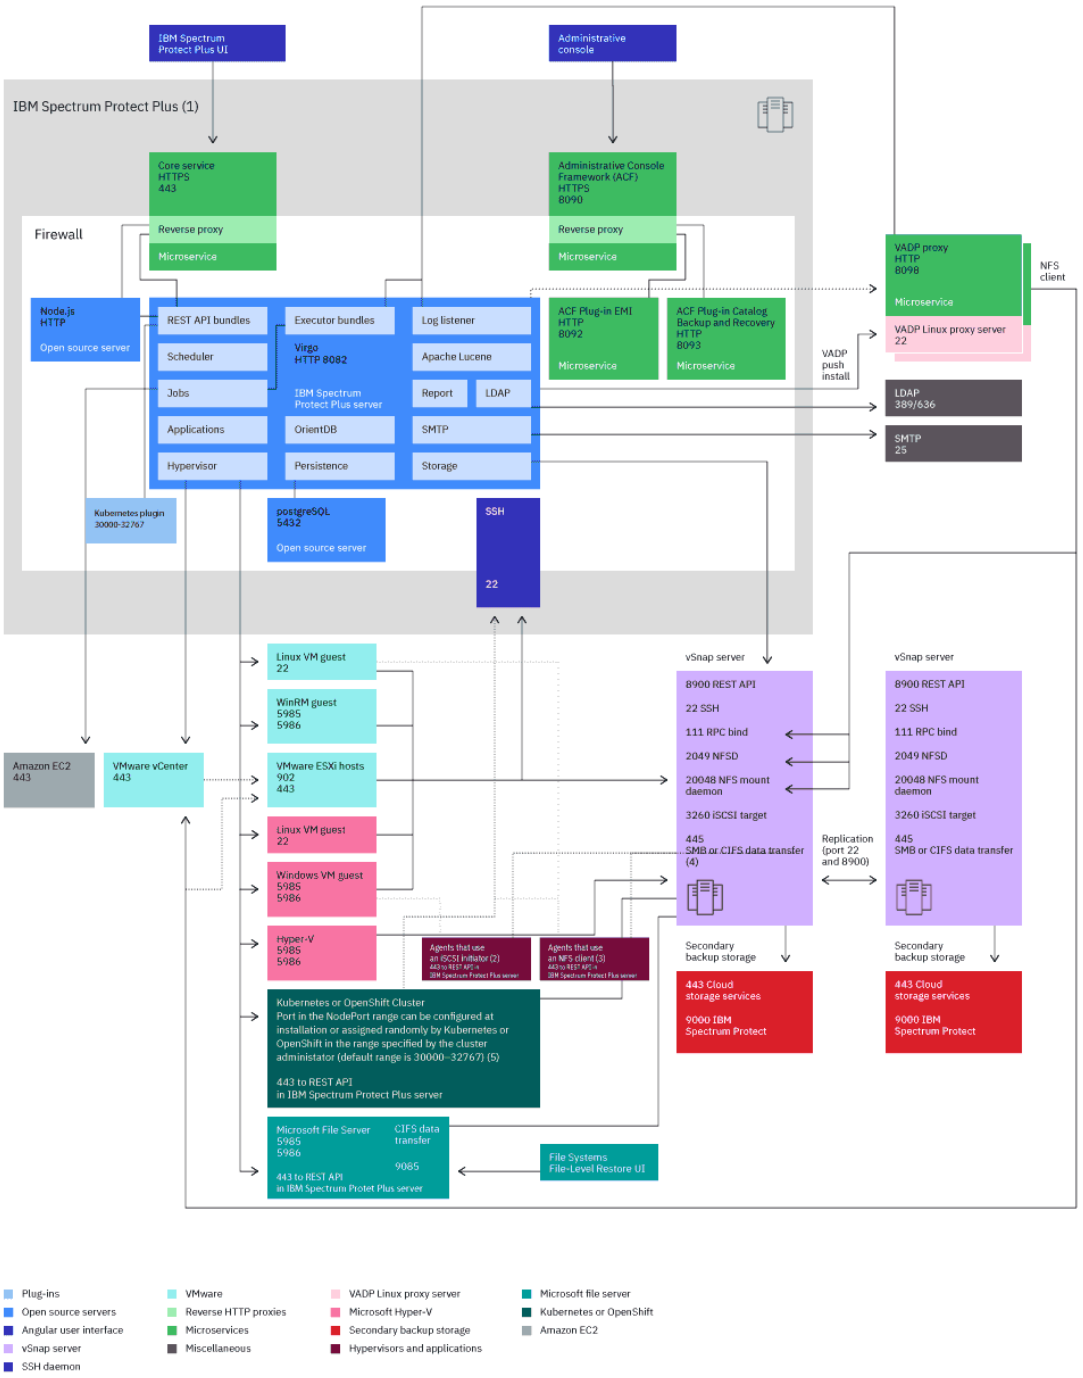


Figure 9. IBM Spectrum Protect Plus communication paths diagram

Component details:

1. The IBM Spectrum Protect Plus contains several base components, for detailed information, see “Product components” on page 6.
2. The following hypervisors and agents use an iSCSI initiator:
   - Hypervisor: Microsoft Hyper-V
   - Agents: Microsoft SQL Server, Microsoft Exchange, and Windows file systems.
3. The following hypervisors and agents use an NFS client:

- Hypervisor: VMware
- Agents: Oracle server, IBM Db2, MongoDB, Container (Kubernetes and OpenShift), and Microsoft 365.

  Note: IBM Db2 agent does not require port 443 to the REST API in the IBM Spectrum Protect Plus server.

4. The following agents use a Server Message Block (SMB) or the Common Internet File System (CIFS) protocol client:
   - Microsoft SQL Server (only for transaction log backup and restore operations)
   - Microsoft Exchange (only for transaction log backup and restore operations)
   - Windows File systems.
5. An SSH port connects the IBM Spectrum Protect Plus server to the Kubernetes Backup Support agent. If you do not select a port, a random port number is selected by the NodePort Services in the default range. If you specify a value for this port, use a port number within the NodePort range that is set by the Kubernetes administrator that is not already in use.

Connectivity requirements

Ensure that the following connectivity requirements are met:

- The secure file transfer protocol (SFTP) subsystem for Secure Shell (SSH) is enabled on the IBM Spectrum Protect Plus server, VADP proxies, and vSnap servers.
- The Secure Shell (SSH) service is running on port 22 on the IBM Spectrum Protect Plus server, VADP proxies, and vSnap servers.
- Firewalls are configured to allow IBM Spectrum Protect Plus components to connect with each other by using SSH.
- VADP proxy servers use the Network File System (NFS) to mount storage volumes for backup and restore operations. On Linux, ensure that the native Linux NFS client is installed.
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment can be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and from the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.
- If DNS is not available, you must add the server to the `/etc/hosts` file on the IBM Spectrum Protect Plus server by using the command line.

Repository server storage requirements

If you plan to use IBM Spectrum Protect as a repository server for copying data to cloud storage, ensure that you are using IBM Spectrum Protect V8.1.12.

Cloud storage requirements

Disk cache area

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

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

For instructions about sizing and installing the cache, see the IBM Spectrum Protect Plus Blueprints.

Multipath

During copy operations to object storage, IBM Spectrum Protect Plus attaches and detaches virtual cloud devices on vSnap servers. If a multipath configuration is enabled on the vSnap server by using `dm-multipath`, the configuration can interfere with the copy operation. To avoid this interference, modify the multipath configuration file and specify a rule to exclude devices whose vendor matches "LIO-ORG". For instructions and examples, go to the Red Hat Customer Portal and see the DM Multipath documentation

Certificates

- **Self-signed certificates**

  If the cloud endpoint or repository server uses a self-signed certificate, you must specify the certificate in Privacy Enhanced Mail (PEM) format when you register the cloud or repository server in the IBM Spectrum Protect Plus user interface.

- **Certificates signed by a private certificate authority**

  If the cloud endpoint or repository server uses a certificate signed by a private certificate authority (CA), the endpoint certificate must be specified (in PEM format) when you register the cloud or repository server in the IBM Spectrum Protect Plus user interface.

  In addition, you must add the root or intermediate certificate of the private CA to the system certificate store in each vSnap server by using the following procedure:

  1. Log in to the vSnap server console as the `serveradmin` user and upload any private CA certificates (in PEM format) to a temporary location.
  2. Copy each certificate file to the system certificate store directory (`/etc/pki/ca trust/source/anchors/`) by running the following command: `$ sudo cp /tmp/private-ca-cert.pem /etc/pki/ca-trust/source/anchors/`
  3. To incorporate the newly added custom certificate and update the system certificate bundle, run the following command: `$ sudo update-ca-trust`

- **Certificates signed by public certificate authority**

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

- **Wildcard certificates**

  If the cloud endpoint uses a wildcard certificate, note that the wildcard applies only to one subdomain level of the domain name. For example, if the certificate is for *.example.com, the certificate will match hostname level1.example.com but will not match level1.level2.example.com. If the bucket name contains periods (for example, "my.bucket") and it is part of the hostname used for registering the cloud endpoint in IBM Spectrum Protect Plus (for example, "my.bucket.example.com"), certificate validation can fail. In such cases, ensure that the bucket name does not contain periods.

Network

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

Table 14. Communication ports when the target is a cloud server or repository server endpoint

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

Table 14. Communication ports when the target is a cloud server or repository server endpoint (continued)

Port	Protocol	Initiator	Target	Description
9000	TCP	vSnap server	Repository server endpoints	Allows the vSnap server to communicate with IBM Spectrum Protect (repository server) endpoints.

Any firewalls or network proxies that inspect TLS or conduct a deep packet inspection of traffic between the vSnap servers and cloud endpoints might interfere with TLS certificate validation on vSnap servers. This interference can also cause cloud copy job failures. To prevent this interference, the vSnap servers must be exempted from TLS interception and inspection in the firewall or proxy configuration.

Cloud provider

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

If the cloud provider uses an TLS certificate that is self-signed or signed by a private certificate authority, see "Connectivity requirements" on page 54.

- **Amazon S3 cloud requirements**
  - **Standard object storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing bucket in one of the supported storage tiers must be specified: S3 Standard, S3 Intelligent-Tiering, S3 Standard-Infrequent Access, or S3 One Zone-Infrequent Access.
  - **Archive object storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing bucket in one of the supported storage tiers must be specified: S3 Standard, S3 Intelligent-Tiering, S3 Standard-Infrequent Access, or S3 One Zone-Infrequent Access. IBM Spectrum Protect Plus directly uploads data files to the Glacier tier. Some small metadata files are stored in the default tier for the bucket. A copy of these metadata files is also placed into the Glacier tier for disaster recovery purposes.
- **IBM Cloud Object Storage requirements**
  - **Standard object storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing bucket must be specified. If the specified bucket has a Write Once Read Many (WORM) policy that locks objects for a certain time period, IBM Spectrum Protect Plus automatically detects the configuration and deletes snapshots after the WORM policy removes the lock. The bucket must have the `Name Index` setting enabled.
  - **Archive object storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing bucket must be specified. If the specified bucket has a WORM policy that locks objects for a certain time period, IBM Spectrum Protect Plus automatically detects the configuration and deletes snapshots after the WORM policy removes the lock. IBM Spectrum Protect Plus creates a single lifecycle management rule on the bucket to migrate data files to the archive tier. The bucket must have the `Name Index` setting enabled.
- **Microsoft Azure requirements**
  - **Standard object storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing container in a hot or cool storage account must be specified.
  - **Archive object storage:** When the cloud provider is registered in IBM Spectrum Protect Plus, an existing container in a hot or cool storage account must be specified. IBM Spectrum Protect Plus moves files between tiers on demand. Data files are immediately moved to the archive tier and temporarily returned to the hot tier only during restore operations. Some small metadata files are

stored in the default tier for the container. A copy of these metadata files is also placed in the archive tier for disaster recovery purposes.

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

Table 15. Copy and archive copy requirements for cloud providers

Operation	Provider	Requirements
Copy	Amazon S3	An existing bucket must be specified from one of the supported storage tiers.
Copy	IBM Cloud Object Storage	An existing bucket must be specified. The bucket must have the `Name Index` setting enabled.
Copy	Microsoft Azure	An existing container must be specified from a hot or cool storage tier.
Copy	IBM Spectrum Protect	IBM Spectrum Protect Plus creates its own unique bucket.
Archive copy	Amazon S3	An existing bucket must be specified from one of the supported storage tiers.
Archive copy	IBM Cloud Object Storage	An existing bucket must be specified from the archive tier. The bucket must have the `Name Index` setting enabled.
Archive copy	Microsoft Azure	An existing container must be specified from the hot storage tier and archive tier.
Archive copy	IBM Spectrum Protect	IBM Spectrum Protect Plus creates its own unique bucket to be copied to IBM Spectrum Protect tape storage.

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

Review the hypervisor requirements for IBM Spectrum Protect Plus.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861

Hyper-V requirements

Configuration

Table 16. Coverage matrix for Microsoft Hyper-V servers supported by IBM Spectrum Protect Plus

IBM Spectrum Protect Plus	Microsoft Hyper-V on Windows Server 2016	Hyper-V Server 2016	Microsoft Hyper-V on Windows Server 2016	Hyper-V Server 2019
V10.1.0	✓	✓		
V10.1.1	✓	✓		
V10.1.2	✓	✓		
V10.1.3	✓	✓	✓	
V10.1.4	✓	✓	✓	✓
V10.1.5	✓	✓	✓	✓
V10.1.6	✓	✓	✓	✓
V10.1.7	✓	✓	✓	✓
V10.1.8	✓	✓	✓	✓

Beginning with IBM Spectrum Protect Plus V10.1.5, you can protect virtual machines (VMs) that are enabled to use the Hyper-V Replica feature. Depending on your Hyper-V environment, you might be required to update some service level agreement (SLA) policies when you update your system environment to IBM Spectrum Protect Plus V10.1.5 or later levels. For more information, see "Additional steps for updating virtual machines in Hyper-V Replica environments" on page 251.

Restrictions

- Windows file indexing and file restore operations on volumes residing on dynamic disks are not supported.
- For Hyper-V data, backup and restore operations are supported only for virtual hard disks (VHDX).
- File indexing and file restore operations are not supported from restore points that were copied to cloud resources or repository servers.
- File cataloging, backup, point-in-time restores, and other operations that invoke the Windows agent will fail if a nondefault local administrator ID is entered as the Guest OS Username when you define a backup job. A nondefault local administrator is any user ID that was created in the guest operating system and has been granted the administrator role.
- Virtual hard disk exclusion for Hyper-V is available only using Hyper-V Server 2019 and later.

Software

Ensure that the newest Hyper-V integration services are installed:

- Ensure that the 64-bit Microsoft Visual C++ 2008 SP1 Redistributable Package is installed on the VM guest machine, before you start restore operation from a backup image.
- For Microsoft Windows environments, see Supported Windows guest operating systems for Hyper-V on Windows Server
- For Linux environments, see Supported Linux and FreeBSD virtual machines for Hyper-V on Windows

Connectivity

Ensure that the following connectivity requirements are met:

- The network adapter that is used for the connection must be configured as a client for Microsoft Networks.
- The Microsoft Windows Remote Management (WinRM) service must be running.
- Firewalls must be configured to enable IBM Spectrum Protect Plus to connect to the server by using WinRM.
- The IP address of the machine that you register must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. The Hyper-V server must have a WinRM service that is listening on port 5985.
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment must be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names. If the Hyper-V server is part of a cluster, all nodes in the cluster must be resolvable by DNS.
- If DNS is not available, you must add the server to the `/etc/hosts` file on the IBM Spectrum Protect Plus server by using the command line. If more than one Hyper-V server is set up in a cluster environment, you must add all of the servers to the `/etc/hosts` file.
- When you are registering the cluster in IBM Spectrum Protect Plus, register the Failover Cluster Manager.
- Ensure that the Microsoft iSCSI Initiator Service is running on all Hyper-V servers, including cluster nodes. In the Services window, set the startup type for the Microsoft iSCSI Initiator Service to **Automatic** so that the service is available when the Hyper-V server or cluster node starts.
- **Troubleshooting tip:** If the IP address of the IBM Spectrum Protect Plus server is changed after an initial Hyper-V base backup is created, the target iSCSI qualified name (IQN) of the Hyper-V resource might be left in a bad state. To correct this issue, from the Microsoft iSCSI Initiator tool, click the **Discovery** tab. Select the old IP address, then click **Remove**. Click the **Target** tab and disconnect the reconnecting session.

Prerequisites and operations

Prerequisites

Ensure that the Software and Connectivity requirements are met.

Operations

Before you start a backup or restore operation, ensure that your system meets the following requirements:

- Register the providers that you want to back up. For instructions, see “Adding a Hyper-V server” on page 360.
- A service level agreement (SLA) policy is configured.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the Accounts pane. Add the user to the local administrator group on the Hyper-V server.

Review the following information about creating backup and restore jobs:

- Use a backup job to back up Hyper-V data with snapshots. For instructions, see "Backing up Hyper-V data" on page 363.
- Hyper-V restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source. For instructions, see "Restoring Hyper-V data" on page 367.

VMware requirements

Configuration

Table 17. Coverage matrix for VMware vSphere versions supported by IBM Spectrum Protect Plus

IBM Spectrum Protect Plus	VMware vSphere 6.0*	VMware vSphere 6.5*	VMware vSphere 6.7*	VMware vSphere 7.0*
V10.1.0	✓	✓		
V10.1.1	✓	✓		
V10.1.2	✓	✓	✓	
V10.1.3	✓	✓	✓	
V10.1.4	✓	✓	✓	
V10.1.5	✓	✓	✓	
V10.1.6	✓	✓	✓	✓
V10.1.7	✓	✓	✓	✓
V10.1.8	(1)	✓	✓	✓

*The base level and later updates and patch levels are supported.

(1)Beginning with IBM Spectrum Protect Plus V10.1.8, VMware VDDK 7.0 is included. This VDDK level does not support vSphere 6.0.

Review the following information about supported functionality:

- Backing up and restoring encrypted VMs is supported with vSphere 6.5 and later.
- IBM Spectrum Protect Plus supports VMware VM tags.
- IBM Spectrum Protect Plus V10.1.5 and later protects VMs that are managed by a VMware Cloud (VMC) on an Amazon Web Services (AWS) Software-Defined Data Center (SDDC). For more information, see IBM Spectrum Protect Plus for VMware Cloud on AWS

Restrictions

- Restored VM templates cannot be powered on after recovery of a VM.
- Secure Shell (SSH) keys are not a valid authorization mechanism for Windows platforms.

- Physical RDM (pRDM) volumes do not support snapshots. VMs that contain one or more raw device-mapping (RDM) volumes that are provisioned in pRDM mode are backed up. However, the pRDM volumes are not processed as part of the VM backup operation.
- File cataloging, backup, point-in-time restores, and other operations that invoke the Windows agent will fail if a nondefault local administrator ID is specified as the Guest OS Username when you define a backup job. A nondefault local administrator is any ID that has been created in the guest operating system and has been granted the administrator role.
- Windows file indexing and file restore operations on volumes residing on dynamic disks are not supported.

Software

- Ensure that the most recent version of VMware Tools is installed on VMware VMs.
- Ensure that the 64-bit Microsoft Visual C++ 2008 SP1 Redistributable Package is installed on the VM guest machine, before you start restore operation from a backup image.

Connectivity

For VMware hypervisor connectivity requirements, see System requirements: IBM Spectrum Protect Plus V10.1.8

Privileges

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

If the user that is associated with the provider is not assigned to the Administrator role for an inventory object, the user must be assigned to a role that has the required privileges, as described in “Virtual machine privileges ” on page 336.

Prerequisites and operations

Prerequisites

Ensure that the Software, Connectivity, and Privileges requirements are met.

Operations

Before you start a backup or restore operation, ensure that your system meets the following requirements:

- Register the providers that you want to back up. For instructions, see “Adding a vCenter Server instance” on page 335.
- A service level agreement (SLA) policy is configured.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the **Accounts** pane.

Review the following information about creating backup and restore jobs:

- Use a backup job to back up VMware resources such as VMs, datastores, folders, vApps, and datacenters with snapshots. For instructions, see “Backing up VMware data” on page 340.
- In IBM Spectrum Protect Plus, you can create proxies to run VMware backup jobs by using vStorage API for Data Protection (VADP) in Linux environments. The proxies reduce demand on system resources by enabling load sharing and load balancing. For instructions, see “Managing VADP backup proxies” on page 345.
- VMware restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source. For instructions about creating VMware restore jobs, see “Restoring VMware data” on page 350

Amazon EC2 requirements

EC2 data is stored in Amazon Web Services (AWS) Elastic Block Store (EBS) snapshots rather than the vSnap server. IBM Spectrum Protect Plus manages these snapshots for backup and restore operations.

Prerequisites and operations

Prerequisites

- To protect Amazon EC2 data, first add an account for your EC2 instances in IBM Spectrum Protect Plus, and then create jobs for backup and restore operations for those instances.
- To add an EC2 account to IBM Spectrum Protect Plus, access keys are required. Access keys are long-term credentials for an Identity and Access Management (IAM) user or the Amazon Web Services (AWS) account root user.
- For information about how to create an IAM user with access keys and the permissions that are required for IBM Spectrum Protect Plus, see "Creating an AWS IAM user " on page 373.
- For increased security, avoid using the AWS account root user for IBM Spectrum Protect Plus. For more information about the root user, see the AWS Identity and Access Management User Guide
- For Amazon EC2 workloads, the IBM Spectrum Protect Plus hostname must consist of only lowercase and alphanumeric characters.

Operations

Before you start a backup or restore operation, ensure that your system meets the following requirements:

- When an Amazon EC2 account is added to IBM Spectrum Protect Plus, an inventory of the instances that are associated with the account is captured. For more information, see "Adding an Amazon EC2 account" on page 375.
- Ensure that one or more SLA policies are configured for the EC2 instances. For instructions, see "Creating an SLA policy for Amazon EC2 instances" on page 328.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the **Accounts** pane.

Review the following information about creating backup and restore jobs:

- You can use a backup job to back up data in an Amazon EC2 instance. For instructions, see "Backing up Amazon EC2 data" on page 376.
- You can use a restore job to restore EC2 data from a backup copy. You can restore data to the original availability zone or to a different availability zone in the same region, with different types of recovery options and configurations. For instructions, see "Restoring Amazon EC2 data" on page 378.

Ports

The following ports are used by IBM Spectrum Protect Plus hypervisors.

Table 18. Communication ports when the target is an IBM Spectrum Protect Plus hypervisor (VMware, Microsoft Hyper-V, or Amazon EC2)

Port	Protocol	Initiator	Target	Description
443	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus server	Hypervisor: VMware Elastic Sky X Integrated (ESXi) host and vCenter	Provides access to ESXi and vCenter for managing operations.
443	TCP	VADP proxy host	Hypervisor: VMware ESXi host	Provides access to ESXi and vCenter for managing operations.
443	TCP	IBM Spectrum Protect Plus server	Hypervisor: Amazon EC2	Provides access to AWS for managing operations.

Table 18. Communication ports when the target is an IBM Spectrum Protect Plus hypervisor (VMware, Microsoft Hyper-V, or Amazon EC2) (continued)

Port	Protocol	Initiator	Target	Description
902	TCP	IBM Spectrum Protect Plus server	Hypervisor: VMware ESXi host	Used for the Network File Copy (NFC) protocol, which provides a file-type-aware File Transfer Protocol (FTP) service for vSphere components. By default, ESXi uses NFC for operations such as copying and moving data between datastores.
902	TCP	VADP proxy host	Hypervisor: VMware ESXi host	Used for the Network File Copy (NFC) protocol, which provides a file-type-aware File Transfer Protocol (FTP) service for vSphere components.
5985	TCP	IBM Spectrum Protect Plus server	Hypervisor: Microsoft Hyper-V	Provides access to the Microsoft WinRM service for Windows-based servers.
5986	TCP	IBM Spectrum Protect Plus server	Hypervisor: Microsoft Hyper-V	Provides access to the Microsoft WinRM service for Windows-based servers.

Table 19. Communication ports when the initiator is an IBM Spectrum Protect Plus hypervisor (VMware, Microsoft Hyper-V, or Amazon EC2)

Port	Protocol	Initiator	Target	Description
22	TCP	Hypervisor	vSnap server	Provides access for troubleshooting and maintenance tasks on vSnap servers by using SSH protocol.

Table 19. Communication ports when the initiator is an IBM Spectrum Protect Plus hypervisor (VMware, Microsoft Hyper-V, or Amazon EC2) (continued)

Port	Protocol	Initiator	Target	Description
111	TCP and User Datagram Protocol (UDP)	Hypervisor: VMware ESXi host	vSnap server	Used for Network File System (NFS) file sharing by the vSnap server.
2049	TCP and UDP	Hypervisor: VMware ESXi host	vSnap server	Used for NFS file sharing by the vSnap server.
3260	TCP	Hypervisor: Microsoft Hyper-V	vSnap server	Used for Internet Small Computer System Interface (iSCSI) data transfer by vSnap servers.
20048	TCP and UDP	Hypervisor: VMware ESXi host	vSnap server	Used for NFS file sharing by the vSnap server.

File indexing and restore requirements

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

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861

General

- For hypervisor operations, IBM Spectrum Protect Plus supports only the operating systems that are available to your hypervisors. For information about supported operating systems, review the hypervisor documentation.
- IBM Spectrum Protect Plus can protect and restore virtual machines (VMs) with file systems that are not listed in this documentation, but only the listed file systems are eligible for file indexing and restore operations.
- Internet Small Computer System Interface (iSCSI) disks that are directly mapped to the guest operating system will not be indexed. Supported volumes include virtual machine disk (VMDK) volumes that are mounted as specified by the configuration of the associated VM.
- The maximum number of files that can be indexed per IBM Spectrum Protect Plus server is 2147483647.
- The amount of free space that is required for the metadata in the catalog depends on the total number of files in the environment. To catalog 1 million files, the catalog volume in the IBM Spectrum Protect Plus server requires roughly 350 MB of free space per retained version.

  Example: 2 virtual machines that each contain 10 million files are backed up daily with a retention of 30 days. The resulting number of catalog entries is 2 * 10 million * 30 days for a total of 600 million entries. Therefore, the approximate space needed is 600 * 350 MB which is 210 GB.

  The space that is used by file indexing metadata is reclaimed when the corresponding backup instances expire.
- During file indexing, results are temporarily stored in **/data2** directory prior to being saved in the catalog. If there are several jobs running concurrently, the size of **/data2** directory may need to be increased from the default size.

- File indexing and file restore are not supported from restore points that were copied to cloud resources or repository servers.
- A file can be restored to an alternative location only if credentials were established for the alternative VM by using the **Guest OS Username** and **Password** option in the backup job definition.

VMware requirements

- Ensure that the most recent version of VMware Tools is installed on VMware VMs.
- In the VM settings under Advanced Configuration, the **`disk.EnableUUID`** parameter must be set to `true`.
- File indexing and restore operations support Small Computer System Interface (SCSI) disks in a VMware environment:

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

Hyper-V requirements

- Ensure that the most recent version Hyper-V Integration Services is installed on your Hyper-V VMs.
- File indexing and restore operations support Small Computer System Interface (SCSI) disks in a Hyper-V environment:
  - Only volumes on SCSI disks are eligible for file cataloging and file restore operations.
  - Integrated Drive Electronics (IDE) disks are not supported.

Windows requirements

Configuration

Table 20. Coverage matrix for Microsoft Windows 64-bit operating systems

IBM Spectrum Protect Plus	Windows Server 2008 R2* Standard and Datacenter editions	Windows Server 2012 R2 and Windows Server 2012R2 core* Standard and Datacenter editions	Windows Server 2016 and Windows Server 2016 core* Standard and Datacenter editions	Windows Server 2019 and Windows Server 2019 core* Standard and Datacenter editions
V10.1.0	✓	✓	✓	--
V10.1.1	✓	✓	✓	--
V10.1.2	✓	✓	✓	--
V10.1.3	✓	✓	✓	✓ (Windows Server 2019 core only)
V10.1.4	✓	✓	✓	✓
V10.1.5	✓	✓	✓	✓

Table 20. Coverage matrix for Microsoft Windows 64-bit operating systems (continued)				
IBM Spectrum Protect Plus	**Windows Server 2008 R2* Standard and Datacenter editions**	**Windows Server 2012 R2 and Windows Server 2012R2 core* Standard and Datacenter editions**	**Windows Server 2016 and Windows Server 2016 core* Standard and Datacenter editions**	**Windows Server 2019 and Windows Server 2019 core* Standard and Datacenter editions**
V10.1.6	✓	✓	✓	✓
V10.1.7	✓	✓	✓	✓
V10.1.8	✓	✓	✓	✓
* The base release and later maintenance levels are supported. Windows Server core refers to Windows Server with the Server Core option				

Table 21. Coverage matrix for supported file systems and disk storage types	
Supported file systems	• New Technology File System (NTFS) • Resilient File System (ReFS) • File allocation table (FAT)
Supported disk storage types	Basic disks with the following partitions: • MBR (Master Boot Record) • GPT (GUID Partition Table) **Restriction:** You cannot back up or restore files on dynamic disks.

Restrictions

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

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

  Files within these directories are not added to the IBM Spectrum Protect Plus inventory and are not available for file recovery.
- Encrypted Windows file systems are not supported for file cataloging or file restore.
- When restoring files in a ReFS environment, restore jobs from newer versions of Windows Server to earlier versions are not supported. For example, you cannot restore a file from Windows Server 2016 to Windows Server 2012.

Disk space

- The `C:\` drive must have sufficient temporary space to save the file indexing results.
- When file systems are indexed, temporary metadata files are generated under the `C:\ProgramData\SPP` directory and are deleted when the indexing is complete. For information about free space see the section General

Software

- File indexing and file restore operations for a Windows VM require that the Windows PowerShell binary path is set in the %PATH% environment variable.
- Ensure that the 64-bit Microsoft Visual C++ 2008 SP1 Redistributable Package is installed on the VM guest machine, before you start restore operation from a backup image.
- Install a supported version of a Windows 64-bit operating system in your environment. Ensure that the most recent patches and updates are installed.

Connectivity

Ensure that your system environment meets the following connectivity requirements:

- The hostname of the IBM Spectrum Protect Plus server should be resolvable from the Windows VM.
- The Internet Protocol (IP) address of the VM that is selected for indexing must be visible to the vSphere client or Hyper-V Manager.
- The Windows VM that is selected for indexing must support outgoing connections to port 22, which uses the Secure Shell (SSH) protocol, on the IBM Spectrum Protect Plus server.
- The network adapter that is used for the connection must be configured as a client for Microsoft Networks.
- The Microsoft® Windows Remote Management (WinRM) service must be running.
- Firewalls must be configured to enable IBM Spectrum Protect Plus to connect to the server by using WinRM.
- The IP address of the machine that you register must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. A Windows guest machine must have a WinRM service that is listening on port 5985.
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment must be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and from the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.
- If DNS is not available, you must add the server to the /etc/hosts file on the IBM Spectrum Protect Plus server by using the command line.

Authentication and privilege requirements

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

- The system login credentials must have the permissions of the local administrator.
- The user identity must have the "Log on as a service" right, which is assigned through the Administrative Tools control panel on the local server (**Local Security Policy** > **Local policies** > **User Rights Assignment** > **Log on as a service**).

  For more information about the "Log on as a service" right, see Add the Log on as a service Right to an Account.
- The default security policy uses the Windows Challenge/Response (NTLM) protocol, and the user identity follows the default `domain\Name` format if the Hyper-V VM is attached to a domain. The format `local administrator` is used if the user is a local administrator. Credentials must be established for the associated VM by using the **Guest OS user name** and **Guest OS password** option within the associated backup job definition.
- File cataloging, backup, point-in-time restores, and other operations that start the Windows agent fail if a nondefault local administrator ID is entered as the Guest OS username when you define a backup job. A nondefault local administrator is any user ID that is created in the guest operating system and is granted the administrator role.

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

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

Linux requirements

Table 22. Coverage matrix for supported operating systems on Linux® x86_64

IBM Spectrum Protect Plus	RHEL 6.4*	RHEL 7.0*	RHEL 8.0*	CentOS 6.4*	CentOS 7.0*	CentOS 8.0*	SLES 12.0*	SLES 15.0*
V10.1.0	✓	✓	--	✓	✓	--	✓	--
V10.1.1	✓	✓	--	✓	✓	--	✓	--
V10.1.2	✓	✓	--	✓	✓	--	✓	--
V10.1.3	✓	✓	--	✓	✓	--	✓	--
V10.1.4	✓	✓	--	✓	✓	--	✓	--
V10.1.5	✓	✓	--	✓	✓	--	✓	--
V10.1.6	✓	✓	✓	✓	✓	✓	✓	✓
V10.1.7	✓	✓	✓	✓	✓	✓	✓	✓
V10.1.8	✓	✓	✓	✓	✓	✓	✓	✓

* The base release and later maintenance levels are supported.

Table 23. Coverage matrix for supported file systems

Supported file systems	• ext2 • ext3 • ext4 • XFS

Restrictions

- A file system that was created on a newer kernel version might not be mountable on a system with a previous kernel version. In this case, restoring files from the newer to the previous system is not supported.
- When files are indexed in a Linux environment, the following directories on the resource are skipped:

  ```
  /tmp
  /usr/bin
  /Drivers
  /bin
  /sbin
  ```

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

Disk space

- The system disk must have sufficient temporary space to save the file indexing results.
- When file systems are indexed, temporary metadata files are generated under the `/tmp` directory and are then deleted when the indexing is complete. For information about free space see the section General.

Software

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

  Tip: The required **bash** and **sudo** packages are included in the supported Linux 86_64 operating system installation packages.
- Ensure that the supported version of Linux x86_64 is installed. Ensure that the most recent patches and updates are installed.
- The International Components for Unicode (**libicu**) RPM-package must be installed for the corresponding version of your operating system.
- In a Linux environment, depending on your version or distribution, ensure that the Linux utility package `util-linux-ng` or `util-linux` package is current.
- **RHEL and CentOS 6 users:** To ensure that the `util-linux-ng` or `util-linux` package is current, run the following command:

  ```
  yum update package_name
  ```

  where *package_name* is the name of the Linux utility package.
- If data resides on Logical Volume Manager (LVM) volumes, ensure that the LVM version is 2.0.2.118 or later.

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

  1. Run the following commands:

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

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

     ```
     use_lvmetad = 0
     ```

     For more information, see The Metadata Daemon (lvmetad).

- If data resides on XFS file systems and the version of the `xfsprogs` package is between 3.2.0 and 4.1.9, the file restore operation can fail due to a known issue in the `xfsprogs` package that causes corruption of a clone or snapshot file system when its Universally Unique Identifier (UUID) is modified. To resolve this issue, update the `xfsprogs` package to version 4.2.0 or later. For more information, see Debian Bug report logs.

Connectivity

Ensure that your system environment meets the following connectivity requirements:

- The secure file transfer protocol (SFTP) subsystem for SSH is enabled.
- The SSH service is running on port 22 on the proxy host server.
- Firewalls are configured to allow IBM Spectrum Protect Plus to connect to the proxy host server by using SSH.
- IBM Spectrum Protect Plus uses the Network File System (NFS) protocol to mount storage volumes for backup and restore operations. Ensure that the native Linux NFS client is installed on the proxy host server.
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment can be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and from the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.
- If DNS is not available, you must add the server to the `/etc/hosts` file on the IBM Spectrum Protect Plus server by using the command line.

Authentication and privileges

IBM Spectrum Protect Plus requires root privileges by using the **sudo** program for various tasks such as discovering storage layouts, mounting and unmounting disks, and managing databases. The credentials for the VM must specify a user with the following **sudo** privileges:

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

The recommended approach is to create a dedicated IBM Spectrum Protect Plus agent user with the privileges that are shown in the sample configuration:

- Create the user by using the command:

  ```
  useradd -m sppagent
  ```

  where *sppagent* specifies the IBM Spectrum Protect Plus agent user.
- Set a password by using the command:

  ```
  passwd sppagent_password
  ```

  where *sppagent_password* specifies the agent password.
- To enable superuser privileges for the agent user, set the `!requiretty` setting. At the end of the `/etc/sudoers` configuration file, add the following lines:

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

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

File system requirements

Before you register Microsoft Windows file systems with IBM Spectrum Protect Plus, ensure that your system environment meets the outlined requirements.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861

The IBM file systems backup and restore requirements for IBM Spectrum Protect Plus are as follows.

General

Beginning with IBM Spectrum Protect Plus V10.1.6, support was added for backing up and restoring Microsoft Windows file systems data.

Before you register a client host where the Microsoft Windows file system is located with IBM Spectrum Protect Plus, ensure that the system environment meets the following requirements.

Configuration

Application versions

Table 24. Coverage matrix for Microsoft Windows file systems supported by IBM Spectrum Protect Plus

IBM Spectrum Protect Plus	**Microsoft Windows Resilient File System** (ReFS)	**Microsoft New Technology File System** (NTFS)
V10.1.6	✓	✓
V10.1.7	✓	✓
V10.1.8	✓	✓

Restriction: Even if other Microsoft Windows file systems, such as File Allocation Table (FAT), are detected during the inventory process, these file systems are not protected by IBM Spectrum Protect Plus and cannot be added to backup or restore jobs.

Operating systems

Table 25. Coverage matrix for supported Microsoft Windows 64-bit operating systems

IBM Spectrum Protect Plus	**Microsoft Windows Server 2012 R2*** Standard and Datacenter editions	**Microsoft Windows Server 2016*** Standard and Datacenter editions	**Microsoft Windows Server 2019*** Standard and Datacenter editions
V10.1.6	✓	✓	✓
V10.1.7	✓	✓	✓
V10.1.8	✓	✓	✓
*The base release and later maintenance levels (64-bit kernel) are supported.			

Browser support

For supported browsers, see the *Browser support* section in System Requirements: IBM Spectrum Protect Plus V10.1.8.

Restrictions

- IBM Spectrum Protect Plus does not protect file system shares or Microsoft cluster volumes.
- Microsoft FAT file systems are not supported.
- Stub files are not supported.
- Network shares are not valid alternative locations for restore jobs.
- Only one application server or file server can be assigned per host. For example, if a host is registered as a Microsoft Windows file system, you cannot register the same host as a Microsoft SQL Server or a Microsoft Exchange Server.
- Backup and restore operations are not supported for a file server instance that was created from a template or cloned from other virtual machines (VMs).
- Alternate data streams are not protected.
- Sparse files are protected like normal files.
- Windows share definitions are not protected.

Software

- A supported version of a Microsoft Windows 64-bit operating system, including the most recent patches and updates, must be installed in your system environment.

Connectivity

Ensure that your system environment meets the following connectivity requirements:

- The network adapter used for the connection must be configured as a client for Microsoft Networks.
- The Microsoft Windows Remote Management (WinRM) service must be running.
- Firewalls must be configured to enable IBM Spectrum Protect Plus to connect to the server by using WinRM.
- Firewalls must be configured to enable the IBM Spectrum Protect Plus File Systems File-Level Restore browser to connect to the restore service.
- The IP address of the client host you register must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. The Windows file systems agent must have a Windows Remote Management service that is listening on port 5985.
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment must be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and from the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.

Authentication and privileges

Authentication

To register a Windows file system, an IBM Spectrum Protect Plus administrative user must register with the client host where the file systems to be protected are located.

Windows file servers can be registered with an Administrative user ID. However, you can register a file server by using a domain user ID, if that user is the domain administrator or a local user with administrator privileges.

Privileges

The user ID for registering Windows file servers can be set up with one of the following Windows configurations:

- For the local system administrator:
  - Ensure that Admin Approval Mode is disabled by completing the following steps:
    1. Click the **Windows System Control Panel > User Account Control Settings**.
    2. Ensure that the Never notify option is enabled.
- For members of the local Administrators group:
  - Disable the Admin Approval Mode security policy setting for a user who is a member of the local Ensure that the Administrators group by completing the following steps:
    1. Log in as a member of the local Administrators group and open the **Windows System Local Security Policy** window.
    2. From the **Security Settings** menu, click **Local Policies > security options > User account Control: Run all administrators** in **Admin Approval Mode** policy.
    3. Disable the **User Account Control: Run all administrators** option.
    4. Ensure that your local Administrator group includes the **Log on as a Service policy** setting.

See also User Account Control Group Policy and registry key settings.

Group Policy Object

For the Network security: LAN Manager authentication level policy setting, click **Computer Configuration** > **Windows Settings** > **Security Settings** > **Local Policies** > **Security Options** and specify one of the following options:

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

Important: For vSnap, you must use one of the security options listed above. Other options are not compatible with the vSnap Common Internet File System (CIFS) and the Server Message Block (SMB) protocol versions, and can cause CIFS authentication problems.

Specify the Group Policy Object (GPO) setting by navigating to **Computer Configuration** > **Windows Settings** > **Security Settings** > **Local Policies** > **Security Options** > **Network security: Restrict NTLM: Incoming NTLM traffic** and **Network security: Restrict NTLM: Outgoing NTLM traffic to remote servers**.

For the NTLM traffic, specify one of the following options:

- **Allow all**
- **Allow all accounts**

Prerequisites and operations

Prerequisites

Ensure that the Software, Connectivity, and Authentication and privileges requirements are met.

The following prerequisites must be met before you start protecting your resources. For details, see Prerequisites for file systems.

- Before you start backing up data that is stored on the registered file system, ensure that you have enough free disk space on the backup host and in the vSnap repository.
- If you plan to restore data to an alternative location, allow for extra space. No files are overwritten during the restore process. When files with identical names are found, you can decide whether to retain both copies or to overwrite data.

- If the IBM Spectrum Protect Plus file systems agent is running, a self-signed certificate and key are created. You can enhance the security of a protected file system by creating a certificate and managing its placement

Operations

Before you start a backup or restore operation, take the following actions:

- Add the file system servers that you want to back up.
- Configure a service level agreement (SLA) policy.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the **Accounts** pane.

Review the following information about creating backup and restore jobs:

- During the initial backup, IBM Spectrum Protect Plus creates a new vSnap volume and CIFS share. During incremental backups jobs, the previously created volume is reused. The IBM Spectrum Protect Plus file system agent mounts the share on the server where the backup is to be completed, as described in Backing up file system data.
- To restore file system data from the vSnap repository, define a job that restores data from either the newest backup or an earlier backup copy. You can restore data to the original location or to an alternative location. You can also specify other recovery options, as described in Restoring file system data.

For an overview about protecting Windows file systems with IBM Spectrum Protect Plus, see Protecting Windows file systems.

Ports

The following ports are used by IBM Spectrum Protect Plus agents users.

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

Port	Protocol	Initiator	Target	Description
5985	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus server	Windows file systems	Provides access to the Microsoft WinRM service for Windows-based servers
5986	TCP	IBM Spectrum Protect Plus server	Windows file systems	Provides access to the Microsoft WinRM service for Windows-based servers
9085	TCP	File Systems File-Level Restore browser	Windows file systems	The File System File-Level Restore browser used during restore operations to connect the user interface and the file server

Table 27. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user

Port	Protocol	Initiator	Target	Description
443	TCP	Windows file systems	IBM Spectrum Protect Plus server	Port that allows the agent to communicate with IBM Spectrum Protect Plus for making representational state transfer application programming interface (REST API) calls.
445	TCP	Windows file systems	vSnap server	Used for SMB or CIFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations

Hardware

Table 28. Minimum hardware requirements

System	Disk space	Memory
x86_64 based hardware that is compatible with one of the Windows operating system versions that is listed in the Software section	A minimum of 500 MB of disk space is required for product installation. The system also requires 1 GB per 1 million files of available disk space for temporary files at run time in the file system to be protected.	2.5 GB RAM per 1 million files for backup operations

Container Backup Support requirements

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

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861

General

Beginning with IBM Spectrum Protect Plus V10.1.5, support was added to protect persistent volume claims that are attached to containers in Kubernetes clusters. Operations were initiated by using the Kubernetes command line. In IBM Spectrum Protect Plus V10.1.6, backup support for containers was extended to the IBM Spectrum Protect Plus user interface. In addition the Container Backup Support package was made available for download from the IBM Helm Charts Repository by using the IBM Entitled Registry.

In IBM Spectrum Protect Plus V10.1.7, support was added to protect container clusters on Red Hat® OpenShift® Container Platform. In addition, container backup support was extended to protect Red Hat OpenShift and Kubernetes cluster-scoped and namespace-scoped resources. Container backup support was extended to include the IBM block storage Container Storage Interface (CSI) driver 1.2.0.

In IBM Spectrum Protect Plus V10.1.8, support was added to protect Red Hat OpenShift Container Platform and Kubernetes container data with added Ceph® File System (CephFS) support. In addition, container backup support was extended to protect Red Hat OpenShift Container Platform and Kubernetes container data with IBM Spectrum Scale.

Before you deploy IBM Spectrum Protect Plus V10.1.8 Container Backup Support in the Red Hat OpenShift or Kubernetes environment, ensure that the system environment meets the requirements and any prerequisite software is up-to-date, with all security-related patches applied.

Configuration

Application versions

Docker containers v17.09.00 and later are supported in Container Backup Support.

Operating systems

Table 29. Coverage matrix for supported Linux x86_64 operating systems for Kubernetes environments

IBM Spectrum Protect Plus	RHEL 7.6 & 7.7	RHEL 7.8	RHEL 7.9	RHEL 8.0 & 8.1 & 8.2 & 8.3
V10.1.8	Beginning with V10.1.5	Beginning with V10.1.6	Beginning with V10.1.7	Beginning with V10.1.8

Cluster requirements

Table 30. Coverage matrix for supported software and systems to protect Kubernetes or OpenShift environments that are attached to clusters

IBM Spectrum Protect Plus cluster requirements	Kubernetes environments	OpenShift environments
Kubernetes(1)	V1.18 and later updates V1.19 and later updates V1.20 and later updates V1.21 and later updates (Beginning with V10.1.8.1)	--
OpenShift Container Platform (OCP)(2)(3)	--	V4.5 and later updates (Beginning with V10.1.7) V4.6 and later updates (Beginning with V10.1.7) V4.7 and later updates (Beginning with V10.1.8)
OpenShift Data Foundation (ODF) (formerly OpenShift Container Storage)	--	V4.6 and later updates (Beginning with V10.1.7) V4.7 and later updates (Beginning with V10.1.8.1)
OpenShift API for Data Protection (4)to install Velero tool	--	V0.1.0, V0.1.1, V0.1.2 (Beginning with V10.1.7) V0.2.0 and V0.2.1 (Beginning with V10.1.8)

Table 30. Coverage matrix for supported software and systems to protect Kubernetes or OpenShift environments that are attached to clusters (continued)

IBM Spectrum Protect Plus cluster requirements	Kubernetes environments	OpenShift environments
Helm[5]	V3.3 and later (Beginning with V10.1.7)	V3.3 and later (Beginning with V10.1.7)
Velero to protect cluster-scoped and namespace-scoped resources [6][7]	V1.4.2, V1.4.3, V1.5.1 (Beginning with V10.1.7) V1.5.2 (Beginning with V10.1.8) V1.6.0 (Beginning with V10.1.8.1)	Installation by using OADP operator only (Beginning with V10.1.7)
Rook.io Ceph Storage	V1.4 and later (Beginning with V10.1.7)	V1.4 and later (Beginning with V10.1.7)
External Ceph File System (CephFS)	on Ceph Storage Cluster V15.2.8 or later (Beginning with V10.1.8)	OpenShift recommended version (Beginning with V10.1.8)
External Ceph Rados Block Device (RBD)	on Ceph Storage Cluster V14.2.2 or later (Beginning with V10.1.7)	OpenShift recommended version (Beginning with V10.1.7)
Ceph Container Storage Interface (CSI) driver with Rados Block Device (RBD) storage [8]	V3.0 (Beginning with V10.1.7)	(Installed with OCS)
Ceph Container Storage Interface (CSI) driver with Ceph FS storage	V3.1.0 (Beginning with V10.1.8)	(Installed with OCS)
IBM block storage CSI for virtualized storage[9]	V1.3 or later (Beginning with V10.1.8) V1.6 or later (Beginning with V10.1.8.2)	V1.3 and later (Beginning with V10.1.7) V1.6 or later (Beginning with V10.1.8.2)
IBM Spectrum Scale CSI driver[10]	V2.2.0 or later updates (Beginning with V10.1.8)	V2.2.0 or later updates (Beginning with V10.1.8)
IBM Spectrum Scale	V5.1.1 or later updates (Beginning with V10.1.8)	V5.1.1 or later updates (Beginning with V10.1.8)

Notes:

- [1] Kubernetes levels supported in earlier IBM Spectrum Protect Plus reached end of life, see Kubernetes Patch Releases
- [2]For information about OpenShift releases, see Red Hat OpenShift Container Platform Life Cycle Policy
- [3]OpenShift Container Platform (OCP) V4.6 support started with OCS V4.6.12 because of issue BZ - 1915007
- [4]To install the OADP operator on OCP 4.5 follow the procedure described in https://github.com/konveyor/oadp-operator/blob/master/docs/olm.md
- [5]For details which Helm level support Kubernetes level, see Helm Version Support Policy
  If the Helm V3.3.3 installation process fails during version validation, see the guidance in Warning after upgrading to 3.3.3. The workaround is to install a later version of Helm.
- [6]For instructions on installing Velero, see:
  - For OpenShift: Installing and configuring Velero by using the OADP Operator
    For details about version mapping between OADP and velero, see velero-version-relationship
  - For Kubernetes: Installing and configuring Velero

- [7]If an instance of Velero is already installed in the cluster, you must install and configure another instance of Velero. For more information and instructions, see Installing a second instance of Velero
- [8]For CSI driver 1.2, 2.0, and 2.1, use IBM Spectrum Protect Plus V10.1.6.
- [9]For IBM block storage CSI driver supported orchestration platforms, see under IBM block storage CSI driver > Release Notes > Compatibility and requirements
- [10]Backup and recovery of Scale Extended Attributes is not supported in this version.

To install and configure Container Backup Support, you must deploy the Container Backup Support software in the Kubernetes or OpenShift cluster environment. For instructions, see Chapter 6, "Installing Container Backup Support," on page 197.

Restrictions

The following restrictions apply to Kubernetes and OpenShift environments:

- OCP V4.5 cannot be installed from the web console. Use the command line to install OCP V4.5.
- Backup operations for raw block device volumes (volumeMode 'Block') are not supported.
- To ensure that a snapshot restore operation request works correctly, do not manually delete any snapshots of volumes that are protected by Container Backup Support.
- You cannot restore a snapshot backup to a different cluster or namespace.
- Container Backup Support protects only persistent storage that was allocated by a storage plug-in that supports the CSI.
- Only formatted volumes can be mounted to the data mover for copy operations.
- The Container Backup Support component is available only in English.

Software

Cluster prerequisites

- Command line tool:
  - Kubernetes environment: The Kubernetes command line tool **`kubectl`** must be accessible on the installation host and in the local path.
  - OpenShift environment: The OpenShift command line tool **oc** must be accessible on the installation host and in the local path.
- Tips for collecting metrics and improving performance:
  - On Kubernetes environment: To help optimize product performance and scalability, ensure that Kubernetes Metrics Server v0.3.5 or later is installed and running on your cluster. For instructions, see Verifying whether the metrics server is running.
  - In an OpenShift environment: The Kubernetes Metrics Server is included and augmented with Prometheus and Prometheus-Adapter for custom metrics. Prometheus and Prometheus-Adapter are part of the OpenShift Cluster Monitoring Operator. Ensure that the OpenShift Cluster Monitoring Operator is installed and running in the environment.
- CSI external-snapshotter:
  - Kubernetes 1.17-1.19 environment: The CSI external-snapshotter v2.1.1 or later is required for snapshots of volumes on a storage system.
  - Kubernetes 1.20 and later environment: The CSI external-snapshotter v4.0.0 or later is required for snapshots of volumes on a storage system.
  - OpenShift environment: The external-snapshotter is part of the installation package. Ensure that the cluster operator csi-snapshot-controller is in the `Available: True` state.
- A storage class must be defined for the persistent volumes that are being protected.
- The target image registry must be accessible from the Kubernetes or OpenShift cluster. The target image registry can be a local image registry or an external image registry.

- The host that is used to install Container Backup Support must be using a `kubeconfig` file with cluster-admin privileges, KUBECONFIG, and the Helm client must be installed.
- To create new cluster-wide resources, you must be logged in to the target cluster as a user with `cluster-admin` privileges.
- Ensure that Container Backup Support secrets that include user IDs, passwords, and keys are encrypted at rest in the `etcd` distributed key-value store. For more information, see Encrypting Secret Data at Rest.

Helm prerequisites

- Helm 3 is an application package manager that runs on Kubernetes or OpenShift. Helm is designed to simplify the definition, storage, and management of applications. The installation process for Container Backup Support uses a Helm 3 chart. The installation script that is provided with the installation package requires that the Helm 3 binary file is renamed to `helm3`. For instructions, see "Installing Helm 3 and renaming the binary file" on page 197.
- The Helm tool must be configured on the target cluster so that a new deployment can be run with the **`helm`** command line. Deploying a package with Helm enables cluster-wide role-based access control (RBAC) rules and role bindings to be generated.

IBM Spectrum Protect Plus prerequisites

The IBM Spectrum Protect Plus server and the IBM Spectrum Protect Plus vSnap server must be provisioned and configured by the IBM Spectrum Protect Plus administrator:

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

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

  You must specify this account name in the SPP_ADMIN_USERNAME parameter in the `baas_options.sh` configuration file before you deploy Container Backup Support. The `baas_options.sh` file is in the installation directory. For instructions, see "Setting up the installation variables" on page 206.
- An IBM Spectrum Protect Plus instance must be deployed in a container environment or as a VMware virtual appliance. Network connectivity must exist to and from the target cluster. The IBM Spectrum Protect Plus Internet Protocol (IP) address and port number must be specified in the `baas-values.yaml` file before you deploy Container Backup Support. Only one port (443) can be specified for use with all IBM Spectrum Protect Plus instances.
- **Optional:** For copy backup and copy restore operations, the IBM Spectrum Protect Plus vSnap server must be provisioned and configured by the IBM Spectrum Protect Plus administrator. An IBM Spectrum Protect Plus vSnap instance must be deployed as a VMware virtual appliance and configured to store backups:
  - Network connectivity must exist to and from the target Kubernetes or OpenShift cluster and the IBM Spectrum Protect Plus vSnap instance.
  - If backups are encrypted at rest, ensure that enough capacity is allocated for encryption on the vSnap server.

Connectivity

Ensure that the following connectivity requirements are met:

- The Secure Shell (SSH) service is running on Kubernetes NodePort services.
- Firewalls are configured to allow IBM Spectrum Protect Plus to connect data mover containers by using SSH over the NodePort port range of the Kubernetes or OpenShift cluster. The NodePort service allows the specific port in the NodePort range to be determined by Kubernetes or OpenShift at run time.

- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment must be registered by using a Fully Qualified Domain Name (FQDN) name or Internet Protocol (IP) address.
- If FQDN names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.
- If FQDN is not available, you must add the server to the `/etc/hosts` file on the IBM Spectrum Protect Plus server by using the command line.

Authentication and privileges

- Specify the username for the IBM Spectrum Protect Plus administrator with the containers role in the `baas_options.sh` configuration file. For more information, see "Setting up the installation variables" on page 206.
- The data mover runs as a privileged container to access the device location on the host system of the volume that is being protected. The application agent also runs as a privileged container to gain access to the `sudo` command to set up the data mover user account in the container at run time. The application agent accesses no host resources.
- Depending on their role, enterprise application developers and backup administrators interact with different user interfaces to protect persistent data in containers, as described in "User roles" on page 406.

Prerequisites and operations

Prerequisites

Ensure that the "Software" on page 78, "Connectivity" on page 79, "Authentication and privileges" on page 80 requirements are met before you start to install Container Backup Support on a Kubernetes or Red Hat OpenShift cluster as described in Chapter 6, "Installing Container Backup Support," on page 197.

Operations

Before you start a backup or restore operation, ensure that your system meets the following requirements:

- After Container Backup Support is installed, the application host for the Container Backup Support container is automatically registered upon startup of the cluster host in Kubernetes or OpenShift. When a cluster is registered with IBM Spectrum Protect Plus, an inventory of the resources in the cluster is automatically captured, enabling you to complete backup and restore jobs and to run reports. If the automatic registration is not successful and your cluster does not appear in the IBM Spectrum Protect Plus user interface, you must manually register the cluster. For instructions, see "Registering a Kubernetes cluster" on page 409 or "Registering an OpenShift cluster" on page 431.
- If you do not plan to use the default SLA policy for containers, ensure that you configure an SLA policy. For instructions, see "Creating an SLA policy for containers" on page 329.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the **Accounts** pane.

Review the following information about creating backup and restore jobs:

- You can use the IBM Spectrum Protect Plus user interface to back up or restore Kubernetes persistent volumes, namespace-scoped resources, and cluster-scoped resources. For instructions, see "Backing up and restoring Kubernetes clusters" on page 408.
- You can use the IBM Spectrum Protect Plus user interface to backup or restore OpenShift resources such as persistent volumes, project-scoped resources, and cluster-scoped resources. For instructions, see "Backing up and restoring OpenShift clusters" on page 430.

For an overview about protecting containers with IBM Spectrum Protect Plus, see Chapter 13, "Protecting containers," on page 401.

Ports

The following ports are used by IBM Spectrum Protect Plus agents.

Table 31. Communication ports when the target is an IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
Assigned by the NodePort service in Kubernetes	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus server	Kubernetes or OpenShift agent	Used by IBM Spectrum Protect Plus to connect to the data mover container to deploy and run agents **Note**: By default, port 30001 is used for SSH connections to the Kubernetes or OpenShift agent containers. This port is configurable and is port-forwarded to port 22. This port is used only when a containerized IBM Spectrum Protect Plus server opens an SSH connection to the Kubernetes or OpenShift agent container. SSH connections are never used within the Container Backup Support containers.

Table 32. Communication ports when the initiator is the IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
111	TCP and User Datagram Protocol (UDP)	Kubernetes or OpenShift agent	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations
443	TCP	Kubernetes or OpenShift agent	IBM Spectrum Protect Plus server	Used for IBM Spectrum Protect Plus issued commands to run backup, restore, inventory, and other operations

Table 32. Communication ports when the initiator is the IBM Spectrum Protect Plus agent (continued)				
Port	**Protocol**	**Initiator**	**Target**	**Description**
2049	TCP and UDP	Kubernetes or OpenShift agent	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations
20048	TCP and UDP	Kubernetes or OpenShift agent	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations

Hardware

The required system resources are based on the default installation parameters. By default, when you use the Helm chart for installation, you start with the containers and required resources that are listed in the table.

Table 33. Minimum resource requirements for Container Backup Support					
Component	**Replica**	**CPU (request)**	**CPU (limit)**	**Memory (request)**	**Memory (limit)**
baas-spp-ageent	1	2	3	800Mi	1000Mi
baas-cert-monitor*	1	250m	1	50Mi	250Mi
baas-datamover	1	100m	500m	500Mi	1000Mi
baas-kafka	1	300m	2	400Mi	1Gi
baas-scheduler	1	100m	750m	150Mi	500Mi
baas-controller	1	250m	1	50Mi	250Mi
baas-MinIO	1	100m	3	600Mi	3Gi
baas-transaction-manager	3	200m	1	100Mi	500Mi
baas-transaction-manager-worker	3	200m	2	250Mi	500Mi
baas-transaction-manager-redis	3	50m	200m	50Mi	250Mi

Table 33. Minimum resource requirements for Container Backup Support (continued)					
Component	**Replica**	**CPU (request)**	**CPU (limit)**	**Memory (request)**	**Memory (limit)**
baas-strimzi-cluster-operator (Kubernetes environment) amq-streams-cluster-operator (OpenShift environment)	1	200m	1	384Mi	384Mi
baas-entity-operator	1	300m	2	400Mi	1Gi
baas-zookeeper	1	300m	2	400Mi	1Gi
* Applicable only in a Kubernetes environment.					

Note: Beginning with IBM Spectrum Protect Plus V10.1.8 the baas-entity-operator is a requirement for Kubernetes and OpenShift environment.

Tip: The CPU resource is measured in Kubernetes *cpu* units. Memory is specified in units of bytes. For more information about CPU units and memory, see Managing Resources for Containers.

Db2 requirements

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

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861

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

General

Beginning with IBM Spectrum Protect Plus V10.1.2, support was added for backing up and restoring Db2 database data.

Before you register a Db2 database with IBM Spectrum Protect Plus, ensure that the system environment meets the following requirements.

Configuration requirements

Application versions

Table 34. Coverage matrix for application levels supported by IBM Spectrum Protect Plus			
IBM Spectrum Protect Plus	**Db2 V10.5*** Enterprise Edition	**Db2 V11.1*** Enterprise Edition	**Db2 V11.5*** Enterprise Edition
V10.1.2	✓	✓	--
V10.1.3	✓	✓	--

Table 34. Coverage matrix for application levels supported by IBM Spectrum Protect Plus (continued)			
V10.1.4	✓	✓	--
V10.1.5	✓	✓	✓
V10.1.6	✓	✓	✓
V10.1.7	✓	✓	✓
V10.1.8	✓	✓	✓
* The base release and later maintenance and modification levels are supported.			

Operating systems

Table 35. Coverage matrix for supported operating systems on IBM PowerPC		
IBM Spectrum Protect Plus	**IBM AIX 7.1***	**IBM AIX 7.2***
V10.1.2	✓	✓
V10.1.3	✓	✓
V10.1.4	✓	✓
V10.1.5	✓	✓
V10.1.6	✓	✓
V10.1.7	✓	✓
V10.1.8	✓	✓
* The base release and later maintenance and modification levels are supported.		

Table 36. Coverage matrix for supported Linux® x86_64 operating systems						
IBM Spectrum Protect Plus	**RHEL 6.8***	**RHEL 7.0***	**RHEL 8.0***	**SLES 11.0 SP4***	**SLES 12.0 SP1***	**SLES 15.0***
V10.1.2	✓	✓	--	✓	✓	--

Table 36. Coverage matrix for supported Linux® x86_64 operating systems (continued)						
V10.1.3	✓	✓	--	✓	✓	--
V10.1.4	✓	✓	--	✓	✓	--
V10.1.5	✓	✓	--	✓	✓	--
V10.1.6	✓	✓	--	✓	✓	--
V10.1.7	✓	✓	--	✓	✓	✓
V10.1.8	✓	✓	✓	✓	✓	✓
*The base release and later maintenance and modification levels are supported.						

Table 37. Coverage matrix for supported operating systems on Linux on Power Systems (little endian)				
IBM Spectrum Protect Plus	**RHEL 7.1***	**RHEL 8.0***	**SLES 12.0 SP1***	**SLES 15***
V10.1.4	✓	--	✓	--
V10.1.5	✓	--	✓	--
V10.1.6	✓	--	✓	--
V10.1.7	✓	--	✓	--
V10.1.8	✓	✓	✓	✓
*The base release and later maintenance and modification levels are supported.				

Restrictions

- The IBM Db2 pureScale® feature is not supported.
- Ensure that your Db2 logical volume setup does not include nested mount points.
- Ensure that the protected Db2 database does not share the same local database directory with other Db2 databases.

  Note: The snapshot backup will fail if an Db2 database is not unique within within its volume directory, this an DB2® restriction.
- If you plan to protect multiple partitions, Db2 must be in parallel backup mode. Parallel backup mode can be enabled by editing Db2 registry variables. For more information, see Prerequisites for Db2. The **DB2_PARALLEL_ACS** registry variable is available only in certain fix pack levels of Db2. If

the **DB2_PARALLEL_ACS** variable is not available in your version, you can meet the requirement by specifying **DB2_WORKLOAD** = **SAP**.

Software

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

  Tip: The required bash and sudo packages are included in the supported Linux86_64 and Linux Power® Systems (little endian) operating system installation packages.
- Install the most recent Db2 patches and updates in your environment.
- Ensure that the supported version of Linux x86_64, Linux Power Systems (little endian), or AIX is installed with the most recent patches and updates.
- The International Components for Unicode (`libicu`) RPM package corresponding to the operating system must be installed.
- Ensure that the user limit value `ulimit -f` is set to `unlimited`, for the IBM Spectrum Protect Plus agent user and the Db2 instance user. Alternatively, set the value sufficiently high to allow copying of the largest database files in your backup and restore jobs. If you change the `ulimit` setting, restart the Db2 instance to finalize the configuration.
- In a Linux environment, depending on your version or distribution, ensure that the Linux utility package `util-linux-ng` or `util-linux` is current.
- **RHEL users:** To ensure that the `util-linux-ng` or `util-linux` package is current, run the following command: `yum update package_name`, where package_name is the name of the Linux utility package.

Connectivity

Ensure that your system environment meets the following connectivity requirements:

- The secure file transfer protocol (SFTP) subsystem for Secure Shell (SSH) is enabled.
- The SSH service is running on port 22 on the proxy host server.
- Firewalls are configured to allow IBM Spectrum Protect Plus to connect to the Db2 server by using SSH.
- IBM Spectrum Protect Plus uses the Network File System (NFS) protocol to mount storage volumes for backup and restore operations:
  - On Linux, ensure that the native Linux NFS client is installed on the Db2 server.
  - On AIX, ensure that NFS communication is configured with reserved ports by using the following command:

    ```
    nfso -p -o nfs_use_reserved_port=1
    ```
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment must be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.
- If DNS is not available, you must add the server to the `/etc/hosts` file on the IBM Spectrum Protect Plus by using the command line.

Authentication and privileges

Authentication

- The Db2 server must be registered with IBM Spectrum Protect Plus by using an operating system user who exists on the Db2 server. The user is then referred to as the *IBM Spectrum Protect Plus agent user*.
- Ensure that the root user password is correctly configured and that the user can log in without other prompts, such as prompts to reset the password.

Privileges

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

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

Prerequisites and Operations

Prerequisites

Ensure that the Software, Connectivity, and Authentication and privileges requirements are met.

The following prerequisites must be met before you start protecting your resources.

- Db2 archive logging is activated and Db2 is in recoverable mode.
- Sufficient space is available in the Db2 database management system, in the volume groups for the backup operation, and on the target volumes for copying files during the restore operation. For more information about space requirements, see Space requirements for Db2 protection.
  - Before you back up Db2 databases, ensure you have enough free disk space on the target and source hosts, and in the vSnap repository. Extra free disk space is required in the volume groups on the source host for creating temporary Logical Volume Manager (LVM) snapshots of the logical volumes that the Db2 database and log files are stored on. To create LVM snapshots of a protected Db2 database, ensure that the volume groups with Db2 data have sufficient free space.
    - For AIX, no more than 15 snapshots can exist for each Enhanced Journaled File System (JFS2). Internal and external JFS2 snapshots cannot exist at the same time for the same file system. Ensure that no internal snapshots exist on the JFS2 volumes, as these snapshots can cause issues when the Db2 agent is creating external snapshots.
    - For every LVM or JFS2 snapshot logical volume that contains data, allow at least 10% of its size as free disk space in the volume group. If the volume group has enough free disk space, the IBM Spectrum Protect Plus Db2 agent reserves up to 25% of the source logical volume size for the snapshot logical volume.
  - When you are restoring data to an alternative location, allocate extra dedicated volumes for copy and restore processes. The data paths for table spaces and logs on the target host are the same as the paths on the original host. This setup supports the copying of data from the mounted vSnap server to the target host. Ensure that protected Db2 database does not share the same local database directory with other Db2 databases.
- Logical volumes holding Db2 table spaces (data and temporary table spaces), the local database directory, and Db2 log files are managed by the Logical Volume Management system (LVM2) on Linux or by JFS2 on AIX. LVM2 on Linux and JFS2 on AIX are used for creating temporary volume

snapshots. As long as the snapshot exists, the logical volume grows in size as the data changes on the source volume. For more information, see LVM2 and JFS2.

Operations

Before you start a backup or restore operation, take the following actions:

- Ensure that the application servers that host the Db2 databases that you want to back up are registered with IBM Spectrum Protect Plus.
- A service level agreement (SLA) policy is configured.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the **Accounts** pane.

Review the following information about creating backup and restore jobs:

- Define regularly scheduled Db2 backup jobs to protect your data. You can also enable continuous backup operations for archive logs so that you can restore a point-in-time copy with rollforward options if required. For instructions, see Backing up Db2 data.
- To restore Db2 data from the vSnap repository, define a job that restores data from either the newest backup or an earlier backup copy. You can restore data to the original instance or to an alternative instance on a different client host. For instructions, see Restoring Db2 data.

Review the following information about creating backup and restore jobs:

Ports

The following ports are used by IBM Spectrum Protect Plus agent users.

Table 38. Communication ports when the target is an IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
22	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus server	Db2	Provides access to troubleshoot and maintain remote proxy host servers running guest application components by using the SSH protocol

Table 39. Communication ports when the initiator is the IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
111	TCP and User Datagram Protocol (UDP)	Db2	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations
2049	TCP and UDP	Db2	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations

Table 39. Communication ports when the initiator is the IBM Spectrum Protect Plus agent (continued)				
Port	**Protocol**	**Initiator**	**Target**	**Description**
20048	TCP and UDP	Db2	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations

Hardware

Table 40. Minimum hardware requirements	
System	**Disk Space**
Compatible hardware that is supported by the operating system and Db2 database server	A minimum of 500 MB under `/tmp` and 800 MB under `/opt` is required for product installation.

Microsoft Exchange Server requirements

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

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861

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

General

Beginning with IBM Spectrum Protect Plus V10.1.3, support was added for backing up and restoring of Microsoft Exchange Server data.

Before you register a Microsoft Exchange Server with IBM Spectrum Protect Plus, ensure that the system environment meets the following requirements.

Configuration

Application versions

Table 41. Coverage matrix for application levels supported by IBM Spectrum Protect Plus			
IBM Spectrum Protect Plus	**Microsoft Exchange Server 2013 CU16*** **Standard and Enterprise editions**	**Microsoft Exchange Server 2016 CU5*** **Standard and Enterprise editions**	**Microsoft Exchange Server 2019*** **Standard and Enterprise editions**
V10.1.3	✓	✓	✓
V10.1.4	✓	✓	✓
V10.1.5	✓	✓	✓

Table 41. Coverage matrix for application levels supported by IBM Spectrum Protect Plus (continued)			
IBM Spectrum Protect Plus	**Microsoft Exchange Server 2013 CU16*** **Standard and Enterprise editions**	**Microsoft Exchange Server 2016 CU5*** **Standard and Enterprise editions**	**Microsoft Exchange Server 2019*** **Standard and Enterprise editions**
V10.1.6	✓	✓	✓
V10.1.7	✓	✓	✓
V10.1.8	✓	✓	✓
* The base release and later cumulative updates and maintenance levels are supported.			

Microsoft Exchange Server database availability groups (DAGs) are supported.

Operating Systems

Table 42. Coverage matrix for supported Windows x64 operating systems			
IBM Spectrum Protect Plus	**Microsoft Windows Server 2012 R2*** **Standard and Datacenter editions**	**Microsoft Windows Server 2016*** **Standard and Datacenter editions**	**Microsoft Windows Server 2019*** **Standard and Datacenter editions**
V10.1.3	✓	✓	✓
V10.1.4	✓	✓	✓
V10.1.5	✓	✓	✓
V10.1.6	✓	✓	✓
V10.1.7	✓	✓	✓
V10.1.8	✓	✓	✓
* The base release and later maintenance levels are supported.			

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

- VMware Elastic Sky X (ESX) guest operating system
- Microsoft Windows Hyper-V guest operating system

To enable write range tracking, see the minimum requirements in “Incremental backups” on page 93.

Restrictions

- Windows Server 2019 with the Server Core option is supported. However, the granular restore feature is not supported by the Server Core installation option.
- The database logs are backed up on the preferred node only. Only one Exchange Server instance at a time can write log backups to the vSnap server.
- When you restore a mailbox item (or an entire mailbox) to an Outlook personal folders (.pst) file, you can use the Mailbox Restore Browser view only with non-Unicode .pst files.
- When you restore a mailbox item (or a mailbox) to a different mailbox, you cannot drag mail items or subfolders in the Recoverable Items folder to a destination mailbox.
- When you restore mail items to a non-Unicode personal folders (.pst) file, each folder can contain a maximum of 16,383 mail items.
- Only one application server or file server can be assigned per host. For example, if a host is registered as a Microsoft Windows file system, you cannot register the same host as a Microsoft SQL Server or a Microsoft Exchange Server.

For restrictions on changed bytes tracking, see "Incremental backups" on page 93.

Software

- Install the most recent Microsoft Exchange Server database patches and updates in your environment.
- A supported version of a Microsoft Windows 64-bit operating system, including the most recent patches and updates, must be installed in your system environment.
- The following software must be installed before you use IBM Spectrum Protect Plus:
  - Windows PowerShell 4 or later
  - Windows Management Framework 4 or later
- If you use Microsoft Exchange Server 2013 with the granular restore feature, the minimum level that is supported for Microsoft Exchange Messaging API (MAPI) Client and Collaboration Data Objects (CDO) is version 6.5.8320.0.
- If you use the granular restore feature with Microsoft Exchange Server 2016 or 2019, Microsoft 32-bit Outlook 2013, Outlook 2016, or Outlook 2019 is required.
- The following software, required by Microsoft, is installed automatically by the IBM Spectrum Protect Plus granular restore feature, if not already present on your virtual machine:
  - 32-bit Microsoft Visual C++ 2012 Redistributable Package
  - 64-bit Microsoft Visual C++ 2012 Redistributable Package
  - 32-bit Microsoft Visual C++ 2017 Redistributable Package
  - 64-bit Microsoft Visual C++ 2017 Redistributable Package
  - Microsoft .NET Framework 4.5
  - Microsoft ReportViewer 2012 SP1 Redistributable Package
  - Microsoft SQL Server 2012 System CLR Types
  - Microsoft SQL Server 2014 System CLR Types
  - Microsoft SQL Server 2016 System CLR Types

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

Connectivity

Ensure that your system environment meets the following connectivity requirements:

- The network adapter that is used for the connection must be configured as a client for Microsoft Networks.

- The Microsoft Windows Remote Management (WinRM) service must be running.
- Firewalls must be configured to enable IBM Spectrum Protect Plus to connect to the server by using WinRM.
- Firewalls must be configured to enable the Exchange Server to communicate with the IBM Spectrum Protect Plus server by using Hypertext Transfer Protocol Secure (HTTPS) via port 443.
- The IP address of any client host that you register must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. The Microsoft Exchange Server must have a WinRM service that is listening on port 5985.
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment must be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and from the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.

Authentication and privileges

Authentication

Register each Microsoft Exchange Server with IBM Spectrum Protect Plus by specifying a fully qualified name or IP address. Ensure that the "Connectivity" on page 91 requirements are met.

The user identity must have sufficient privileges to install and start the IBM Spectrum Protect Plus Tools Service on the node. For more information, see the Microsoft article: Add the Log on as a service Right to an Account

Privileges

To use an Exchange Server database, an IBM Spectrum Protect Plus agent user must have the appropriate privileges. For instructions about assigning privileges, see Exchange Server Privileges.

Review the following information about privileges and restrictions:

- To manage Exchange Server role groups by using the Exchange Admin Center (EAC) or Exchange Powershell Cmdlets, the username must be authorized by the security policy.
- The Encrypting File System (EFS) must be enabled in the local or group domain policy, and a valid Domain Data Recovery Agent (DRA) certificate must be available.
- To use the mailbox browser for granular restore operations, Exchange digital certificates must be installed and configured.

  Tip: With Microsoft Exchange Server 2016 and 2019, the Exchange Server is configured to use Transport Layer Security (TLS) by default. The TLS protocol encrypts communication between internal Exchange servers, and between Exchange services on the local server.

Group Policy Object

For the Network security: LAN Manager authentication level policy setting, click **Computer Configuration** > **Windows Settings** > **Security Settings** > **Local Policies** > **Security Options** and specify one of the following options:

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

Important: For vSnap, you must use one of the security options listed above. Other options are not compatible with the vSnap Common Internet File System (CIFS) and the Server Message Block (SMB) protocol versions, and can cause CIFS authentication problems.

Specify the Group Policy Object (GPO) setting by navigating to **Computer Configuration** > **Windows Settings** > **Security Settings** > **Local Policies** > **Security Options** > **Network security: Restrict NTLM: Incoming NTLM traffic** and **Network security: Restrict NTLM: Outgoing NTLM traffic to remote servers**.

For the NTLM traffic, specify one of the following options:

- **Allow all**
- **Allow all accounts**

Prerequisites and operations

Prerequisites

Ensure that the "Software" on page 91, "Connectivity" on page 91, and "Authentication and privileges" on page 92 requirements are met.

Operations

Before you start a backup or restore operation, take the following actions:

- Ensure that the application servers that contain the Exchange Server databases that you want to back up are registered with IBM Spectrum Protect Plus.
- A service level agreement (SLA) policy is configured.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the **Accounts** pane.

Review the following information about creating backup and restore jobs:

- To protect Exchange Server databases, you can define a backup job that runs continuously to create incremental backups. You can also run on-demand backup jobs. Review the information in "Backing up Exchange databases" on page 527
- IBM Spectrum Protect Plus provides a backup strategy called incremental forever. Rather than scheduling periodic full backup jobs, this backup solution requires only one initial full backup. Afterward, an ongoing sequence of incremental backup jobs occurs. For detailed requirements and restrictions that apply to backup jobs, see section "Incremental backups" on page 93.
- If data in an Exchange Server database is lost or corrupted, you can restore the data from a backup copy. Use the Restore wizard to set up a restore job schedule or an on-demand restore operation. You can define a job that restores data to the original instance or to an alternative instance. Various recovery options and configurations are available as described in "Restoring Exchange databases " on page 530
- You can access the Exchange database files by using the instant access restore type and mount the database files from the vSnap volume to an application server.

For an overview about protecting Exchange Server databases with IBM Spectrum Protect Plus, see Protecting Exchange Server.

Incremental backups

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

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

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

The following technologies are not supported for changed bytes tracking:

- Resilient File System (ReFS)
- Server Message Block (SMB) 3.0 protocol

- SMB Transparent Failover (TFO)
- SMB 3.0 with Scale Out file shares

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

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

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

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

Ports

The following ports are used by IBM Spectrum Protect Plus agent users.

Table 43. Communication ports when the target is an IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
5985	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus server	Microsoft Exchange Server	Provides access to the Microsoft WinRM service for Windows-based servers
5986	TCP	IBM Spectrum Protect Plus server	Microsoft Exchange Server	Provides access to the Microsoft WinRM service for Windows-based servers

Table 44. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user

Port	Protocol	Initiator	Target	Description
3260	TCP	Microsoft Exchange Server	vSnap server	Used for Microsoft Internet Small Computer System Interface (iSCSI) data transfer to and from LUNs mounted from vSnap servers during backup and restore operations. To enable this functionality, the Microsoft Microsoft iSCSI Initiator Service is required on the specified node

Table 44. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user (continued)				
Port	**Protocol**	**Initiator**	**Target**	**Description**
443	TCP	Microsoft Exchange Server	IBM Spectrum Protect Plus server	Port that allows the agent to communicate with IBM Spectrum Protect Plus to send alerts in case of log backup failures
445	TCP	Microsoft Exchange Server agent	vSnap server	Used for SMB or CIFS data transfer to and from file systems mounted from vSnap servers during transaction log backup and restore operations

Ports update:

- For Microsoft Exchange Server, port 443 is available in IBM Spectrum Protect Plus V10.1.4 and later.
- In earlier versions, ports 137, 138, and 139 on the vSnap server were used by application agents that use SMBv1. Beginning with IBM Spectrum Protect Plus V10.1.6, the SMBv1 protocol is not used. All agents use SMBv2 or later, which does not require ports 137, 138, or 139.

Hardware

Table 45. Minimum hardware requirements		
System	**Disk space**	**Memory**
Compatible hardware that is supported by the 64-bit operating system and Microsoft Exchange Server	A minimum of 500 MB of disk space is required for product installation Additional 600MB disk space for dump files is required. The default dmp path is in dump sub folder of agent installation directory, for example: `C:\Program Files\IBM\IBM Spectrum Protect Plus\tools\dumps`. **For granular restore operations:** At least 2.1 GB of disk space is required for mandatory Microsoft software, which is installed automatically	16 GB Random Access Memory (RAM)

MongoDB requirements

Beginning with IBM Spectrum Protect Plus V10.1.3, support was added for backing up and restoring of MongoDB database data. Before you register a MongoDB application server with IBM Spectrum Protect Plus, ensure that the system environment meets the following requirements.

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861

General

Beginning with IBM Spectrum Protect Plus V10.1.3, support was added for backing up and restoring of MongoDB database data.

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

Configuration

Application versions

Table 46. Coverage matrix for application levels supported by IBM Spectrum Protect Plus

IBM Spectrum Protect Plus	MongoDB V3.6* Community Server and Enterprise Server editions	MongoDB V4.0* Community Server and Enterprise Server editions	MongoDB V4.2* Community Server and Enterprise Server editions	MongoDB V4.4* Community Server and Enterprise Server editions
V10.1.3	✓	✓	--	--
V10.1.4	✓	✓	--	--
V10.1.5	✓	✓	--	--
V10.1.6	✓	✓	✓	✓
V10.1.7	✓	✓	✓	✓
V10.1.8	✓	✓	✓	✓

* The base release and later maintenance and modification levels are supported.

Operating systems

Table 47. Coverage matrix for supported Linux x86_64 operating systems

IBM Spectrum Protect Plus	RHEL 6.8*	RHEL 7.0*	RHEL 8.0*	CentOS 6.8*	CentOS 7.0*	CentOS 8.0*	SLES 12.0 SP1*	SLES 15.0
V10.1.3	✓	✓	--	✓	✓	--	✓	--
V10.1.4	✓	✓	--	✓	✓	--	✓	--
V10.1.5	✓	✓	--	✓	✓	--	✓	--
V10.1.6	✓	✓	--	✓	✓	--	✓	--
V10.1.7	✓	✓	✓	✓	✓	✓	✓	✓
V10.1.8	✓	✓	✓	✓	✓	✓	✓	✓

* The base release and later maintenance and modification levels are supported.

Table 48. Coverage matrix for supported operating systems on Linux on Power Systems (little endian)

IBM Spectrum Protect Plus	RHEL 7.1*	RHEL 8.0*	CentOS 7.0*	CentOS 8.0*
V10.1.4	✓	--	✓	--
V10.1.5	✓	--	✓	--
V10.1.6	✓	--	✓	--
V10.1.7	✓	✓	✓	✓
V10.1.8	✓	✓	✓	✓

* The base release and later maintenance and modification levels are supported.

IBM Spectrum Protect Plus can also protect the MongoDB application server installed on a VMware or Kernel-based Virtual Machine (KVM) when MongoDB is running on a supported operating system.

Restrictions

- On Linux on Power Systems (little endian), only the MongoDB Enterprise Server Edition is supported.
- MongoDB shared cluster configurations are detected when you run an inventory, but these resources are not eligible for backup or restore operations.

- On MongoDB, Transport Layer Security (TLS)-based encryption and certificate-based authentication are not supported.
- Ensure that your MongoDB setup does not include nested mount points.

Software

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

  Tip: The required bash and sudo packages are included in the supported Linux x86_64 and Linux on Power Systems (little endian) operating system installation packages.
- Install the most recent MongoDB patches and updates in your environment.
- Ensure that a supported version of Linux x86_64 or Linux on Power Systems (little endian) is installed with the most recent patches and updates.
- The International Components for Unicode (**`libicu`**) RPM package corresponding to the operating system must be installed.
- Ensure that the user limit value `ulimit -f` is set to `unlimited`, for the IBM Spectrum Protect Plus agent user and the MongoDB instance user. Alternatively, set a sufficiently high value to support copying of the largest database files in your backup and restore jobs. If you change the `ulimit` setting, restart the MongoDB instance to finalize the configuration.
- In a Linux environment, depending on your version or distribution, ensure that the Linux utility package `util-linux-ng` or `util-linux` is current.
- **RHEL and CentOS 6 users:** To ensure that the `util-linux-ng` or `util-linux` package is current, run the following command:

  ```
  yum update package_name
  ```

  where package_name is the name of the Linux utility package.
- **RHEL and CentOS 6 users:** If the MongoDB application server is running on RHEL 6 or CentOS 6, ensure that the openSSL package is at version 1.0.1e-57 or later. To update the version, run the following command:

  ```
  yum update openssl
  ```

Connectivity

Ensure that your system environment meets the following connectivity requirements:

- The secure file transfer protocol (SFTP) subsystem for Secure Shell (SSH) is enabled.
- The SSH service is running on port 22 on the MongoDB server.
- Firewalls are configured to allow IBM Spectrum Protect Plus to connect to the proxy host server by using SSH.
- Firewalls must be configured to enable the MongoDB server to communicate with IBM Spectrum Protect Plus server using HTTPS via port 443.
- IBM Spectrum Protect Plus uses the Network File System (NFS) protocol to mount storage volumes for backup and restore operations. Ensure that the native Linux NFS client is installed on the proxy host server.
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment must be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.

- If DNS is not available, you must add the server to the `/etc/hosts` file on the IBM Spectrum Protect Plus server by using the command line.

Authentication and privileges

Authentication

- The MongoDB server must be registered with IBM Spectrum Protect Plus by using an operating system user that exists on the MongoDB server. The user is then referred to as the IBM Spectrum Protect Plus agent user.
- Ensure that the root user password is correctly configured and that the user can log in without facing any other prompts, such as prompts to reset the password.
- With the MongoDB Enterprise Server Edition, only the encrypted storage engine is supported.

Privileges

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

- Privileges to run commands as the root user and as a MongoDB software owner user by using sudo. IBM Spectrum Protect Plus requires these privileges for various tasks such as discovering storage layouts, mounting and unmounting disks, and managing databases.
  - The `sudoers` configuration must allow the IBM Spectrum Protect Plus agent user to run commands without a password.
  - The `!requiretty` setting must be specified, as described in "Setting sudo privileges" on page 564.
- Privileges to read the standard MongoDB server module `/usr/local/bin/mongod`. IBM Spectrum Protect Plus requires these privileges to use the PyMongo application programming interface (API) to connect to the MongoDB servers by using the DNS name that is assigned to the instance or the IP address and port number. This mechanism is used to gather information about MongoDB instances and databases.
- If the MongoDB server is protected by role-based authentication, you must set up the appropriate privileges, as described in "Roles for MongoDB" on page 562.

Prerequisites and operations

Prerequisites

Ensure that the "Software" on page 98, "Connectivity" on page 98, and "Authentication and privileges" on page 99 requirements are met.

The following prerequisites must be met before you start protecting your resources.

- The MongoDB is configured as a stand-alone instance or replica set. Backup operations for MongoDB sharded cluster instances are not supported. A backup job always includes all databases in the instance.
- The MongoDB instance is configured to use the WiredTiger Storage Engine.
- Each MongoDB instance to be protected must be registered with IBM Spectrum Protect Plus. After the instances are registered, IBM Spectrum Protect Plus runs an inventory to detect MongoDB resources. Ensure that all instances that you want to protect are detected and listed correctly.
- The user who is registered with IBM Spectrum Protect Plus for the MongoDB application server must be able to retrieve server information and status from the MongoDB admin database.
- Ensure that you have enough free space on the target and source hosts, and in the vSnap repository. Extra space is required to store temporary Logical Volume Manager (LVM) backups of logical volumes where the MongoDB data is located. These temporary backups, known as LVM snapshots, are created automatically by the MongoDB agent. For each LVM snapshot logical volume, at least 10% of free space must be allocated in the volume group. If the volume group has enough free space, the IBM Spectrum Protect Plus MongoDB agent reserves up to 25% of the source logical volume size for the snapshot logical volume. For more information, see "Space prerequisites for MongoDB protection" on page 563.

- Ensure that enough disk space is allocated at the target server for restore operations.
- Logical volumes of MongoDB data and log paths are managed by Linux Logical Volume Manager (LVM2). LVM2 is used to create temporary volume snapshots. The database files and the journal must be on a single volume. The logical volume grows in size with data as the data changes on the source volume while the snapshot exists. For more information, see “Linux LVM2 ” on page 563.

Operations

Before you start a backup or restore operation, take the following actions:

- Add the application servers that you want to back up.
- Configure a service level agreement (SLA) policy.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the **Accounts** pane.

Review the following information about creating backup and restore jobs:

- To regularly back up your data, define a backup job that includes an SLA policy. For instructions, see “Backing up MongoDB data” on page 568.
- To restore data, define a job that restores data to the latest backup or select an earlier backup copy. You can restore data to the original instance or to an alternative instance on a different machine, creating a cloned copy. Define and save the restore job to run as an ad hoc operation, or to run regularly as a scheduled job. For instructions, see “Restoring MongoDB data ” on page 573.

For an overview about protecting MongoDB databases with IBM Spectrum Protect Plus, see “MongoDB ” on page 561.

Ports

The following ports are used by IBM Spectrum Protect Plus agent users.

Table 49. Communication ports when the target is an IBM Spectrum Protect Plus agent

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

Table 50. Communication ports when the initiator is the IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
111	TCP	MongoDB	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.

Table 50. Communication ports when the initiator is the IBM Spectrum Protect Plus agent (continued)				
Port	**Protocol**	**Initiator**	**Target**	**Description**
443	TCP	MongoDB	IBM Spectrum Protect Plus	Port that allows the agent to communicate with IBM Spectrum Protect Plus for making representational state transfer application programming interface (REST API) calls
2049	TCP and UDP	MongoDB	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.
20048	TCP and UDP	MongoDB	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations.

Hardware

Table 51. Minimum hardware requirements	
System	**Disk Space**
Compatible hardware that is supported by the operating system and MongoDB.	A minimum of 500 MB under `/tmp` and 800 MB under `/opt` is required for product installation.

Microsoft 365 requirements

This document details the Microsoft 365 backup and restore requirements for IBM Spectrum Protect Plus Version 10.1.8.

General

Beginning with IBM Spectrum Protect Plus V10.1.5, support was added for backing up and restoring Microsoft 365 data.

Product name update: Microsoft Corporation announced new product names, effective 21st April 2020, for its Office 365 offerings for small and medium businesses. With this announcement, all small and medium business plans transitioned to the new Microsoft 365 brand. In IBM Spectrum Protect Plus V10.1.6, the user interface and documentation use the original product name, Office 365. For more information, see New Microsoft 365 offerings for small and medium-sized businesses.

If you choose to protect Microsoft 365 data with IBM Spectrum Protect Plus, you must purchase the IBM Spectrum Protect Plus for Microsoft 365 Entity ID Monthly License. For more information about this entitlement, see the IBM Spectrum Protect V10.1.5 announcement letter.

Before you start protecting Microsoft 365 data with IBM Spectrum Protect Plus, ensure that the system environment meets the following requirements.

Configuration

Cloud Service

To protect a Microsoft 365 application, you must register the application with Azure Active Directory and grant appropriate permissions. Before you begin, you must have the following items:

- An active Microsoft 365 subscription
- A Microsoft 365 administrative user ID and password

For instructions about registering the Microsoft 365 application with Azure Active Directory, see “Registering with Azure Active Directory ” on page 491.

If you have a Microsoft 365 administrative account, you can add users to ensure that they have valid licenses. For instructions, see Microsoft 365 in Visual Studio subscriptions.

Ensure that you remember the Microsoft 365 administrative user IDs and passwords or maintain this information in a secure location.

Application versions

Table 52. Coverage matrix for application levels supported by IBM Spectrum Protect Plus

IBM Spectrum Protect Plus	**Microsoft 365 Business** Basic, Business Standard, Business Premium editions	**Office 365 for Enterprise** E1, E3, and E5 editions	**Office365 for Education** A1, A3, and A5 editions	**Office 365 for Firstline Workers** F3 edition	**Microsoft 365 for Enterprise** E3 and E5 editions
	Former product name: Office 365 Business: Business, Essentials, and Business Premium editions		**Former product name**: Office 365 Education edition	**Former product name:** Microsoft 365 F1	
V10.1.5	✓	✓	✓	✓	✓
V10.1.6	✓	✓	✓	✓	✓
V10.1.7	✓	✓	✓	✓	✓
V10.1.8	✓	✓	✓	✓	✓

Operating systems

Table 53. Coverage matrix for supported Linux x86_64 operating systems				
IBM Spectrum Protect Plus	**RHEL 7.0***	**RHEL 8.0***	**CentOS 7.0***	**CentOS 8.0***
V10.1.5	✓	--	✓	--
V10.1.6	✓	✓	✓	--
V10.1.7	✓	✓	✓	✓
V10.1.8	✓	✓	✓	✓
* The base release and later maintenance and modification levels are supported.				

IBM Spectrum Protect Plus supports proxy host servers running on physical (bare metal) servers and in virtualized environments.

Restrictions

- The Microsoft 365 tenant must be in a global region as defined by Microsoft. National regions are not supported. For more information about regions, see National cloud deployments.
- Ensure that vSnap is not used as a proxy server in production mode.

Software

- Ensure that Java™ 8 is installed.
- The bash and sudo packages must be installed. Sudo must be at version 1.7.6p2 or later. Run sudo -V to check the version.

  Tip: The required bash and sudo packages are included in the supported Linux x86_64 operating system.
- Install the most recent Microsoft 365 patches and updates in your environment.
- Ensure that a supported version of Linux x86_64 is installed with the most recent patches and updates.
- The International Components for Unicode (libicu) RPM package must be installed for the corresponding version of your operating system.
- Ensure that the user limit value **`ulimit -f`** value, which specifies the effective file size for the IBM Spectrum Protect Plus agent, is set to `unlimited`. Alternatively, set the value sufficiently high to support copying of the largest Microsoft 365 files in your backup and restore jobs.
- In a Linux environment, depending on your version or distribution, ensure that the Linux utility package, `util-linux-ng` or `util-linux`, is current.

Connectivity

Ensure that your system environment meets the following connectivity requirements:

- The secure file transfer protocol (SFTP) subsystem for Secure Shell (SSH) is enabled.
- The SSH service is running on port 22 on the proxy host server.
- Firewalls are configured to allow IBM Spectrum Protect Plus to connect to the proxy host server by using SSH.
- Firewalls must be configured to enable the proxy host server to communicate with the IBM Spectrum Protect Plus server by using Hypertext Transfer Protocol Secure (HTTPS) via port 443.

- IBM Spectrum Protect Plus uses the Network File System (NFS) protocol to mount storage volumes for backup and restore operations. Ensure that the native Linux NFS client is installed on the proxy host server.
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment must be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.
- If DNS is not available, you must add the server to the /etc/hosts file on the IBM Spectrum Protect Plus server by using the command line.

Authentication and privileges

Authentication

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

Privileges

The IBM Spectrum Protect Plus agent user must have privileges to run commands a root user by using sudo. The **sudoers** configuration must allow the IBM Spectrum Protect Plus agent user to run commands without a password.

Prerequisites and operations

Prerequisites

Ensure that the “Software” on page 103,“Connectivity” on page 103, and “Authentication and privileges” on page 104 requirements are met.

The following prerequisites must be met before you start protecting your resources:

- To protect a Microsoft 365 application, you must register the application with Azure Active Directory and grant appropriate permissions. When you register a new application with Azure Active Directory, the application credentials such as application ID and application secret are made available on the Azure Active Directory portal. For instructions, see “Registering with Azure Active Directory ” on page 491.
- To ensure that the IBM Spectrum Protect Plus agent can connect to the Microsoft 365 tenant, you must register the Microsoft tenant credentials and the proxy host server with IBM Spectrum Protect Plus. For instructions, see “Registering the Microsoft 365 tenant with IBM Spectrum Protect Plus” on page 492.

Operations

Before you start a backup or restore operation, take the following actions:

- Apply a service level agreement (SLA) policy.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the Accounts pane.
- **Performance tip:** To help enhance the performance of backup operations, set the number of parallel sessions to a number in the range 10 - 40.

Review the following information about creating backup and restore jobs:

- To back up Microsoft 365 email, calendars, contacts, and data on OneDrive cloud storage, see Backing up Office 365 data.

- To restore Microsoft 365 data from backup copies on vSnap servers or remote storage, see Restoring Office 365 data.
- For an overview about protecting Microsoft 365 with IBM Spectrum Protect Plus, see "Microsoft 365" on page 491.

Ports

The following ports are used by IBM Spectrum Protect Plus agents users.

Table 54. Communication ports when the target is an IBM Spectrum Protect Plus agent user

Port	Protocol	Initiator	Target	Description
22	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus server	Proxy host server	Provides access to troubleshoot and maintain remote proxy host servers that are running guest application components by using the SSH protocol

Table 55. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user

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

Hardware

Table 56. Minimum hardware requirements		
System	**Disk Space**	**Memory and CPU**
Compatible hardware with quad-core processors that are supported by the operating system	A minimum of 500 MB under `/tmp` and 800 MB under `/opt` is required for product installation. 5 GB of available disk space for temporary files at run time	16 GB of random access memory (RAM) and 8 processors

Oracle Server database backup and restore requirements

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

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861

General

Beginning with IBM Spectrum Protect Plus V10.1.1, support was added for backing up and restoring Oracle Database data.

Before you register an Oracle Server with IBM Spectrum Protect Plus, ensure that the system environment meets the following requirements.

Configuration

Application versions

Table 57. Coverage matrix for application levels supported by IBM Spectrum Protect Plus					
IBM Spectrum Protect Plus	**Oracle 11g R2* Enterprise Edition**	**Oracle 12c R1* Enterprise Edition**	**Oracle 12c R2* Enterprise Edition**	**Oracle 18c* Enterprise Edition**	**Oracle 19c* Enterprise Edition**
V10.1.1	✓	✓	✓	--	--
V10.1.2	✓	✓	✓	--	--
V10.1.3	✓	✓	✓	✓	--
V10.1.4	✓	✓	✓	✓	--
V10.1.5	✓	✓	✓	✓	✓
V10.1.6	✓	✓	✓	✓	✓
V10.1.7	✓	✓	✓	✓	✓

Table 57. Coverage matrix for application levels supported by IBM Spectrum Protect Plus (continued)					
IBM Spectrum Protect Plus	**Oracle 11g R2* Enterprise Edition**	**Oracle 12c R1* Enterprise Edition**	**Oracle 12c R2* Enterprise Edition**	**Oracle 18c* Enterprise Edition**	**Oracle 19c* Enterprise Edition**
V10.1.8	✓	✓	✓	✓	✓
* The base release and later maintenance and modification levels are supported.					

Tip: For multitenant databases in Oracle 12c and later, IBM Spectrum Protect Plus supports protection and recovery of the container database, including all pluggable databases (PDBs) within the container database. You can recover specific PDBs by using the Oracle Recovery Manager (RMAN) with an Instant Disk Restore recovery operation.

Operating systems

Table 58. Coverage matrix for supported operating systems on IBM PowerPC			
IBM Spectrum Protect Plus	**IBM AIX 6.1 TL9***	**IBM AIX 7.1***	**IBM AIX 7.2***
V10.1.1	✓	✓	--
V10.1.2	✓	✓	--
V10.1.3	✓	✓	--
V10.1.4	✓	✓	--
V10.1.5	✓	✓	--
V10.1.6	✓	✓	--
V10.1.7	✓	✓	✓
V10.1.8	✓	✓	✓
* The base release and later maintenance and modification levels are supported.			

Table 59. Coverage matrix for supported Linux® x86_64 operating systems									
IBM Spectrum Protect Plus	**RHEL 6.5***	**RHEL 7.0***	**RHEL 8.0***	**CentOS 6.5***	**CentOS 7.0***	**CentOS 8.0***	**SLES 11.0 SP4***	**SLES 12.0 SP1***	**SLES 15.0***
V10.1.1	✓	✓	--	✓	✓	--	✓	✓	--

Table 59. Coverage matrix for supported Linux® x86_64 operating systems (continued)									
IBM Spectrum Protect Plus	**RHEL 6.5***	**RHEL 7.0***	**RHEL 8.0***	**CentOS 6.5***	**CentOS 7.0***	**CentOS 8.0***	**SLES 11.0 SP4***	**SLES 12.0 SP1***	**SLES 15.0***
V10.1.2	✓	✓	--	✓	✓	--	✓	✓	--
V10.1.3	✓	✓	--	✓	✓	--	✓	✓	--
V10.1.4	✓	✓	--	✓	✓	--	✓	✓	✓
V10.1.5	✓	✓	--	✓	✓	--	✓	✓	✓
V10.1.6	✓	✓	✓	✓	✓	✓	✓	✓	✓
V10.1.7	✓	✓	✓	✓	✓	✓	✓	✓	✓
V10.1.8	✓	✓	✓	✓	✓	✓	✓	✓	✓
* The base release and later maintenance and modification levels are supported.									

Restrictions

- Oracle Data Guard is not supported.
- Databases must be in ARCHIVELOG mode. IBM Spectrum Protect Plus cannot protect databases running in NOARCHIVELOG mode.
- Oracle Database must be running to be protected by IBM Spectrum Protect Plus.
- Only multi-node Oracle Real Application Clusters (RAC not configured as RAC One Node) database is supported.
- Real Application Cluster (RAC) database recovery operations are not server pool-aware. IBM Spectrum Protect Plus can recover databases to a RAC, but not to specific server pools.
- RAC databases must be configured so that the RMAN Snapshot Control File location points to shared storage that is accessible to all cluster instances.
- When restoring an Oracle Database that was configured for multithreading at the time of backup, the restored database is non-multithreaded. The restored database must be manually reconfigured to use multithreading.
- Point-in-time recovery is not supported when one or more data files are added to the database in the period between the chosen point in time and the time when the preceding backup job ran.

Network File System (NFS)

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

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

For Direct NFS to work correctly, the executable file, *oracle_home*/bin/oradism, in each Oracle home directory must be owned by the root user and have setuid privileges. Typically, the binary file is preconfigured by the Oracle installer, but on certain systems, this binary file might not have the required privileges. To set the correct privileges, you must run appropriate commands. In the following examples, OINSTALL specifies the group that owns the installation, and ORACLE_HOME specifies the Oracle home directory.

- `chown root:oinstall oracle_home/bin/oradism`
- `chmod 750 oracle_home/bin/oradism`

Database discovery

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

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

If a database is set to automatically start, but the database is not running at the time of discovery, the discovery process will be partially completed. In this situation, databases that no longer exist might appear in the IBM Spectrum Protect Plus user interface. To resolve the issue, either start the database or set the autostart column to N for that database in the `/etc/oratab` file.

Block change tracking

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

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

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

Software

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

  Tip: The required **bash** and **sudo** packages are included in the supported Linux x86_64 operating system installation packages.
- Install the most recent Oracle Server patches and updates in your environment.
- Ensure that a supported version of Linux x86_64 or AIX is installed with the most recent patches and updates.
- The International Components for Unicode (`libicu`) RPM package must be installed for the corresponding version of your operating system.
- Ensure that the user limit value **ulimit -f** is set to `unlimited`, for the IBM Spectrum Protect Plus agent user and the Oracle instance user. Alternatively, set the value to a sufficiently high value to allow

copying of the largest database files in your backup and restore jobs. If you change the **ulimit** setting, restart the Oracle instance to finalize the configuration.

- In a Linux environment, depending on your version or distribution, ensure that the Linux utility package `util-linux-ng` or `util-linux` is current.
- RHEL and CentOS 6 users: To ensure that the `util-linux-ng` or `util-linux` package is current, run the following command:

```
yum update package_name
```

where *package_name* is the name of the Linux utility package.

Connectivity

Ensure that your system environment meets the following connectivity requirements:

- The secure file transfer protocol (SFTP) subsystem for Secure Shell (SSH) is enabled.
- The SSH service is running on port 22 on the Oracle Server.
- Firewalls are configured to allow IBM Spectrum Protect Plus to connect to the Oracle server by using SSH.
- Firewalls must be configured to enable the Oracle Server to communicate with IBM Spectrum Protect Plus server using HTTPS via port 443.
- IBM Spectrum Protect Plus uses the Network File System (NFS) protocol to mount storage volumes for backup and restore operations. Ensure that the native Linux NFS client is installed on the Oracle server.
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment must be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and from the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.
- If DNS is not available, you must add the server to the `/etc/hosts` file on the IBM Spectrum Protect Plus server by using the command line.
- Oracle RAC nodes are registered by their physical IP or name. Do not use a virtual name or Single Client Access Name (SCAN).
- The Oracle server must have a consistent fully qualified domain name (FQDN) or the user must set the **overrideHostname** flag to the desired FQDN in the `unixagent` section of the `/etc/guestapps.conf` file.

Authentication and privileges

Authentication

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

Privileges

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

- Privileges to run commands as the root user and as an Oracle software owner user (for example, `oracle` or `grid`) by using **sudo**. These privileges are required for tasks such as discovering storage layouts, mounting and unmounting disks, and managing databases and Automatic Storage Management (ASM).

  - The sudoers configuration must allow the IBM Spectrum Protect Plus agent user to run commands without a password.
  - The **!requiretty** setting must be specified.
  - The ENV_KEEP setting must allow the ORACLE_HOME and ORACLE_SID environment variables to be retained.
- Privileges to read the Oracle inventory. These privileges are required for tasks such as discovering and collecting information about Oracle home directories and databases.

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

For information about creating a user with the required privileges, see "Sample configuration of an IBM Spectrum Protect Plus agent user" on page 111.

Sample configuration of an IBM Spectrum Protect Plus agent user

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

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

```
useradd -m sppagent
```

- Set a password:

```
passwd sppagent_password
```

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

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

- Add the user to the Oracle installation and to the operating system (OSDBA) group:

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

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

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

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

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

Prerequisites and operations

Prerequisites

Ensure that the "Software" on page 109, "Connectivity" on page 110, and "Authentication and privileges" on page 110 requirements are met.

Operations

Before you start a backup or restore operation, take the following actions:

- Register the providers that you want to back up. For more information, see "Adding an Oracle application server" on page 588
- Configure a service level agreement (SLA) policy.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the **Accounts** pane.

Review the following information about creating backup and restore jobs:

- You can use backup jobs to back up Oracle environments with snapshots. Review the information in "Backing up Oracle data" on page 590
- You can use restore jobs to restore an Oracle environment from snapshots. IBM Spectrum Protect Plus creates a vSnap clone from the version that is selected during the job definition creation and creates a NFS share. The IBM Spectrum Protect Plus agent then mounts the share on the Oracle Server where the restore job is to be run. For Oracle Real Application Clusters (RAC), the restore job is run on all nodes in the cluster. Review the information in "Restoring Oracle data" on page 593

For an overview about protecting Oracle databases with IBM Spectrum Protect Plus, see "Backing up and restoring Oracle data" on page 588

Log backup

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

Ports

The following ports are used by IBM Spectrum Protect Plus agent users.

Table 60. Communication ports when the target is an IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
22	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus server	Oracle Server	Provides access to troubleshoot and maintain remote proxy host servers running guest application components by using the SSH protocol

Table 61. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user

Port	Protocol	Initiator	Target	Description
111	TCP and User Datagram Protocol (UDP)	Oracle Server	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations

Table 61. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user (continued)				
Port	**Protocol**	**Initiator**	**Target**	**Description**
443	TCP	Oracle Server	IBM Spectrum Protect Plus server	Port that allows the agent to communicate with IBM Spectrum Protect Plus for sending alerts in case of log backup failures
2049	TCP and UDP	Oracle Server	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations
20048	TCP and UDP	Oracle Server	vSnap server	Used for NFS data transfer to and from file systems mounted from vSnap servers during backup and restore operations

Hardware

Table 62. Minimum hardware requirements	
System	**Disk Space**
Compatible hardware that is supported by the operating system and Oracle Server	A minimum of 500 MB under `/tmp` and 800 MB under `/opt` is required for product installation.

Microsoft SQL Server database backup and restore requirements

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

To help ensure that backup and restore operations can be run successfully, your system must meet the hardware and software requirements. Use the following requirements as a starting point. For the most current requirements, which might include updates, see technote 304861

General

Beginning with IBM Spectrum Protect Plus V10.1.1, support was added for backing up and restoring Microsoft SQL Server data.

Before you register a Microsoft SQL Server with IBM Spectrum Protect Plus, ensure that the system environment meets the following requirements.

Configuration

Application versions

Table 63. Coverage matrix for application levels supported by IBM Spectrum Protect Plus

IBM Spectrum Protect Plus	Microsoft SQL Server 2008 R2 SP3* Standard and Enterprise editions	Microsoft SQL Server 2012* Standard and Enterprise editions	Microsoft SQL Server 2014* Standard and Enterprise editions	Microsoft SQL Server 2016* Standard and Enterprise editions	Microsoft SQL Server 2017* Standard and Enterprise editions	Microsoft SQL Server 2019* Standard and Enterprise editions
V10.1.1	✓	✓	✓	✓	✓ Beginning with V10.1.1 patch 1	--
V10.1.2	✓	✓	✓	✓	✓	--
V10.1.3	✓	✓	✓	✓	✓	--
V10.1.4	✓	✓	✓	✓	✓	--
V10.1.5	✓	✓	✓	✓	✓	✓ Beginning with V10.1.5 patch 1
V10.1.6	✓	✓	✓	✓	✓	✓
V10.1.7	✓	✓	✓	✓	✓	✓
V10.1.8	✓	✓	✓	✓	✓	✓

* The base release and later cumulative updates and maintenance levels are supported.

Operating systems

Table 64. Coverage matrix for supported Microsoft Windows 64-bit operating systems

IBM Spectrum Protect Plus	Microsoft Windows Server 2012 R2* Standard and Datacenter editions	Microsoft Windows Server 2016* Standard and Datacenter editions	Microsoft Windows Server 2019* Standard and Datacenter editions
V10.1.1	✓	✓	--
V10.1.2	✓	✓	--

Table 64. Coverage matrix for supported Microsoft Windows 64-bit operating systems (continued)			
IBM Spectrum Protect Plus	**Microsoft Windows Server 2012 R2* Standard and Datacenter editions**	**Microsoft Windows Server 2016* Standard and Datacenter editions**	**Microsoft Windows Server 2019* Standard and Datacenter editions**
V10.1.3	✓	✓	✓
V10.1.4	✓	✓	✓
V10.1.5	✓	✓	✓
V10.1.6	✓	✓	✓
V10.1.7	✓	✓	✓
V10.1.8	✓	✓	✓
* The base release and later maintenance levels are supported.			

Restrictions

- IBM Spectrum Protect Plus does not support log backup operations for simple recovery models.
- SQL Server does not support log backup operations for system databases and databases with the simple recovery model. You cannot use IBM Spectrum Protect Plus to back up the logs of the specified SQL Server databases.
- Failover of an SQL cluster instance during backup operations is not supported.
- The Volume Shadow Copy Service (VSS) restore file path is limited to 256 or fewer characters. If the original path exceeds 256 characters, consider using a customized file path for production restore jobs to reduce the length.
- Due to limitations of the VSS framework, leading spaces, trailing spaces, and unprintable characters should not be used in database names. For more information, see Backing up a SQL Server database using a VSS backup application may fail for some databases
- You cannot restore data to a New Technology File System (NTFS) or file allocation table (FAT) compressed volume because of SQL Server database restrictions. For more information, see Description of support for SQL Server databases on compressed volumes
- Microsoft SQL Server must be configured to use Windows Authentication, sometimes called trusted connections. For more information about SQL Server authentication modes and instructions to change the SQL Server authentication mode, see Choose an Authentication Mode
- Only one application server or file server can be assigned per host.

  For example, if a host is registered as a Microsoft Windows file system, you cannot register the same host as a Microsoft SQL Server or a Microsoft Exchange Server.

Software

- Install the most recent Microsoft SQL Server patches and updates in your environment.
- A supported version of a Microsoft Windows 64-bit operating system, including the most recent patches and updates, must be installed in your system environment.

- If the SQL Server is configured with Transport Layer Security (TLS) 1.2, a compatible version of the Open Database Connectivity (ODBC) driver must be installed on the SQL Server. For information about compatibility, see System Requirements, Installation, and Driver Files - SQL Server.

Connectivity

Ensure that your system environment meets the following connectivity requirements:

- The network adapter that is used for the connection must be configured as a client for Microsoft Networks.
- The Microsoft Windows Remote Management (WinRM) service must be running.
- Firewalls must be configured to enable the SQL Server to communicate with the IBM Spectrum Protect Plus server by using the Hypertext Transfer Protocol Secure (HTTPS) protocol via port 443.
- Firewalls must be configured to enable IBM Spectrum Protect Plus to connect to the server by using WinRM.
- The IP address of any client host that you register must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. The SQL Server must have a WinRM service that is listening on port 5985.
- All servers, proxies, applications, and hypervisors that are added to the IBM Spectrum Protect Plus environment must be registered by using a Domain Name System (DNS) name or Internet Protocol (IP) address.
- If DNS names are used, they must be resolvable over the network by the IBM Spectrum Protect Plus server and from the vSnap server. All IBM Spectrum Protect Plus components must also be resolvable by their DNS names.

Authentication and privileges

Authentication

Register each Microsoft SQL Server with IBM Spectrum Protect Plus by specifying a fully qualified name or IP address. When you register an SQL Server cluster node, register each node by name or IP address. Ensure that the “Connectivity” on page 116 requirements are met.

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

If the SQL Server is attached to a domain, the user identity follows the default `domain\Name` format. If the user is a local administrator, the user identity matches the name of the local administrator.

Privileges

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

- Microsoft SQL Server `public` and `sysadmin` permissions
- Windows local administration permissions, which are required by the VSS framework, and volume and disk access
- Permissions to access cluster resources in an SQL Server Always On and SQL Server failover clustering instance (FCI) environment

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

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

The target Microsoft SQL Server service account must have permissions to access Microsoft SQL Server restore files. For more information about requirements and restrictions, see the Administrative Considerations section in the Microsoft article: Securing Data and Log Files

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

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

This message is issued when a network access group policy setting is enabled. For instructions about disabling the setting, see the Microsoft Support article: Task Scheduler Error "A specified logon session does not exist"

Group Policy Object

For SQL log backup and restore operations, the vSnap server must be configured to use the NT LAN Manager V2 (NTLM V2) authentication protocol on the Windows server. You must specify the network security and Group Policy Object (GPO) settings that are described in this section.

For the Network security: LAN Manager authentication level policy setting, click **Computer Configuration** > **Windows Settings** > **Security Settings** > **Local Policies** > **Security Options** and specify one of the following options:

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

Important: For vSnap, you must use one of the security options listed above. Other options are not compatible with the vSnap Common Internet File System (CIFS) and the Server Message Block (SMB) protocol versions, and can cause CIFS authentication problems.

Specify the Group Policy Object (GPO) setting by navigating to **Computer Configuration** > **Windows Settings** > **Security Settings** > **Local Policies** > **Security Options** > **Network security: Restrict NTLM: Incoming NTLM traffic** and **Network security: Restrict NTLM: Outgoing NTLM traffic to remote servers**.

For the NTLM traffic, specify one of the following options:

- **Allow all**
- **Allow all accounts**

In addition to the configuration of the SQL Server, you must also enable the NTLM V2 authentication protocol at the enterprise level.

Prerequisites and operations

Prerequisites

Ensure that the "Software" on page 115, "Connectivity" on page 116, and "Authentication and privileges" on page 116 requirements are met.

The following prerequisites must be met before you start protecting your resources:

- An Internet Small Computer Interface (iSCSI) route must be enabled between the Microsoft SQL Server system and vSnap server. For more information, see iSCSI Initiator Step-by-Step Guide
- The Windows PowerShell binary path must be set in the %PATH% environment variable.
- Ensure that remote PowerShell invocation is enabled for the SQL Server by using Windows Remote Management (WinRM) from the IBM Spectrum Protect Plus server.
- If you plan to back up databases that were restored in test mode, use the global preference to limit the size of backup target volumes to less than 64 TB. You must set this global preference before you run the first backup for the service level agreement (SLA) that protects the databases. If the size of the backup target volumes is 64 TB or more, the backup job fails.

Operations

Before you start a backup or restore operation, take the following actions:

- Register the SQL Servers that you want to back up. When an SQL Server application server is added, an inventory of the instances and databases that are associated with the application server is captured and added to IBM Spectrum Protect Plus. The inventory is required for backup and restore jobs and to run reports.
- Configure service level agreement (SLA) policies.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the **Accounts** pane.

Review the following information about creating backup and restore jobs:

- You can use a backup job to back up SQL Server environments with snapshots. For instructions, see "Backing up SQL Server data" on page 602.
- You can use a restore job to restore a Microsoft SQL Server environment from snapshots. After you run IBM Spectrum Protect Plus Instant Disk Restore jobs, your SQL Server clones can be used immediately. IBM Spectrum Protect Plus catalogs and tracks all cloned instances, as described in "Restoring SQL Server data" on page 607.

For an overview about protecting SQL Server databases with IBM Spectrum Protect Plus, see "Backing up and restoring SQL Server data" on page 600.

In-Memory online transaction processing

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

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

Configuring Always On availability groups

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

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

When a secondary replica is preferred, and more than one secondary replica is available, the IBM Spectrum Protect Plus job executor selects a secondary replica depending on the backup priority and exclude replica options of the availability group.

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

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

Incremental backups

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

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

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

The following technologies are not supported for changed bytes tracking:

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

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

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

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

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

Log backups

To ensure that SQL log backup works properly, you might have to update the Windows Group Policy Object settings. For more information, see Group Policy Object.

Ports

The following ports are used by IBM Spectrum Protect Plus agent users.

Table 65. Communication ports when the target is an IBM Spectrum Protect Plus agent

Port	Protocol	Initiator	Target	Description
5985	Transmission Control Protocol (TCP)	IBM Spectrum Protect Plus server	Microsoft SQL Server	Provides access to the Microsoft WinRm service for Windows-based servers
5986	TCP	IBM Spectrum Protect Plus server	Microsoft SQL Server	Provides access to the Microsoft WinRm service for Windows-based servers

Table 66. Communication ports when the initiator is an IBM Spectrum Protect Plus agent user

Port	Protocol	Initiator	Target	Description
3260	TCP	Microsoft SQL Server	vSnap server	Used for Microsoft Internet Small Computer System Interface (iSCSI) data transfer to and from Logical Unit Numbers (LPNs) mounted from vSnap servers during backup and restore operations. To enable this functionality, the Microsoft iSCSI Initiator Service is required on the specified node
443	TCP	Microsoft SQL Server agent	IBM Spectrum Protect Plus server	Port that allows the agent to communicate with IBM Spectrum Protect Plus for SQL log backup, restore, and purge operations, and for sending alerts in case of log backup failures
445	TCP	Microsoft SQL Server agent	vSnap server	Used for SMB or CIFS data transfer to and from file systems mounted from vSnap servers during transaction log backup and restore operations

Ports update

- For Microsoft SQL Server, port 443 is available in IBM Spectrum Protect Plus V10.1.4 and later.
- In earlier versions, ports 137, 138, and 139 on the vSnap server were used by application agents that use SMBv1. Beginning with IBM Spectrum Protect Plus V10.1.6, the SMBv1 protocol is not used. All agents use SMBv2 or later, which does not require ports 137, 138, or 139.

Hardware

Table 67. Minimum hardware requirements

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

Post installation tasks

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

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

Assigning a static IP address

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

Procedure

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

1. Log on to the virtual machine console with the user ID `serveradmin`.

   The initial password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in "Start IBM Spectrum Protect Plus" on page 257.
2. From a CentOS command line, enter `nmtui` to open the interface.
3. From the main menu, select **Edit a connection**, and then click **OK**.
4. Select the network connection, then click **Edit**.
5. On the **Edit Connection** screen, enter an available static IP address that is not already in use.
6. Save the static IP configuration by clicking **OK**, then restart the IBM Spectrum Protect Plus appliance.

Related tasks

"Installing IBM Spectrum Protect Plus as a VMware virtual appliance" on page 126
To install IBM Spectrum Protect Plus in a VMware environment, deploy an Open Virtualization Format (OVF) template. Deploying an OVF template creates a virtual appliance containing the application on a VMware host such as an ESXi server.

"Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance" on page 128
To install IBM Spectrum Protect Plus in a Microsoft Hyper-V environment, import the IBM Spectrum Protect Plus for Hyper-V template. Importing a template creates a virtual appliance containing the IBM Spectrum Protect Plus application on a Hyper-V virtual machine.

Uploading the product key

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

Before you begin

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

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

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

Procedure

To upload a license file, complete the following steps:

1. In the IBM Spectrum Protect Plus user interface, click the user menu in the menu bar, and then click **Upgrade license**.
   The number of days until the trial license expires is shown.
2. Review the licensing information, and then click **Proceed to upload**.
3. Browse to select the license file, and then click **Upgrade license**.
4. When the license file is uploaded, close the upgrade notification window.

What to do next

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

Action	How to
Start IBM Spectrum Protect Plus from a supported web browser.	See "Start IBM Spectrum Protect Plus" on page 257.

Editing firewall ports

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

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

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

Use the following command to list the open ports:

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

Use the following command to list zones:

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

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

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

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

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

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

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

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

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

Use the following command to open a range of ports:

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

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

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

SUSE Linux Enterprise Server 12

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

Firewall configurations that use IP tables

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

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

Use the following command to list the current firewall policies:

```
sudo iptables -S
```

```
sudo iptables -L
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

```
sudo iptables-save
```

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

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

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

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

Related concepts

"Post installation tasks" on page 121

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

Overview of IBM Spectrum Protect Plus virtual appliance deployment

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

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

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

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

Obtaining the IBM Spectrum Protect Plus installation package

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

Before you begin

For the list of installation packages by component, and the links to the download site for the files, see technote 6415111.

Procedure

Download the appropriate installation file.

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

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

Related concepts

"Updating IBM Spectrum Protect Plus components" on page 239

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

Related tasks

"Installing IBM Spectrum Protect Plus as a VMware virtual appliance" on page 126

© Copyright IBM Corp. 2017, 2021

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

"Installing IBM Spectrum Protect Plus as a Hyper-V virtual appliance" on page 128
To install IBM Spectrum Protect Plus in a Microsoft Hyper-V environment, import the IBM Spectrum Protect Plus for Hyper-V template. Importing a template creates a virtual appliance containing the IBM Spectrum Protect Plus application on a Hyper-V virtual machine.

"Installing a vSnap server" on page 153
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.

Installing IBM Spectrum Protect Plus as a VMware virtual appliance

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

Before you begin

Important: Beginning with IBM Spectrum Protect Plus V10.1.6, the hostname for the OVA deployment must not use a name that includes an underscore (_).

Complete the following tasks:

- Review the IBM Spectrum Protect Plus system requirements in "Component requirements " on page 27 and "Hypervisor (Microsoft Hyper-V and VMware) and cloud instance (Amazon EC2) backup and restore requirements " on page 57.
- Download the virtual appliance template installation file *<part_number>*.ova from Passport Advantage® Online. For information about downloading files, see technote 6415111.
- Verify the MD5 checksum of the downloaded template installation file. Ensure that the generated checksum matches the one provided in the MD5 Checksum file, which is part of the software download.
- During deployment, you will be prompted to enter network properties from the VMware user interface. You can enter a static IP address configuration, or leave all fields blank to use a DHCP configuration.
- To reassign a static IP address after deployment, you can use the NetworkManager Text User Interface (nmtui) tool. For more information, see "Assigning a static IP address" on page 121.

Note the following considerations:

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

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

Procedure

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

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

   Important: If you are using the vSphere Web Client (FLEX), verify that `disk.enableUUID = true` presents in **Extra Configuration**. If that is not the case or if you are using the vSphere Client (HTML5), proceed with the installation steps and enable this option from the vSphere Web Client at a later time.
6. Read and accept the End User License Agreement. Check **I accept all license agreements** for vSphere Client or click **Accept** for vSphere Web Client. Click **Next**.
7. Select the storage to which the virtual appliance is to be installed. The datastore of this storage must be configured with the destination host. The virtual appliance configuration file and the virtual disk files will be stored in it. Ensure the storage is large enough to accommodate the virtual appliance including the virtual disk files associated with it. Select a disk format of the virtual disks. Thick provisioning allows for better performance of the virtual appliance. Thin provisioning uses less disk space at the expense of performance. Click **Next**.
8. Select networks for the deployed template to use. Several available networks on the ESXi server might be available by clicking **Destination Network**. Select a destination network that allows you to define the appropriate IP address allocation for the virtual machine deployment. Click **Next**.
9. Enter the property values for the virtual appliance: Hostname, DNS, Default Gateway, Domain, Network IP Address and Network Prefix. A static IP address can be provided. If left blank, a dynamic IP address assigned by a DHCP server will be used. The network prefix must be entered using Classless Inter-Domain Routing (CIDR) notation where valid values are 1 - 24. Click **Next**.

   Note: These properties can be configured using the NetworkManager Text User Interface (`nmtui`) tool. Additionally, information for the Search Domain field can be added using this command. For more information, see Assigning a static IP address.
10. Review your template settings. Click **Finish** to exit the wizard and to start deployment of the OVF template.
11. After the OVF template is deployed, power on your newly created VM. You can power on the VM from the vSphere Client.

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

What to do next

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

Action	How to
Connect to the console of the IBM Spectrum Protect Plus virtual appliance by using VMware Remote Console or SSH. Set up network configurations using the NetworkManager Text User Interface (`nmtui`).	See Assigning a static IP address.
Upload the product key.	See “Uploading the product key” on page 121.
Start IBM Spectrum Protect Plus from a supported web browser.	See “Start IBM Spectrum Protect Plus” on page 257.

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

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

Before you begin

Complete the following tasks:

- Review the IBM Spectrum Protect Plus system requirements in "Component requirements " on page 27 and "Hypervisor (Microsoft Hyper-V and VMware) and cloud instance (Amazon EC2) backup and restore requirements " on page 57.
- Download the installation file *<part_number>*.exe from Passport Advantage Online. For information about downloading files, see technote 6415111.
- Review additional Hyper-V system requirements. See System requirements for Hyper-V on Windows Server.
- Verify the MD5 checksum of the downloaded template installation file. Ensure that the generated checksum matches the one provided in the MD5 Checksum file, which is part of the software download.
- If IBM Spectrum Protect Plus will be used to protect Amazon EC2 workloads, the specified hostname must consist of only lowercase, alphanumeric characters.
- If the hostname of the IBM Spectrum Protect Plus virtual appliance changes after deployment, either through user intervention or if a new IP address is acquired through DNS, the IBM Spectrum Protect Plus virtual appliance must be restarted.
- All Hyper-V servers, including cluster nodes, must have the Microsoft iSCSI Initiator Service running in their Services lists. Set startup type of this service to Automatic so that it starts running when the server starts.
- Administrative privileges may be required to complete certain steps during the installation process.

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

Procedure

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

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

   Important: Do not import multiple IBM Spectrum Protect Plus virtual appliances on a single Hyper-V server.
8. In the **Connect Network** step, set Connection to the virtual switch to use. Click **Next**.
9. In the **Summary** step, review the Description. Click **Finish** to close the Import Virtual Machine wizard.
10. In Hyper-V Manager, locate the new virtual machine named **SPP-{release}**. Right-click this virtual machine and click **Settings**.

11. The Settings dialog for this virtual machine will open. In the navigation panel, click **Hardware > IDE Controller 0 > Hard Drive**.
12. In the Media section, ensure that the correct virtual hard disk is selected. Note the file name of the original virtual disk. Click **Edit**.
13. The Edit Virtual Hard Disk Wizard will open. Go to the **Choose Action** step.
14. In the **Choose Action** step, click **Convert** and then click **Next**.
15. In the **Choose Disk Format** step, ensure that **VHDX** is selected. Click **Next**.
16. For the **Choose Disk Type** step, click **Fixed Size**. Click **Next**.
17. For the **Configure Disk** step, locate the folder to store the virtual disk file of the IBM Spectrum Protect Plus virtual appliance. Reuse the same file name that was noted in Step 12. If the same installation directory from Step 12 is reused, use a different name. Click **Next**.

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

    For more information, see “Assigning a static IP address” on page 121.

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

What to do next

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

Action	How to
Restart the virtual appliance.	Refer to the documentation for the virtual appliance.
Upload the product key.	See “Uploading the product key” on page 121.
Start IBM Spectrum Protect Plus from a supported web browser.	See “Start IBM Spectrum Protect Plus” on page 257.

Chapter 4. Installing IBM Spectrum Protect Plus in a container environment

Before you install IBM Spectrum Protect Plus on Red Hat OpenShift Container Platform, review the components to be deployed, the prerequisites, and the installation procedure.

Related concepts

“Post installation tasks” on page 121

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

Overview of IBM Spectrum Protect Plus container deployment

IBM Spectrum Protect Plus can be installed on a Red Hat OpenShift cluster environment. The installation process uses the IBM Spectrum Protect Plus operator, which deploys and manages all the IBM Spectrum Protect Plus components on Red Hat OpenShift.

The IBM Spectrum Protect Plus operator is a Docker image that uses Ansible Operator technology. The image contains the Kubernetes configuration files that are necessary to deploy and upgrade IBM Spectrum Protect Plus.

If you plan to install IBM Spectrum Protect Plus in the IBM Cloud Pak for Multicloud Management environment, you must use the IBM Spectrum Protect Plus operator for IBM Cloud Pak for Multicloud Management. The operator also works in environments that do not have IBM Cloud Pak for Multicloud Management installed.

The following figure is an example of how IBM Spectrum Protect Plus containers are deployed in OpenShift Container Platform.

© Copyright IBM Corp. 2017, 2021

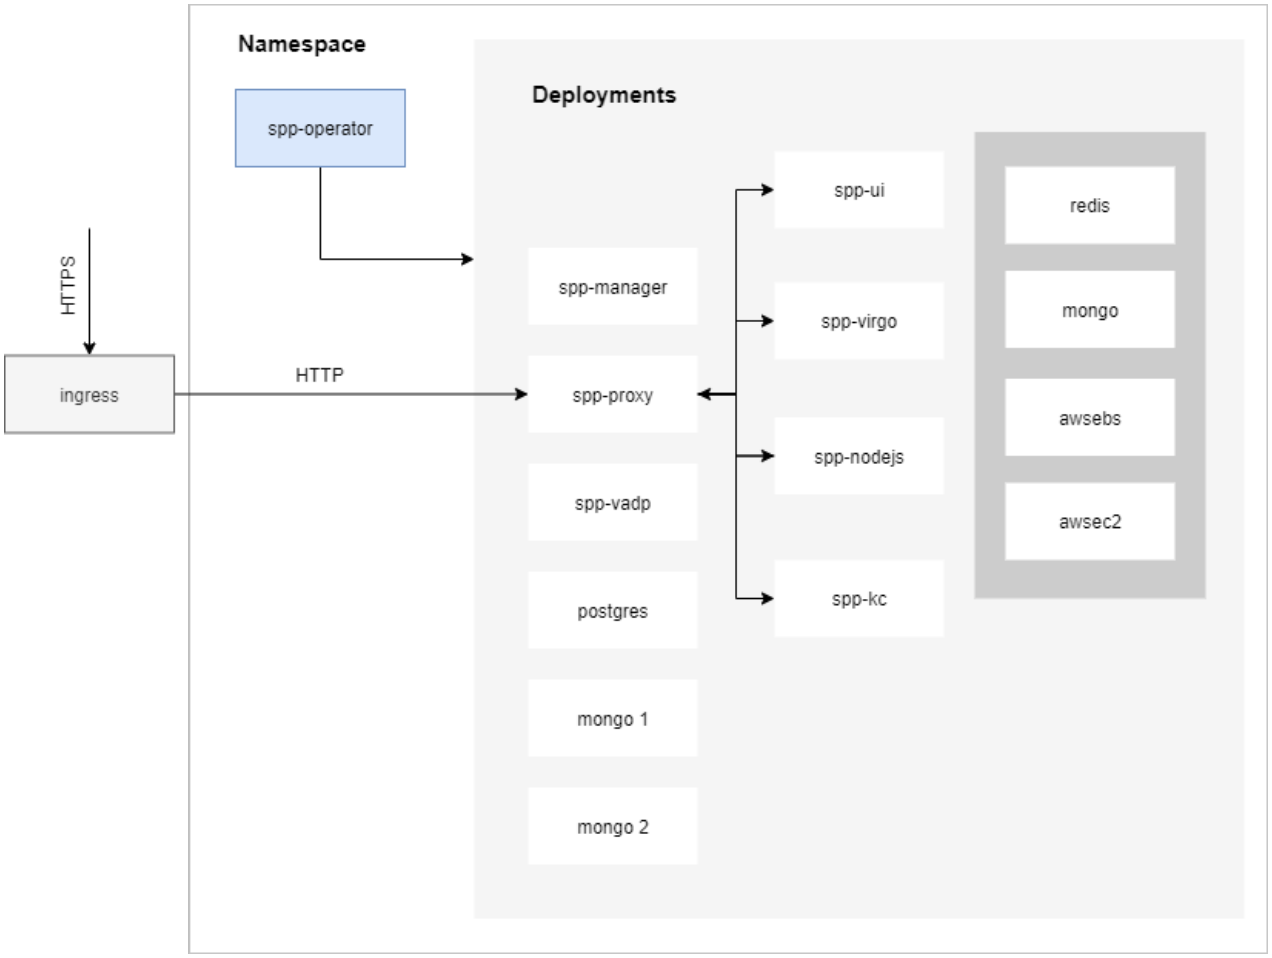


Figure 10. IBM Spectrum Protect Plus container deployment diagram

For an overview of the container components, persistent storage requirements, networking requirements, and CPU and resource requirements, see IBM Spectrum Protect Plus as a set of containers requirements.

Support matrix

For the supported clouds, architectures, and platforms for IBM Spectrum Protect Plus as a set of containers, see IBM Spectrum Protect Plus as a set of containers requirements.

Storage requirements

Storage requirements for IBM Spectrum Protect Plus containers are provided.

Storage provisioning

In order for IBM Spectrum Protect Plus to run on an OpenShift cluster, persistent storage is required.

The IBM Spectrum Protect Plus operator submits requests for storage by using persistent volume claims (PVCs). The OpenShift cluster completes these requests by using an existing storage driver.

For information about the storage capacity that is provisioned for each PVC, see IBM Spectrum Protect Plus as a set of containers requirements.

Storage access modes for persistent volumes

For the storage volume access mode, all PVC access modes are set to RWO (`readwriteonce`).

For more information about the access modes for persistent volumes that are associated with IBM Spectrum Protect Plus containers, see IBM Spectrum Protect Plus as a set of containers requirements.

Roles in IBM Spectrum Protect Plus

IBM Spectrum Protect Plus provides cluster roles (ClusterRole objects) that define the permissions for role-based access control (RBAC) for IBM Spectrum Protect Plus in a Red Hat OpenShift cluster.

A ServiceAccount named `spp-operator` is created to run the pods for IBM Spectrum Protect Plus.

A default ClusterRole named `spp-operator` is defined to the `spp-operator` ServiceAccount. Permissions for this ClusterRole are granted at the cluster-wide level.

A ClusterRoleBinding named `spp-operator` is created to grant cluster-wide access.

The `role_binding.yaml` configuration file is used to apply the permissions.

To apply the `role_binding.yaml` configuration file, issue the following commands in sequence as part of the installation of the IBM Spectrum Protect Plus operator:

```
$ oc apply -f /tmp/spp-operator/operator-group.yaml
$ oc apply -f /tmp/spp-operator/role.yaml
$ oc apply -f /tmp/spp-operator/service_account.yaml
$ oc apply -f /tmp/spp-operator/role_binding.yaml
$ oc apply -f /tmp/spp-operator/spp.ibm.com_ibmspps_crd.yaml
$ oc apply -f /tmp/spp-operator/spp-operator.voperator_version.clusterserviceversion.yaml
```

For more information about installing the IBM Spectrum Protect Plus operator, see “Installing the IBM Spectrum Protect Plus operator in an airgap environment” on page 139.

Installing IBM Spectrum Protect Plus on OpenShift Container Platform

To install IBM Spectrum Protect Plus on an OpenShift cluster, you must prepare your environment, install the IBM Spectrum Protect Plus operator, and create an instance of the IBM Spectrum Protect Plus server.

Before you begin

Ensure that Red Hat OpenShift Container Platform is running on your cloud environment, and you have cluster administrator privileges for your OpenShift cluster.

Ensure that you review the system requirements in IBM Spectrum Protect Plus as a set of containers requirements.

About this task

You can install the IBM Spectrum Protect Plus server as a set of containers by using one of the following methods:

Table 68. Methods for installing IBM Spectrum Protect Plus as a set of containers	
Installation method	**Steps**
By installing the product from an online environment	You can install the product by pulling container images from an online registry, such as the IBM Entitled Registry. Internet access is required to pull containers at deployment time. Complete the following steps: 1. "Preparing to install the IBM Spectrum Protect Plus operator from the IBM Entitled Registry" on page 134 2. "Installing the IBM Spectrum Protect Plus operator in an online environment" on page 137 3. "Creating an IBM Spectrum Protect Plus instance" on page 147 You can also use this method to install IBM Spectrum Protect Plus in the IBM Cloud Pak for Multicloud Management environment.
By installing the product in an airgap environment	You can install the product by using a self-contained product package. Internet access is not required at deployment time. Complete the following steps: 1. "Installing the IBM Spectrum Protect Plus operator in an airgap environment" on page 139 2. "Creating an IBM Spectrum Protect Plus instance" on page 147 You can also use this method to install IBM Spectrum Protect Plus in the IBM Cloud Pak for Multicloud Management environment.
By installing the IBM Spectrum Protect Plus operator at the command line in the IBM Cloud Pak for Multicloud Management environment	Install the IBM Spectrum Protect Plus operator for IBM Cloud Pak for Multicloud Management at the command line in an online environment. Then, create an IBM Spectrum Protect Plus instance in the OpenShift web console. Complete the following steps: 1. "Installing the IBM Spectrum Protect Plus operator for IBM Cloud Pak for Multicloud Management at the command line" on page 145 2. "Creating an IBM Spectrum Protect Plus instance" on page 147

Preparing to install the IBM Spectrum Protect Plus operator from the IBM Entitled Registry

You must complete prerequisite tasks before you can install the IBM Spectrum Protect Plus operator from an online registry such as the IBM Entitled Registry.

For the IBM Cloud Pak for Multicloud Management environment, ensure that you complete the steps on the same hub cluster on which IBM Cloud Pak for Multicloud Management is installed.

1. "Creating an image pull secret for IBM Spectrum Protect Plus" on page 135
2. "Adding the catalog source for IBM Spectrum Protect Plus" on page 135
3. "Creating a project for IBM Spectrum Protect Plus" on page 136

Creating an image pull secret for IBM Spectrum Protect Plus

Create an image pull secret to enable the OpenShift cluster to authenticate with the IBM Entitled Registry. The image pull secret provides the credentials for pulling Docker images from the IBM Entitled Registry.

Before you begin

Before you can pull Docker images from the IBM Entitled Registry, you must obtain an entitlement key for accessing your container software. To obtain an entitlement key:

1. Log in to the IBM Container software library with the IBMid and password that is associated with the entitled software.
2. Click **Get entitlement key**.
3. In the **Access your container software** page, click **Copy key** to copy the generated entitlement key.
4. Save the key to a secure location for later use.

Procedure

To create an image pull secret, complete the following steps:

1. Log on to the OpenShift web console as the cluster administrator.
2. In the navigation panel, click **Workloads** > **Secrets**.
3. On the **Secrets** page, click **Create** > **Image Pull Secret**.
4. Ensure that the project is created in the **openshift-marketplace** project by clicking **Project** > **openshift-marketplace**.
5. On the **Create Image Pull Secret** page, enter **ibmspp-image-secret** for the name of the secret.
6. Create the image pull secret by using one of the following methods:
   - Upload a configuration file that contains the credentials for the IBM Entitled Registry:
     a. In the **Authentication Type** list, click **Upload Configuration File**.
     b. Browse and select an existing configuration file, or copy and paste the credentials from an existing configuration file into the text box.
     c. Click **Create**.
   - Manually enter the credentials for the IBM Entitled Registry:
     a. In the **Authentication Type** list, click **Image Registry Credentials**.
     b. In the **Registry Server Address** field, enter the address for the IBM Entitled Registry: `cp.icr.io/cp/sppserver`
     c. In the **Username** field, enter `cp`.
     d. In the **Password** field, enter the entitlement key that you obtained.
     e. Click **Create**.

Adding the catalog source for IBM Spectrum Protect Plus

Before you can install the IBM Spectrum Protect Plus operator, you must add the catalog source for the operator to the OpenShift web console.

Procedure

Apply the following configuration by using the command line or the OpenShift web console:

```
apiVersion: operators.coreos.com/v1alpha1
kind: CatalogSource
metadata:
  name: ibm-spp-operator
  namespace: openshift-marketplace
spec:
  displayName: IBM SPP Operator
```

```
  image: 'image_path'
  publisher: IBM
  secrets:
    - ibmspp-image-secret
  sourceType: grpc
  updateStrategy:
    registryPoll:
      interval: 45m
```

where:

image_path
Specifies the location for the operator images. Ensure that the image path is enclosed in single quotation marks.

- If you are installing the operator in a standard OpenShift cluster environment, specify: `ibmcom/spp-operator-catalog:latest`
- If you are installing the operator in the IBM Cloud Pak for Multicloud Management environment, specify: `ibmcom/ibm-mcm-spp-operator-catalog:latest`

Tip: To add the catalog source at the command line, create a YAML file that contains the configuration. Then, issue the following command:

```
oc apply -f filename.yaml
```

where *filename* specifies the name of the YAML file that you created.

Creating a project for IBM Spectrum Protect Plus

You must create a project (namespace) for IBM Spectrum Protect Plus in the OpenShift web console. Then, create the image pull secret for IBM Spectrum Protect Plus (**ibmspp-image-secret**) in this project.

About this task

The project is where the operator and instance of the IBM Spectrum Protect Plus server will be installed.

Procedure

1. Log on to the OpenShift web console as the cluster administrator.
2. In the navigation panel, click **Home** > **Projects** > **Create project**.
3. In the **Create Project** window, enter a name for the project.
4. Optional: Enter a display name and description for the project.
5. Click **Create** to create the project.
6. Create the image pull secret for IBM Spectrum Protect Plus (**ibmspp-image-secret**) in this project. This secret is the same pull secret that you created in “Creating an image pull secret for IBM Spectrum Protect Plus” on page 135.
   a) In the navigation panel, click **Workloads** > **Secrets**.
   b) On the **Secrets** page, click **Create** > **Image Pull Secret**.
   c) Ensure that the project you created for IBM Spectrum Protect Plus is selected in the **Project** menu.
   d) On the **Create Image Pull Secret** page, enter **ibmspp-image-secret** for the name of the secret.
   e) Create the image pull secret by using the credentials for the IBM Entitled Registry.
      For more information, see Step “6” on page 135 of “Creating an image pull secret for IBM Spectrum Protect Plus” on page 135.

Results

The IBM Spectrum Protect Plus operator is available in the OpenShift web console in the project that you created for IBM Spectrum Protect Plus.

What to do next

Install the IBM Spectrum Protect Plus operator. For instructions, see “Installing the IBM Spectrum Protect Plus operator in an online environment” on page 137.

Installing the IBM Spectrum Protect Plus operator in an online environment

Install the IBM Spectrum Protect Plus operator so that all IBM Spectrum Protect Plus components can be deployed and managed by the operator on Red Hat OpenShift Container Platform.

Before you begin

The information that is provided applies to installing the operator in an online environment. If the OpenShift cluster operates is an airgap environment, follow the instructions in “Installing the IBM Spectrum Protect Plus operator in an airgap environment” on page 139.

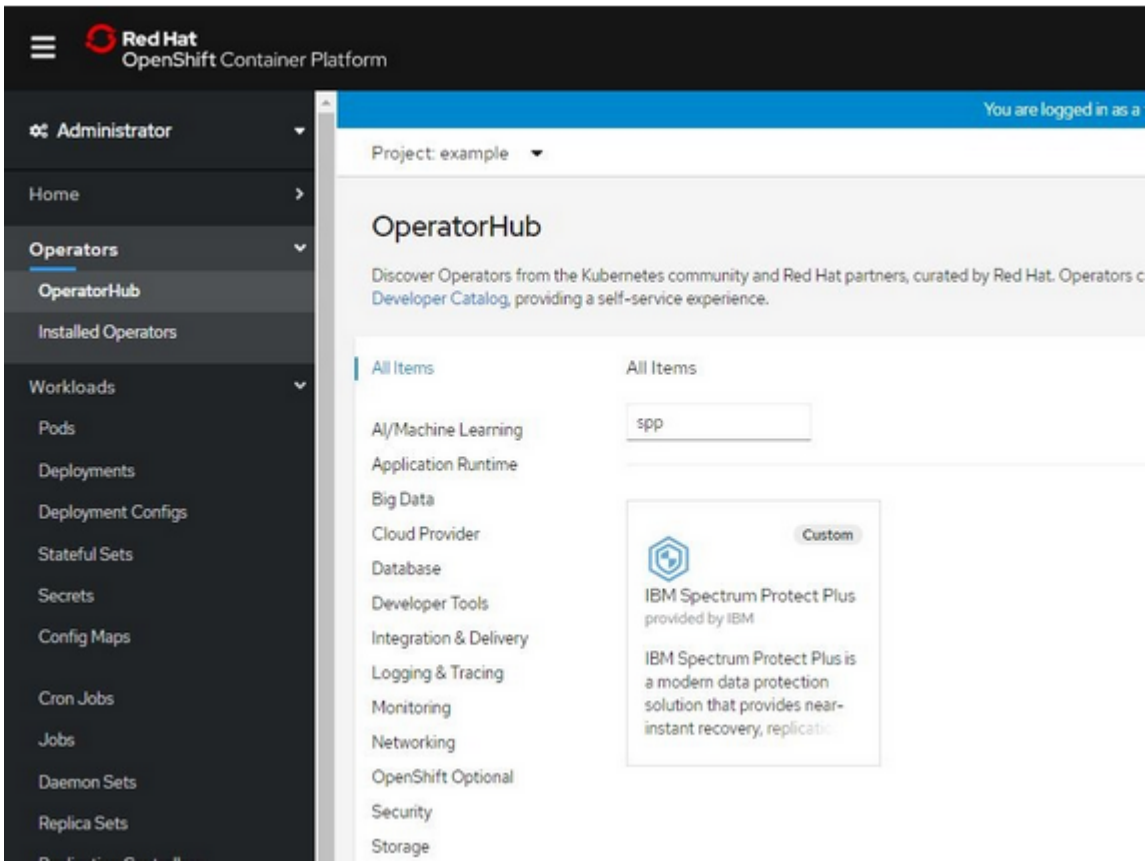
Ensure that you complete the prerequisite steps in “Preparing to install the IBM Spectrum Protect Plus operator from the IBM Entitled Registry” on page 134.

The IBM Spectrum Protect Plus operator is namespace-scoped so that it can be installed in any namespace in an OpenShift cluster. However, the IBM Spectrum Protect Plus instance must be deployed in the namespace in which the IBM Spectrum Protect Plus operator is installed.

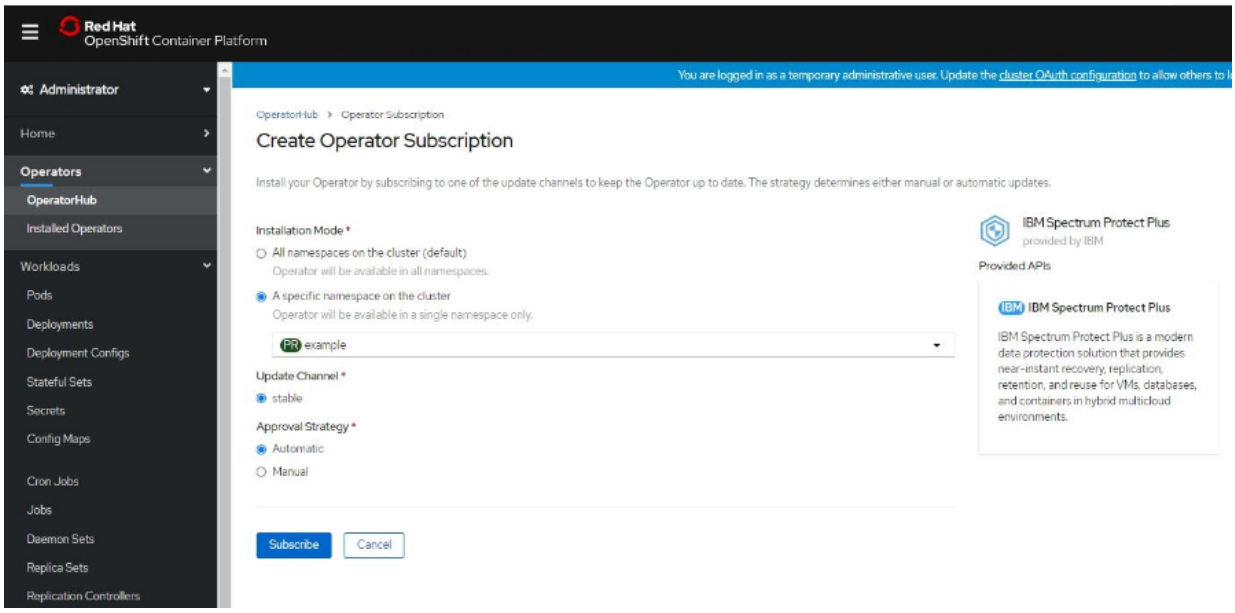
For the IBM Cloud Pak for Multicloud Management environment, ensure that you complete the steps on the same hub cluster on which IBM Cloud Pak for Multicloud Management is installed.

Procedure

1. Log on to the OpenShift web console as the cluster administrator.
2. In the navigation panel, click **Operators** > **OperatorHub**.
3. On the **OperatorHub** page, ensure that **all namespaces** is selected in the **Project** list.
4. In the search field under **All items**, enter **spp**.



5. Click the **IBM Spectrum Protect Plus** card and click **Install**.
6. On the **Create Operator Subscription** page, click **A specific namespace on the cluster** and specify the project that you created for IBM Spectrum Protect Plus.



7. In the **Approval Strategy** section, select whether to update the operator automatically or manually.
8. Click **Subscribe**.

   The **Installed Operators** page is displayed.
9. Wait a few moments for the operator to install. When the operator is successfully installed, the `InstallSucceeded` status is displayed.

Results

The IBM Spectrum Protect Plus operator is installed as a container. You can verify that the operator is installed by completing the following steps:

1. In the navigation panel, click **Workloads** > **Pods**.
2. On the **Pods** page, verify that the IBM Spectrum Protect Plus operator container is running. The container name begins with the prefix **spp-operator**.
3. In the navigation panel, click **Operators** > **Installed Operators**.
4. From the list of installed operators, click IBM Spectrum Protect Plus and review information about the installed operator.

What to do next

Create an instance of the IBM Spectrum Protect Plus server by using the IBM Spectrum Protect Plus operator that you installed. For instructions, see "Creating an IBM Spectrum Protect Plus instance" on page 147.

Installing the IBM Spectrum Protect Plus operator in an airgap environment

If you plan to deploy an instance of the IBM Spectrum Protect Plus server as a set of containers in an airgap environment, you must prepare your environment and install the IBM Spectrum Protect Plus operator. You can install the operator by using IBM Passport Advantage Online or by using a Container Application Software for Enterprises (CASE) bundle.

Installing the IBM Spectrum Protect Plus operator in an airgap environment by using IBM Passport Advantage Online

You can install the IBM Spectrum Protect Plus in an air gap environment by using a package that you downloaded from the IBM Passport Advantage Online website. You can use this procedure in either a standard OpenShift cluster environment or the IBM Cloud Pak for Multicloud Management environment.

Before you begin

Ensure that the following prerequisites are met:

- All command and scripts must be run on a Linux operating system.
- The OpenShift command-line interface (**oc**) must be installed in your environment.
- The OpenShift cluster must be able to pull images from the private registry.
- For the IBM Cloud Pak for Multicloud Management environment, ensure that you complete the steps on the same hub cluster on which IBM Cloud Pak for Multicloud Management is installed.
- Ensure that the **skopeo** v1.2.2 utility is installed in your environment. The **skopeo** utility is used by the `load_files.sh` script to copy the installation images from the IBM Entitled Registry to your private registry. For information about downloading the **skopeo** v1.2.2 utility, see https://github.com/containers/skopeo.
- Ensure that you have access with the proper permissions to the IBM Entitled Registry (`cp.icr.io`). To pull Docker images from the IBM Entitled Registry, you must obtain an entitlement key for accessing your container software. To obtain an entitlement key:
  1. Log in to the IBM Container software library with the IBMid and password that is associated with the entitled software.
  2. Click **Get entitlement key**.
  3. In the **Access your container software** page, click **Copy key** to copy the generated entitlement key.
  4. Save the key to a secure location for later use.

Procedure

1. Log in to the `cp.icr.io` registry by providing the entitlement key that you obtained. Issue the following **`docker login`** command:

   ```
   docker login cp.icr.io -u cp
   ```

   When prompted for the password, enter the entitlement key.

2. On the server where Docker images are loaded, create a directory for downloading the package for installing IBM Spectrum Protect Plus as a set of containers. Issue the following command:

   ```
   mkdir -p sppimages
   ```

   where *sppimages* is the name of the directory where the installation package is downloaded.

3. Download the installation package from IBM Passport Advantage Online to the *sppimages* directory that you created.

   - For a standard OpenShift cluster environment, download the following package: SPP_V*sppversion*_FOR_OpenShift.tar.gz
   - For the IBM Cloud Pak for Multicloud Management environment, download the following package: SPP_V*sppversion*_FOR_OpenShift_MCM.tar.gz

   where *sppversion* specifies the version of IBM Spectrum Protect Plus that you are installing, such as 10.1.8.

   For the list of installation packages by component, and the links to the download site for the files, see technote 6415111.

4. Extract the installation package by issuing the following commands:

   ```
   $ cd sppimages
   $ tar -xvf spp_installation_package
   ```

   where *sppimages* specifies the directory where the installation package is downloaded, and *spp_installation_package* specifies the name of the installation package that you downloaded.

   The *sppimages* directory is populated with the contents in the archive package. The contents in the package are similar to the following example:

   ```
   .
   ├── operator_version
   │   ├── catalog-source.yaml
   │   ├── image-content-source-policy.yaml
   │   ├── operator-group.yaml
   │   ├── role_binding.yaml
   │   ├── role.yaml
   │   ├── service_account.yaml
   │   ├── spp.ibm.com_ibmspps_crd.yaml
   │   └── spp-operator.voperator_version.clusterserviceversion.yaml
   ├── images
   │   ├── mongodb-36-rhel7.tar
   │   ├── nginx-116-rhel7.tar
   │   ├── postgresql-96.tar
   │   ├── Readme
   │   ├── redis-5-rhel7.tar
   │   ├── spp-awsebs.tar
   │   ├── spp-awsec2.tar
   │   ├── spp-kc.tar
   │   ├── spp-manager.tar
   │   ├── spp-node.tar
   │   ├── spp-operator.tar
   │   ├── spp.tar
   │   ├── spp-uinginx.tar
   │   ├── spp-vadp.tar
   │   └── version.yaml
   └── load_images.sh
   ```

   where *operator_version* is the IBM Spectrum Protect Plus operator version, for example, "1.0.0-570".

5. Export the Docker registry to your private registry by issuing the following command:

```
$ export DOCKER_REGISTRY=registry.example.com:5000
```

where *registry.example.com* is the Docker registry address.

6. If you have not already done so, log in to the private Docker registry by issuing the following command:

```
$ docker login registry.example.com
```

where *registry.example.com* is the Docker registry address.

7. Run the `load_images.sh` script to tag and push the IBM Spectrum Protect Plus images to the private Docker registry:

```
$ cd sppimages
$ ./load_images.sh
```

where *sppimages* specifies the directory where the installation package was extracted.

8. From the *sppimages* directory, create a copy of the configuration directory to a temporary location by issuing the following command:

```
$ cp -R ./operator_version /tmp/spp-operator
```

where the folder name, *operator_version*, specifies the IBM Spectrum Protect Plus operator version. For example, issue the following command:

```
$ cp -R ./1.0.0-570 /tmp/spp-operator
```

9. Log in to the OpenShift cluster by issuing the following command:

```
$ oc login
```

10. Create a namespace (project) where the operator and IBM Spectrum Protect Plus instance will be installed:

```
$ oc create namespace your_namespace
```

11. Update the namespace values in the configuration files with the newly created namespace by issuing the following command:

```
$ sed -i 's#{{ NAMESPACE }}#your_namespace#g' /tmp/spp-operator/*
```

where *your_namespace* is the namespace that you created for the operator and IBM Spectrum Protect Plus instance.

12. Update the image registry values in the configuration files with your image registry by issuing the following command:

```
$ sed -i 's#{{ IMAGE_REGISTRY }}#your_image_registry#g' /tmp/spp-operator/*
```

where *your_image_registry* is your private Docker registry address.

13. Optional: Create an image pull secret to authenticate with your private Docker registry by taking the following actions:

a) Create an image pull secret at the command line. For example, you can issue the following command:

```
$ oc create -n your_namespace secret docker-registry your_image_pull_secret \
    --docker-server=your_image_registry \
    --docker-username=your_docker_username \
    --docker-password=your_docker_password \
    --docker-email=your_docker_email
```

where *your_image_pull_secret* is the name of the image pull secret and *your_image_registry* is your private Docker registry address.

b) In the following files:

```
service_account.yaml
spp-operator.voperator_version.clusterserviceversion.yaml
```

where *operator_version* specifies the IBM Spectrum Protect Plus operator version, modify the **imagePullSecrets** field by replacing the default name with the name of the image pull secret that you created.

```
...
spec:
  imagePullSecrets:
  - name: your_image_pull_secret
...
```

14. Apply the configuration files to install the operator by issuing the following commands in the order that is listed:

```
$ oc apply -f /tmp/spp-operator/operator-group.yaml
$ oc apply -f /tmp/spp-operator/role.yaml
$ oc apply -f /tmp/spp-operator/service_account.yaml
$ oc apply -f /tmp/spp-operator/role_binding.yaml
$ oc apply -f /tmp/spp-operator/spp.ibm.com_ibmspps_crd.yaml
$ oc apply -f /tmp/spp-operator/spp-operator.voperator_version.clusterserviceversion.yaml
```

where *operator_version* specifies the IBM Spectrum Protect Plus operator version, for example, "1.0.0-570".

Results

The IBM Spectrum Protect Plus operator is installed in the namespace that you created for IBM Spectrum Protect Plus.

What to do next

Create an instance of the IBM Spectrum Protect Plus server. For instructions, see “Creating an IBM Spectrum Protect Plus instance” on page 147.

Installing the IBM Spectrum Protect Plus operator in an airgap environment by using a CASE bundle

You can install the IBM Spectrum Protect Plus operator in an airgap environment by using a Container Application Software for Enterprises (CASE) bundle. You can use this procedure in either a standard OpenShift cluster environment or the IBM Cloud Pak for Multicloud Management environment.

Before you begin

Ensure that prerequisites are met and preliminary tasks are completed:

- Ensure that the following tools are installed and updated to the required version:
  - IBM Cloud Pak command-line interface (**cloudctl**) v3.5.0 or later
  - Kubernetes command-line tool (**kubectl**) v1.16.0 or later
  - OpenShift command-line tool (**oc**) v4.5.0 or later
- All commands must be run in the Linux operating system.
- For the IBM Cloud Pak for Multicloud Management environment, ensure that you complete the steps on the same hub cluster on which IBM Cloud Pak for Multicloud Management is installed.
- Download the IBM Spectrum Protect Plus operator CASE bundle.
  - For the IBM Cloud Pak for Multicloud Management environment, obtain the SPP_V*sppversion*_FOR_OpenShift_MCM.tar.gz package from IBM Passport Advantage Online, where *sppversion* specifies the version of IBM Spectrum Protect Plus that you are installing, such as 10.1.8.

– For the standard OpenShift environment, obtain the SPP_V*sppversion*_FOR_OpenShift.tar.gz package from IBM Passport Advantage Online, where *sppversion* specifies the version of IBM Spectrum Protect Plus that you are installing, such as 10.1.8.

For the list of installation packages by component, and the links to the download site for the files, see technote 6415111.

About this task

You must prepare your cluster and then install the operator.

The following table shows the descriptions of the variables that you must set before you begin the procedure:

Table 69. Variables for installing the IBM Spectrum Protect Plus operator by using a CASE bundle

Variable	Description
$CASE_BUNDLE_DIRECTORY	The name of the CASE bundle directory. For IBM Cloud Pak for Multicloud Management, specify: `ibm-spp-server-bundle` For standard OpenShift clusters, specify: `ibm-spp-server-std-bundle`
$CASE_BUNDLE_NAME	The name of the CASE bundle. For IBM Cloud Pak for Multicloud Management, specify: `ibm-spp-server-case-bundle` For standard OpenShift clusters, specify: `ibm-spp-server-std-case-bundle`
$CASE_NAME	The CASE name. For IBM Cloud Pak for Multicloud Management, specify: `ibm-spp-server` For standard OpenShift clusters, specify: `ibm-spp-server-std`
$SOURCE_REGISTRY	The environment variable for the IBM Entitled Registry address. The value must be set to `cp.icr.io/cp/sppserver`.
$SOURCE_REGISTRY_USER	The environment variable for the username for the IBM Entitled Registry. The value must be set to `cp`.
$SOURCE_REGISTRY_PASS	The environment variable for the entitlement key that you obtained from the IBM Container software library.
$TARGET_REGISTRY	The environment variable for the private Docker registry where images are loaded.
$TARGET_REGISTRY_USER	The environment variable for the username for the private Docker registry.
$TARGET_REGISTRY_PASS	The environment variable for the private Docker registry password.
project_name	The project that you created for IBM Spectrum Protect Plus.

Procedure

To install the IBM Spectrum Protect Plus operator by using a CASE bundle in an airgap environment, complete the following steps:

1. Log in to the IBM Cloud Pak environment. Issue the following **`cloudctl`** command to log in to your cluster:

```
cloudctl login -a cluster_url -u username -p password -n kube-system --skip-ssl-validation
```

2. Create a project for IBM Spectrum Protect Plus by issuing the following command:

```
oc new-project project_name
```

   where *project_name* is the name of the project where the operator and instance of the IBM Spectrum Protect Plus server will be installed.

3. Save the CASE bundle. Prepare the CASE bundle for installation by issuing the following commands:

```
cd $CASE_BUNDLE_DIRECTORY/stable/$CASE_BUNDLE_NAME
cloudctl case save --case case/$CASE_NAME/ --outputdir /tmp/ibm-spp-server-case
```

4. Configure the credentials for the source registry from which images are pulled. Add the registry credentials for the source registry by issuing the following commands:

```
cd $CASE_BUNDLE_DIRECTORY/stable/$CASE_BUNDLE_NAME
cloudctl case launch --case case/$CASE_NAME/ --namespace project_name --inventory
sppserverSetup --action configure-creds-airgap --args "--registry $SOURCE_REGISTRY --user
$SOURCE_REGISTRY_USER --pass $SOURCE_REGISTRY_PASS"
```

5. Configure the credentials for the target registry where images are loaded. Add the registry credentials for the target registry by issuing the following commands:

```
cd $CASE_BUNDLE_DIRECTORY/stable/$CASE_BUNDLE_NAME
cloudctl case launch --case case/$CASE_NAME/ --namespace project_name --inventory
sppserverSetup --action configure-creds-airgap --args "--registry $TARGET_REGISTRY --user
$TARGET_REGISTRY_USER --pass $TARGET_REGISTRY_PASS"
```

6. Mirror the images from the source registry to the target registry by issuing the following commands:

```
cd $CASE_BUNDLE_DIRECTORY/stable/$CASE_BUNDLE_NAME
cloudctl case launch --case case/$CASE_NAME/ --namespace project_name --inventory
sppserverSetup --action mirror-images --args "--registry $TARGET_REGISTRY --inputDir /tmp/
ibm-spp-server-case"
```

7. Configure your cluster for an airgap installation by adding a global pull secret and setting an image content source policy. Issue the following commands:

```
cd $CASE_BUNDLE_DIRECTORY/stable/$CASE_BUNDLE_NAME
cloudctl case launch --case case/$CASE_NAME/ --namespace project_name --inventory
sppserverSetup --action configure-cluster-airgap --args "--registry $TARGET_REGISTRY --
inputDir /tmp/ibm-spp-server-case --dryRun"
```

8. Install the IBM Spectrum Protect Plus operator by completing the following steps:

   a) Make the IBM Spectrum Protect Plus operator available for installation by adding the catalog source. Issue the following commands:

```
> cd $CASE_BUNDLE_DIRECTORY/stable/$CASE_BUNDLE_NAME
> cloudctl case launch --case case/$CASE_NAME/ --namespace openshift-marketplace --
inventory sppserverSetup --action installCatalog \
    --args "--registry registry_server --secret pull_secret_name"
```

   where *registry_server* is the registry server address for the IBM Entitled Registry (`cp.icr.io/cp/sppserver`) and *pull_secret_name* is the name of the image pull secret that you created.

   b) Install the IBM Spectrum Protect Plus operator by issuing the following commands:

```
> cd $CASE_BUNDLE_DIRECTORY/stable/$CASE_BUNDLE_NAME
> cloudctl case launch --case case/$CASE_NAME/ --namespace project_name --inventory
sppserverSetup --action installOperator
```

   where *project_name* is the project that you created for IBM Spectrum Protect Plus.

What to do next

Create an instance of the IBM Spectrum Protect Plus server by using the OpenShift web console. For instructions, see “Creating an IBM Spectrum Protect Plus instance” on page 147.

Installing the IBM Spectrum Protect Plus operator for IBM Cloud Pak for Multicloud Management at the command line

If you plan to run IBM Spectrum Protect Plus in the IBM Cloud Pak for Multicloud Management environment, you can install the IBM Spectrum Protect Plus operator at the command line instead of using the OpenShift web console.

Installing the IBM Spectrum Protect Plus operator for IBM Cloud Pak for Multicloud Management at the command line in an online environment

Before you can create an instance of the IBM Spectrum Protect Plus server, you must install the IBM Spectrum Protect Plus operator for IBM Cloud Pak for Multicloud Management.

Before you begin

Ensure that prerequisites are met and preliminary tasks are completed:

- Ensure that the following tools are installed and updated to the required version:
  - IBM Cloud Pak command-line interface (**cloudctl**) v3.5.0 or later
  - Kubernetes command-line tool (**kubectl**) v1.16.0 or later
  - OpenShift command-line tool (**oc**) v4.5.0 or later
- All commands must be run in the Linux operating system.
- For the IBM Cloud Pak for Multicloud Management environment, ensure that you complete the steps on the same hub cluster on which IBM Cloud Pak for Multicloud Management is installed.
- Download the IBM Spectrum Protect Plus operator Container Application Software for Enterprises (CASE) bundle.

  Obtain the SPP_V*version*_FOR_OpenShift_MCM.tar.gz package from IBM Passport Advantage Online, where *sppversion* specifies the version of IBM Spectrum Protect Plus that you are installing, such as 10.1.8. For the list of installation packages by component, and the links to the download site for the files, see technote 6415111.

About this task

Use the following procedure to install the IBM Spectrum Protect Plus operator in an online environment by using a CASE bundle.

Procedure

1. Because you will be pulling images from the IBM Entitled Registry, you must obtain an entitlement key for accessing your container software. To obtain an entitlement key:
   a) Log in to the IBM Container software library with the IBMid and password that is associated with the entitled software.
   b) Click **Get entitlement key**.
   c) In the **Access your container software** page, click **Copy key** to copy the generated entitlement key.
   d) Save the key to a secure location for later use.
2. Log in to the IBM Cloud Pak environment. Issue the following **cloudctl** command to log in to your cluster:

   ```
   cloudctl login -a cluster_url -u username -p password -n kube-system --skip-ssl-validation
   ```

3. Create a project for IBM Spectrum Protect Plus by issuing the following command:

   ```
   oc new-project project_name
   ```

where *project_name* is the name of the project where the operator and instance of the IBM Spectrum Protect Plus server will be installed.

4. Add an image pull secret to pull Docker images from the IBM Entitled Registry. Issue the following commands to create an image pull secret and add the secret to the default service account in your cluster:

```
> kubectl create secret docker-registry pull_secret_name \
    --docker-server=registry_server \
    --docker-username=user_name \
    --docker-password=password \
    --docker-email=email \
    --namespace openshift-marketplace
> kubectl patch serviceaccount default -n openshift-marketplace -p '{"imagePullSecrets":
[{"name": "pull_secret_name"}]}'
```

where:

pull_secret_name
Specifies a name for the image pull secret that you create.

registry_server
Specifies the registry server address for the IBM Entitled Registry: `cp.icr.io/cp/sppserver`

user_name
Specifies the username for the IBM Entitled Registry: `cp`

password
Specifies the entitlement key that you obtained from the IBM Container software library.

5. Install the IBM Spectrum Protect Plus operator by completing the following steps:

a) Make the IBM Spectrum Protect Plus operator available for installation by adding the catalog source. Issue the following commands:

```
> cd ibm-spp-server-bundle/stable/ibm-spp-server-case-bundle
> cloudctl case launch --case case/ibm-spp-server/ --namespace openshift-marketplace --
inventory sppserverSetup --action installCatalog \
    --args "--registry registry_server --secret pull_secret_name"
```

where *registry_server* is the registry server address for the IBM Entitled Registry (`cp.icr.io/cp/sppserver`) and *pull_secret_name* is the name of the image pull secret that you created.

b) Install the IBM Spectrum Protect Plus operator by issuing the following commands:

```
> cd ibm-spp-server-bundle/stable/ibm-spp-server-case-bundle
> cloudctl case launch --case case/ibm-spp-server/ --namespace project_name --inventory
sppserverSetup --action installOperator
```

where *project_name* is the project that you created for IBM Spectrum Protect Plus.

What to do next

Create an instance of the IBM Spectrum Protect Plus server by using the OpenShift web console. For instructions, see “Creating an IBM Spectrum Protect Plus instance” on page 147.

Uninstalling the IBM Spectrum Protect Plus operator for IBM Cloud Pak for Multicloud Management at the command line

You must uninstall the IBM Spectrum Protect Plus operator and then remove the catalog source.

Procedure

1. Log in to the IBM Cloud Pak environment. Issue the following **`cloudctl`** command to log in to your cluster:

```
cloudctl login -a cluster_url -u username -p password -n kube-system --skip-ssl-validation
```

2. Uninstall the IBM Spectrum Protect Plus operator by issuing the following commands:

```
> cd ibm-spp-server-bundle/stable/ibm-spp-server-case-bundle
> cloudctl case launch --case case/ibm-spp-server-case/ --namespace project_name --inventory
sppserverSetup --action uninstallOperator
```

where *project_name* is the project that you created for IBM Spectrum Protect Plus.

3. Remove the catalog source by issuing the following commands:

```
> cd ibm-spp-server-bundle/stable/ibm-spp-server-case-bundle
> cloudctl case launch --case case/ibm-spp-server-case/ --namespace openshift-marketplace
--inventory sppserverSetup --action uninstallCatalog
```

Creating an IBM Spectrum Protect Plus instance

After the IBM Spectrum Protect Plus operator is installed, create an instance of the IBM Spectrum Protect Plus server from the installed operator page. The instance is required for running IBM Spectrum Protect Plus in a container environment.

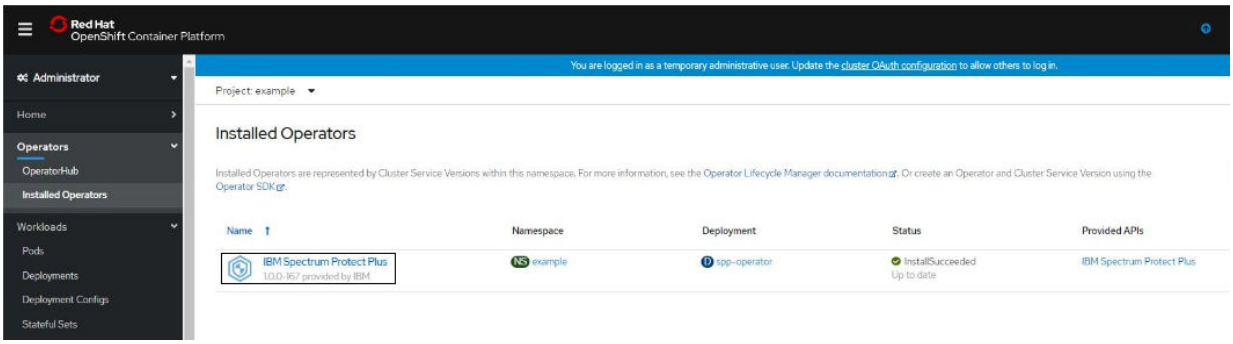
Before you begin

Ensure that you create the IBM Spectrum Protect Plus instance in the same namespace in which the IBM Spectrum Protect Plus operator is installed.

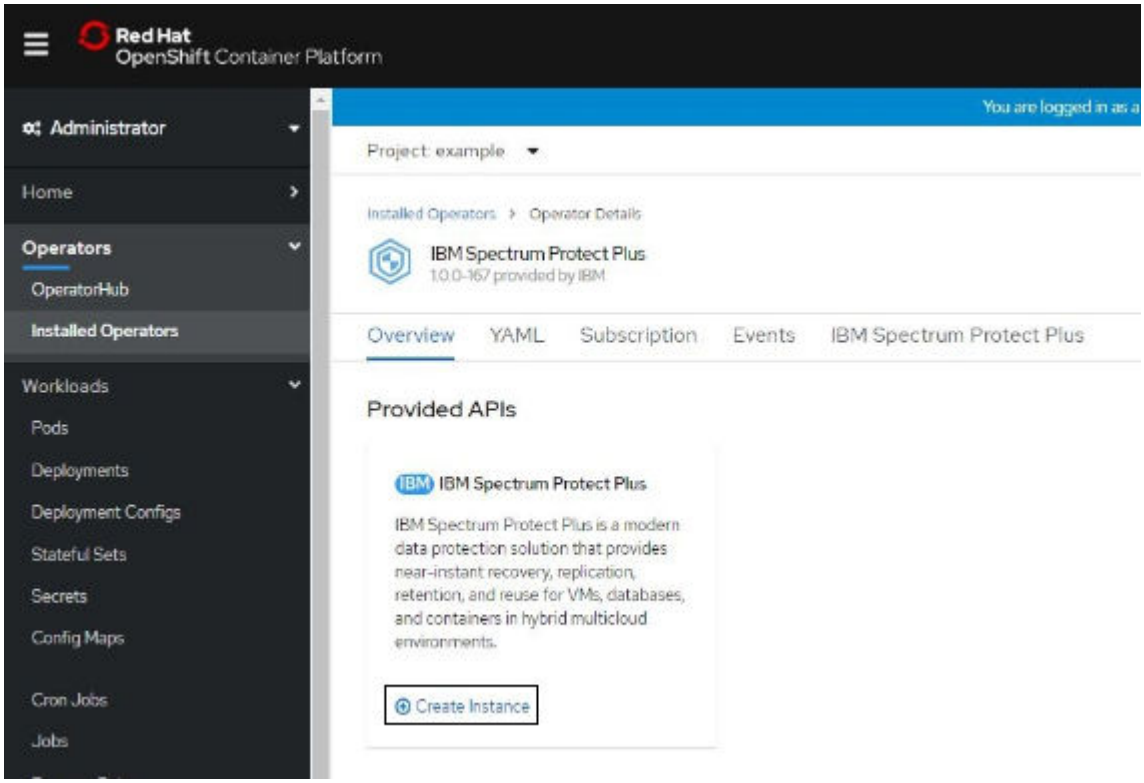
For the IBM Cloud Pak for Multicloud Management environment, ensure that you complete the steps on the same hub cluster on which IBM Cloud Pak for Multicloud Management is installed.

Procedure

1. Log on to the OpenShift web console as the cluster administrator.
2. In the navigation panel, click **Operators** > **Installed Operators**.
3. On the **Installed Operators** page, ensure that the project that you created for IBM Spectrum Protect Plus is selected in the **Project** list.
4. Select the newly created IBM Spectrum Protect Plus operator.



5. On the IBM Spectrum Protect Plus operator details page, click **Create Instance**.



6. On the **Create IBMSPP** page, click **Edit form**.
7. Update the configurations by completing the following fields:

Field name	Description
Name	Specify a name for the IBM Spectrum Protect Plus instance.
Image Registry	Specify the path for pulling images. For an online installation from the IBM Entitled Registry, specify **cp.icr.io/cp/sppserver**. For an airgap installation, specify the name of the Docker registry in your environment where container images are pulled from.
Image Pull Secret	Select the image pull secret that you created for IBM Spectrum Protect Plus.
Hostname	Specify a public hostname for the instance. The hostname must be in the following format: `instancename-spp.routerCanonicalHostname` where *instancename* is the name of the IBM Spectrum Protect Plus instance (as specified in the **Name** field), and *routerCanonicalHostname* is the external host name for the OpenShift router. This hostname is also used to access the IBM Spectrum Protect Plus user interface.
Storage Class	Select the storage class to that is used for volumes. The storage class is already configured with the OpenShift environment.
Version	Select a version of IBM Spectrum Protect Plus to install.
License	Review and accept the license agreement.
Use same registry for RedHat images	For airgap installations only: Set this option to **true** to pull Red Hat images from the same registry that you specified for IBM Spectrum Protect Plus images (in the **Image Registry** field). For online installations, leave the setting as **false**.

8. Modify default configuration values as needed by switching from **Form view** to **YAML view** in the **Create IBMSPP** page.

   Important: For the IBM Spectrum Scale storage class, you must set the default value for the `postgres` persistent volume claim (PVC) to at least 10Gi. To modify the default value, click **YAML view**. In the YAML editor, locate the **`postgres`** field and change the value in the **`size`** field to `10Gi` as shown in the following example:

   ```
   postgres:
     name: sppdbpostgres
     image: postgresql-96
     service: pg
     version: latest
     size: 10Gi
   ```

9. Install the instance of IBM Spectrum Protect Plus by clicking **Create**.

   The operator details page for the IBM Spectrum Protect Plus operator shows the progress of the installation.

What to do next

Verify that the IBM Spectrum Protect Plus instance was created successfully and then start the IBM Spectrum Protect Plus user interface. For instructions, see "Verifying the installation of the IBM Spectrum Protect Plus instance" on page 149.

Related reference

"Troubleshooting installation issues for IBM Spectrum Protect Plus as a set of containers" on page 659
Review the available information to resolve installation issues with the IBM Spectrum Protect Plus server in a container environment.

Verifying the installation of the IBM Spectrum Protect Plus instance

You can verify that the IBM Spectrum Protect Plus instance was created successfully and start the IBM Spectrum Protect Plus user interface.

About this task

To verify the installation, you must complete the following tasks:

- Ensure that all pods are running in the project that you created for IBM Spectrum Protect Plus
- Check the status of the IBM Spectrum Protect Plus instance
- Ensure that the spp route is created and the IBM Spectrum Protect Plus user interface can be started

Procedure

1. Ensure that all pods are in the running state in the project that you created for IBM Spectrum Protect Plus:
   a) In the navigation panel of the OpenShift web console, click **Workloads** > **Pods**.
   b) In the **Project** list on the **Pods** page, select the project that you created for IBM Spectrum Protect Plus.
   c) Select the container for the IBM Spectrum Protect Plus operator. The container name begins with the prefix `spp-operator`.
   d) On the pod details page for the IBM Spectrum Protect Plus operator, view the operator logs by clicking **Logs** > **operator**. You can view the IBM Spectrum Protect Plus components as they are being installed.
   e) After the containers are installed, wait for the containers to reach the `Running` state. For each container that is shown on the **Pods** page, you can click the name of the container to show the detailed logs. Ensure that the SPP`-proxy` and `spp-virgo` containers are running. The `spp-virgo` container can take approximately 5 - 10 minutes to be up and running.

2. Ensure that the status of the IBM Spectrum Protect Plus instance is successful:
   a) In the navigation panel of the OpenShift web console, click **Operators** > **Installed Operators**.
   b) Click the IBM Spectrum Protect Plus operator from the list of installed operators.
   c) On the **Operator Details** page, click the IBM Spectrum Protect Plus tab. The list of running instances is displayed.
   d) Click the name of an instance to show its status.
   e) Scroll to the **Conditions** section of the page and review the status. The instance is created successfully when the `True` message is displayed in the **Status** column and the `Successful` message is displayed in the **Reason** column.
3. Ensure that the `spp` route is created and then start the IBM Spectrum Protect Plus user interface:
   a) In the navigation panel of the OpenShift web console, click **Networking** > **Routes**.
   b) On the **Routes Operators** page, ensure that the project that you created for IBM Spectrum Protect Plus is selected in the **Project** list.
   c) Ensure that the `spp` route is created.
   d) In the **Location** column, click the URL of the hostname that you configured when you created the IBM Spectrum Protect Plus instance.

      The IBM Spectrum Protect Plus user interface opens in a web browser.

      Starting the user interface requires the `virgo` server to load all the necessary bundles. This process can take approximately 10 - 12 minutes to complete. Make sure that you wait for the IBM Spectrum Protect Plus login page to be displayed.

What to do next

After you verify that the IBM Spectrum Protect Plus instance is running, you must create an IBM Spectrum Protect Plus vSnap server for your primary backup destination. The vSnap server is installed outside of your container environment. For instructions, see Chapter 5, “Installing and managing vSnap servers,” on page 153.

Related reference

“Troubleshooting installation issues for IBM Spectrum Protect Plus as a set of containers” on page 659
Review the available information to resolve installation issues with the IBM Spectrum Protect Plus server in a container environment.

Uninstalling IBM Spectrum Protect Plus containers

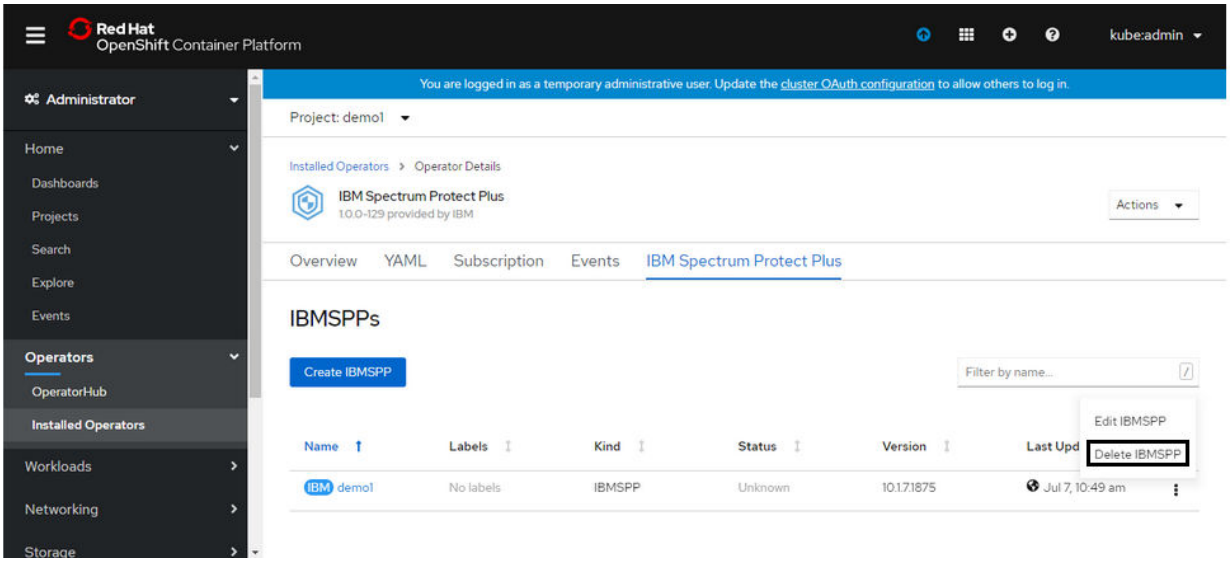
To uninstall IBM Spectrum Protect Plus from your container environment, you must uninstall the IBM Spectrum Protect Plus instance, and then uninstall the IBM Spectrum Protect Plus operator.

Uninstalling the IBM Spectrum Protect Plus instance

Uninstall the instance of IBM Spectrum Protect Plus from the OpenShift web console.

Procedure

1. Log on to the OpenShift web console as the cluster administrator.
2. In the navigation panel, click **Operators** > **Installed Operators**.
3. On the **Installed Operators** page, ensure that the namespace that you created for IBM Spectrum Protect Plus is selected in the **Project** list.
4. Click the IBM Spectrum Protect Plus operator.
5. On the operator details page for IBM Spectrum Protect Plus, click the **IBM Spectrum Protect Plus** tab.
6. Locate the installed IBM Spectrum Protect Plus instance in the table. Click the options button and click **Delete IBMSPP**.



What to do next

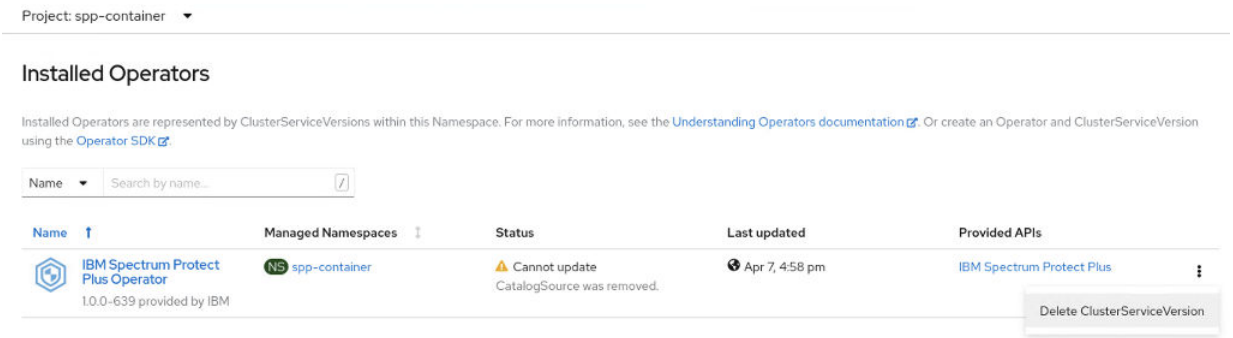
Uninstall the IBM Spectrum Protect Plus operator. For instructions, see "Uninstalling the IBM Spectrum Protect Plus operator" on page 151.

Uninstalling the IBM Spectrum Protect Plus operator

After you uninstall the IBM Spectrum Protect Plus instance, you can uninstall the IBM Spectrum Protect Plus operator from the OpenShift web console.

Procedure

1. Log on to the OpenShift web console as the cluster administrator.
2. In the navigation panel, click **Operators** > **Installed Operators**.
3. On the **Installed Operators** page, ensure that the namespace that you created for IBM Spectrum Protect Plus is selected in the **Project** list.
4. Locate the IBM Spectrum Protect Plus operator in the table. Click the options button and click **Delete ClusterServiceVersion**.



Chapter 5. Installing and managing 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 in “Component requirements ” on page 27.
2. Download the installation package. Different installation files are provided for installation on physical or virtual machines. Ensure that you download the correct files for your environment. For more information about downloading files and other useful information, see technote 567387.

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

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

Installing a physical vSnap server

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

Procedure

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

   See vSnap server physical installation 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:

© Copyright IBM Corp. 2017, 2021

```
$ 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 “Managing vSnap servers” on page 157.

Installing a virtual vSnap server in a VMware environment

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

Before you begin

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

For instructions, see “Assigning a static IP address” on page 121, Work with your network administrator when configuring network properties.

Procedure

1. Download the vSnap server template file *<part_number>*.ova from Passport Advantage Online. For information about downloading files, see technote 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 “Managing vSnap servers” on page 157.
For easier network administration, assign a static IP address for the virtual machine. Use the NetworkManager Text User Interface (nmtui) tool to assign the IP address.	For instructions, see “Assigning a static IP address” on page 121. 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” on page 121.

What to do next

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

Action	How to
Add the vSnap server to IBM Spectrum Protect Plus and configure the vSnap environment.	See “Managing vSnap servers” on page 157.

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 159 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 257.

   You can also use a user ID that has vSnap administrator privileges that you create by using the **`vsnap user create`** command. For more information about using console commands, see “vSnap server administration reference ” on page 176.
2. Run the following commands:

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

   ```
   $ rm -rf /etc/vsnap
   $ rm -rf /etc/nginx
   $ rm -rf /etc/uwsgi.d
   $ rm -f /etc/uwsgi.ini
   ```
4. Reboot the system to ensure kernel modules are unloaded and detach the data disks containing vSnap pool data.

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

Results

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

Managing vSnap servers

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 “Configuring backup storage options” on page 160.

If the vSnap server will also be registered as a VADP proxy, the account added in the **Storage Properties** field for the vSnap must have **sudo** privileges for the VADP proxy registration to succeed. For more information, see “Permission types ” on page 650.

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 257.
2. Run the **`vsnap user create`** command to create a user name and password for the vSnap server.
3. Start the IBM Spectrum Protect Plus user interface by entering the host name or IP address of the virtual machine where IBM Spectrum Protect Plus is deployed in a supported browser.
4. In the navigation panel, click **System Configuration** > **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 158.

   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 “Initializing the vSnap server” on page 169.
Expand the vSnap storage pool.	See “Configuring backup storage partners” on page 164.
If necessary, configure and administer certain aspects of vSnap, such as network configuration or storage pool management.	See “Configuring backup storage options” on page 160

Related tasks

“Start IBM Spectrum Protect Plus” on page 257
Start IBM Spectrum Protect Plus to begin using the application and its features.

Editing settings for a vSnap server

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

Procedure

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

1. In the navigation panel, click **System Configuration** > **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**.

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 159.

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 158.
   - Subsequent runs of backup jobs will automatically create volumes on this or another vSnap server in the same site and will perform new base backups.

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

- Do not unregister the vSnap server from IBM Spectrum Protect Plus. The deletion process will cause backups to be rebased.
- Replace the vSnap server. For information about replacing a failed, primary vSnap server, see this section "Troubleshooting vSnap servers" on page 183.

Procedure

To unregister a vSnap server, complete the following steps:

1. In the navigation panel, 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.

Configuring backup storage options

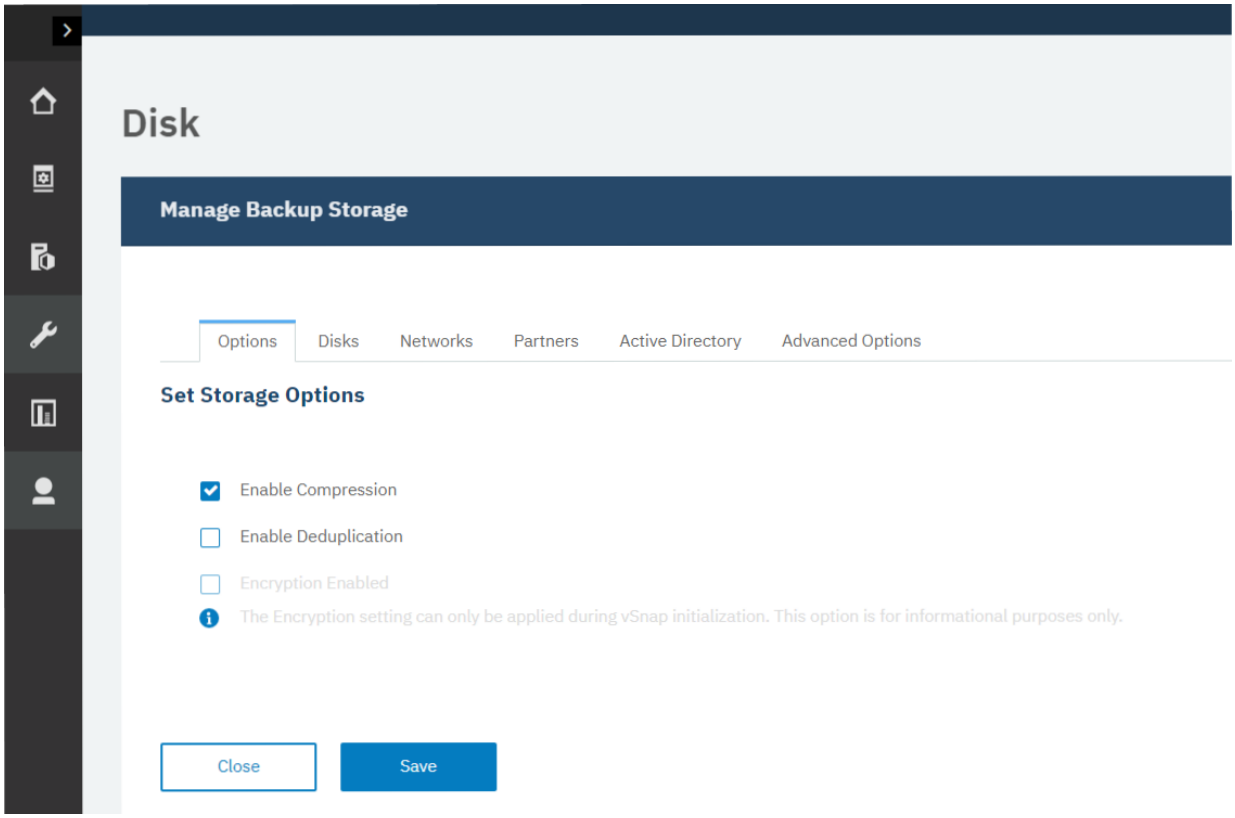
You can configure additional storage-related options for your primary and secondary backup storage hosts.

Procedure

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

1. In the navigation panel, 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**.

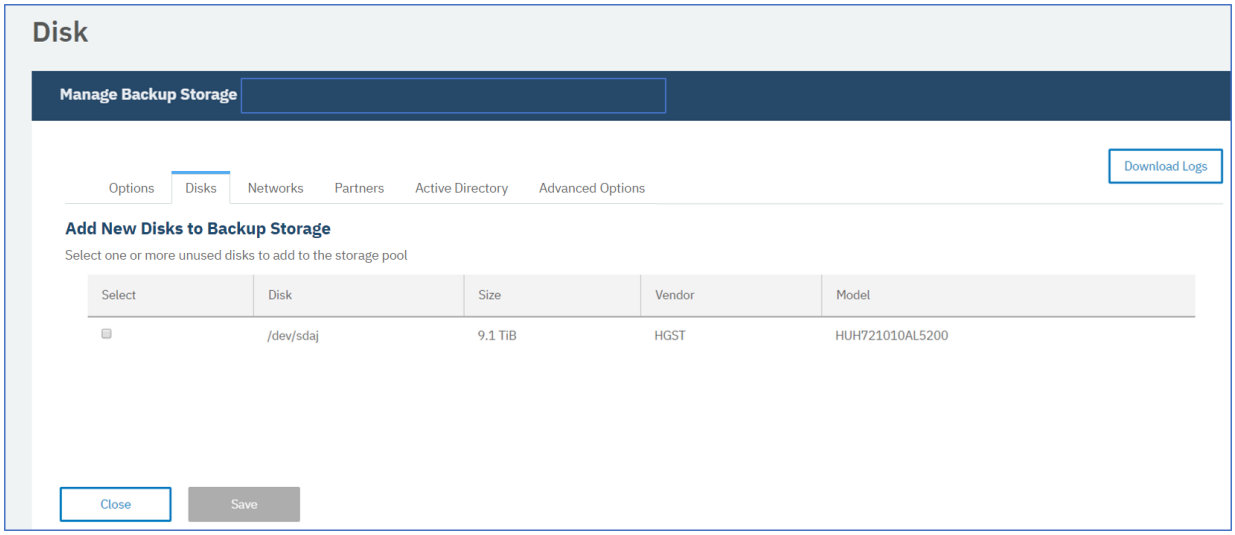
Adding new disks to backup storage

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.



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 panel, 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 panel, 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 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 248.

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.

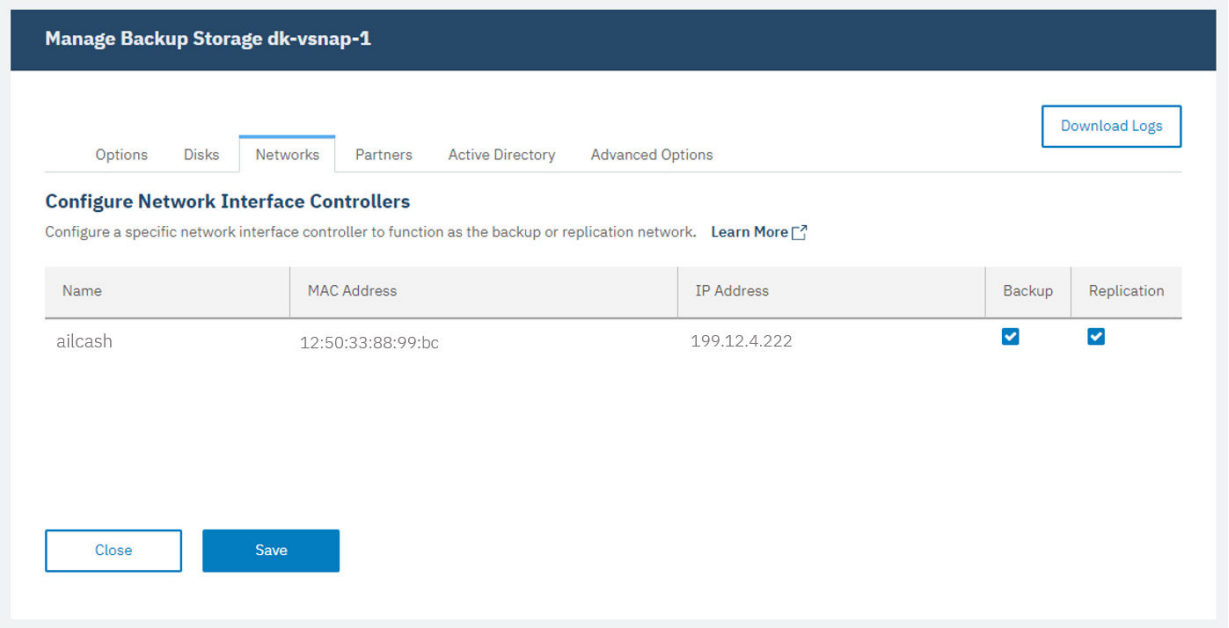
In the case that a second network interface card (NIC) is used, export `RMQ_SERVER_HOST` in the `/home/virgo/.bashrc` file on the IBM Spectrum Protect Plus appliance. The provided IP will be used for VADP to connect back to IBM Spectrum Protect Plus. The *<ip_address>* variable is the IP assigned to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



3. Click **Save**.

Configuring backup storage partners

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

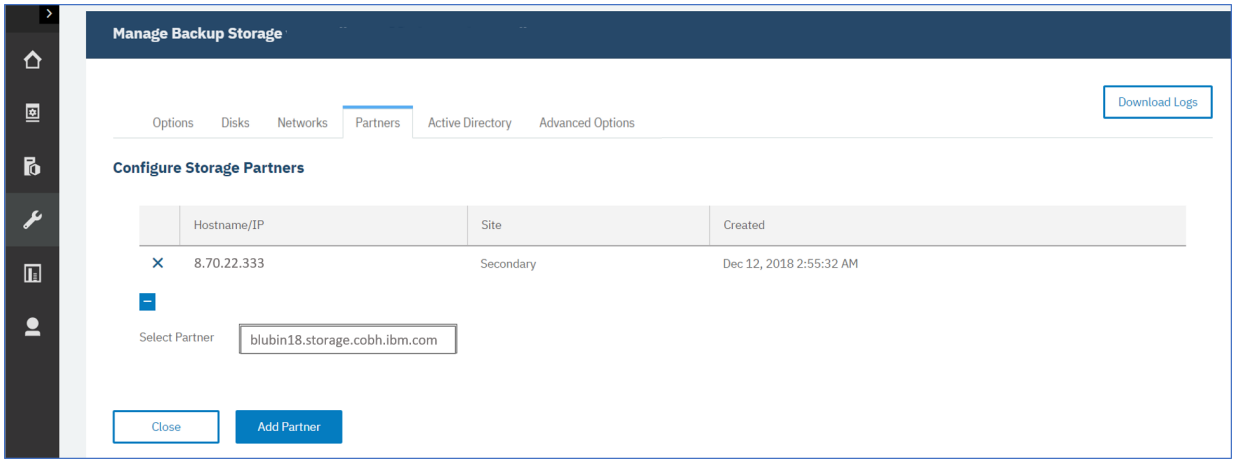
Before you begin

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

Procedure

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

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



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

Configuring an Active Directory

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

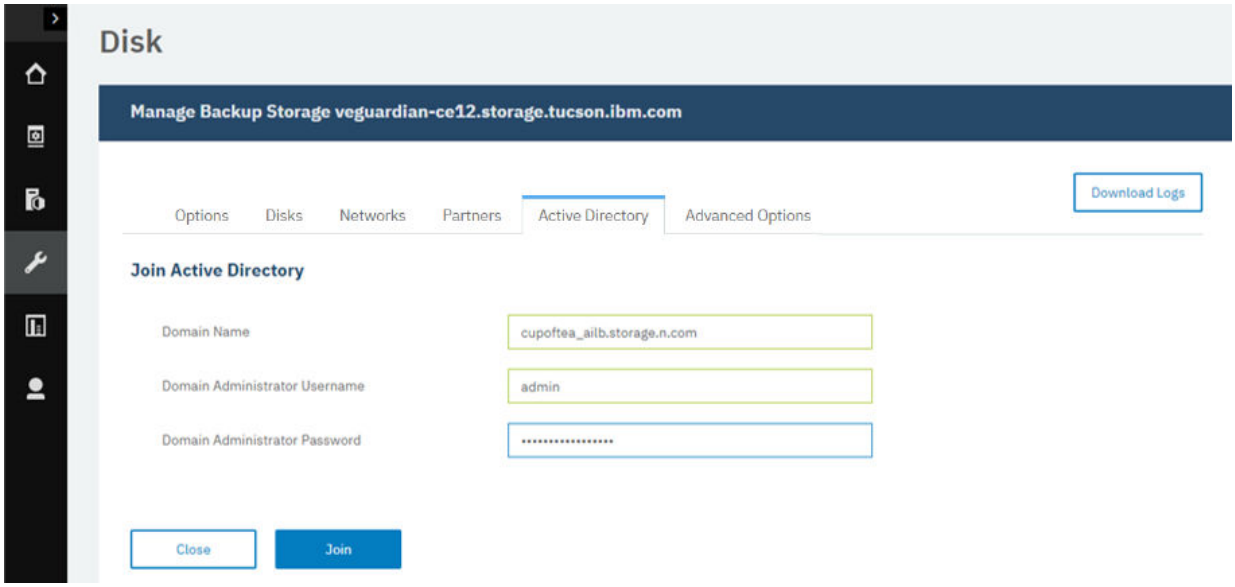
Before you begin

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

Procedure

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

1. In the navigation panel, 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.



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 panel, 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:

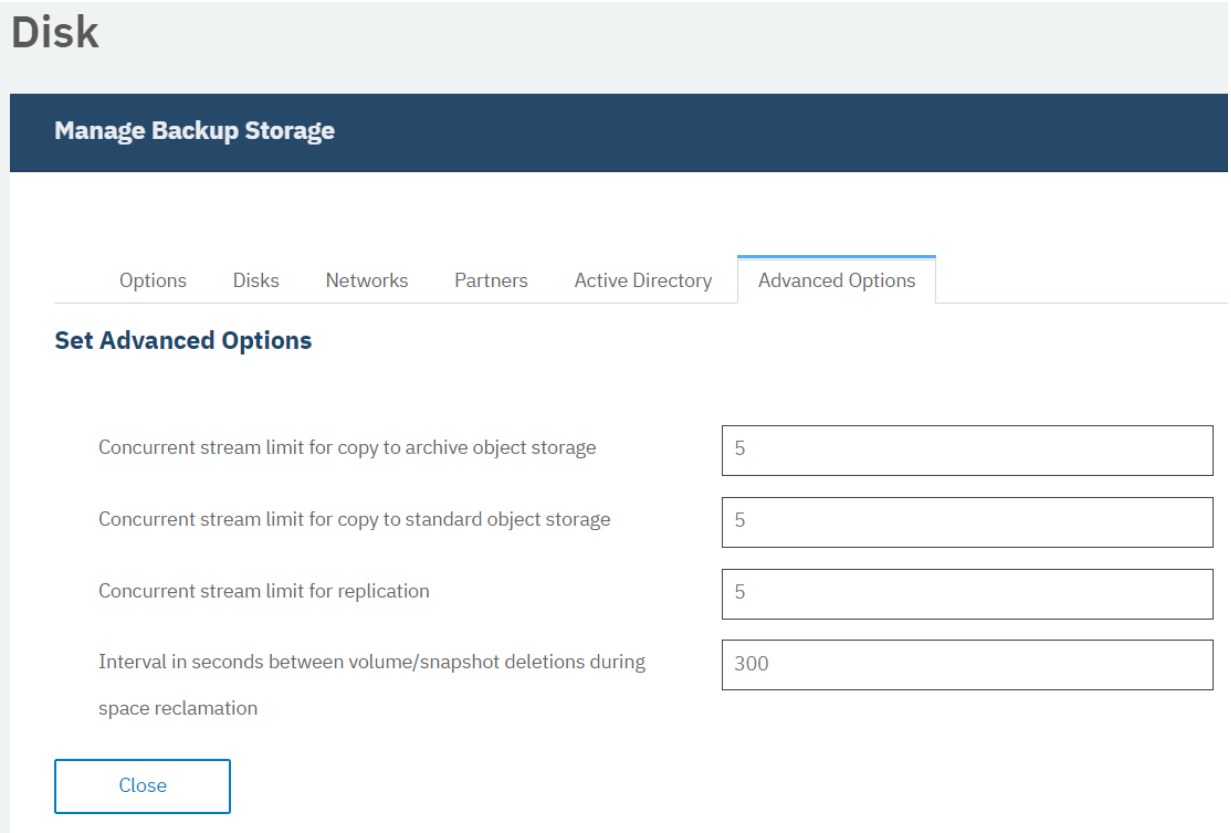


Figure 11. Manage backup storage advanced options.

- **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.

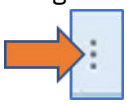
4. Click **Close**.

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 159.
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” on page 619.
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_pool`. 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     |
KNAME | NAME
-------------------------------------------------------------------------------------------
--------------
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 158.

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 178.

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

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

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

Completing a simple initialization

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

Procedure

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

1. In the navigation panel, click **System Configuration** > **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 257. If deployed physically, use the password that you created for the `serveradmin` account during installation.

   You can also use a user ID that has vSnap privileges that was previously created using the **`vsnap user create`** command. For more information about using console commands, see "vSnap server administration reference " on page 176.
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 178.

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 159. 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 250 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 154.
- 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 250.

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_pool`. 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 154.
   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 164.

j) Release schedules for all jobs that were paused in Step 3. Navigate to **Jobs and Operations** > **Schedule** and then click **Release All Schedules**.

9. Remove the vSnap software from the IBM Spectrum Protect Plus server.

a) Using secure shell (SSH), log in to the IBM Spectrum Protect Plus server as the **serveradmin** user.

b) Execute the **yum remove** commands to remove the vSnap server software from the IBM Spectrum Protect Plus server:

```
$ sudo yum remove vsnap
```

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

Results

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

Expanding a vSnap storage pool

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

Before you begin

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

Note:

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

Procedure

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

1. In the navigation panel, click **System Configuration** > **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 162.

Changing the throughput rate

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

Procedure

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

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

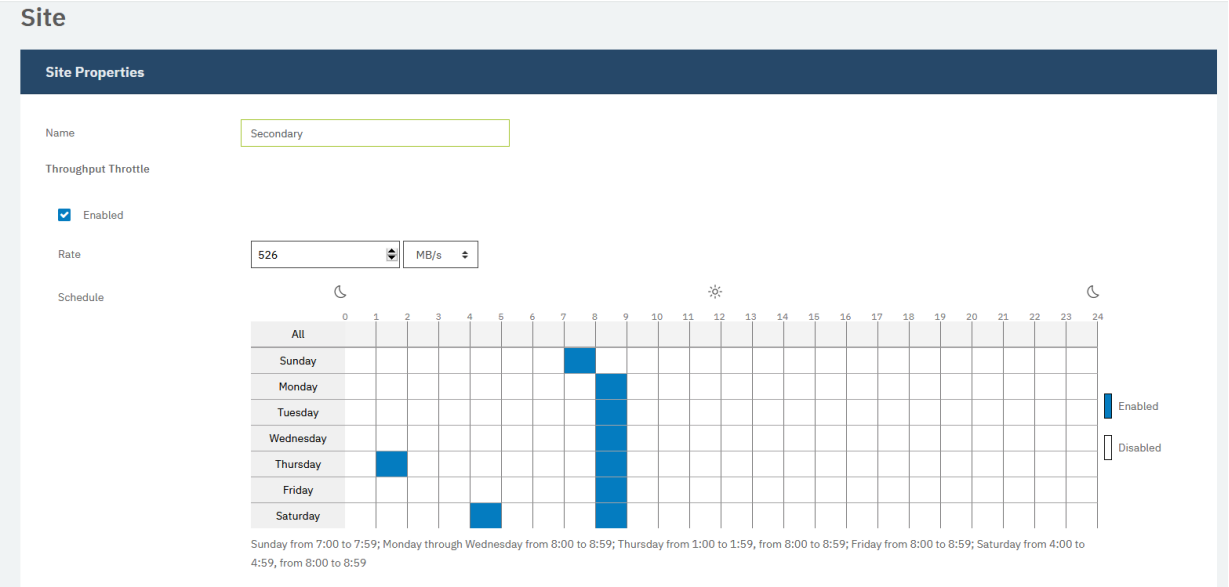


Figure 12. Enabling different throttles for different times to improve throughput

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

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

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

Replacing a failed vSnap server

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

Before you begin

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

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

About this task

The procedure for replacing a failed vSnap server is documented in technote 1103847.

Installing iSCSI initiator utilities

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

About this task

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

Procedure

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

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

vSnap server administration reference

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

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

Managing vSnap by using the command line interface

The vSnap server can be managed through the command-line interface and is the primary means of administering a vSnap server. Run the **vsnap** command from the vSnap server's interface after connecting through SSH using the user ID `serveradmin` or any other operating system user who has been assigned vSnap admin privileges. The initial `serveradmin` password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in “Start IBM Spectrum Protect Plus” on page 257.

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 panel. Click the manage icon for a vSnap server to edit its settings.

Related tasks

"Managing vSnap servers" on page 157
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 165
You can set advanced storage-related options for the primary or secondary backup storage in your environment.

User management

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

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

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

```
$ vsnap user create
```

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

--username *<username>*
Enter the username of the user.

--password *<password>*
Enter the password of the user.

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

```
$ vsnap user grant
```

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

--username *<username>*
Enter the username of the user.

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

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

```
$ vsnap user revoke
```

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

--username ***<username>***
Enter the username of the user.

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

```
$ vsnap user show
```

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

```
$ vsnap user update
```

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

--username ***<username>***
Enter the username of the user.

--password ***<old_password>***
Enter the old password of the user.

--new_password ***<new_password>***
Enter the new password of the user.

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 panel, 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 170, 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 179. 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 179.

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 panel. 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`**

Network management

Configure and administer network services for a vSnap server.

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

Showing network interface information

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

```
$ vsnap network show
```

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

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

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

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

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

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

Modifying services associated with network interfaces

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

```
$ vsnap network update
```

The following options are required:

--id <id>
Enter the ID of the interface to update.

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

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

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

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 257.
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 *xxxxxxxxx*`.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 xxxxxxxxx.rpm
   ```

Results

The kernel header or tool is installed.

Troubleshooting vSnap servers

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

Preventing job failures by synchronizing vSnap and CIFS passwords

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

About this task

For information about how to synchronize passwords, see "User management" on page 177.

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.

For information about how to set up and run a VMware restore job, see “Restoring VMware data” on page 350.

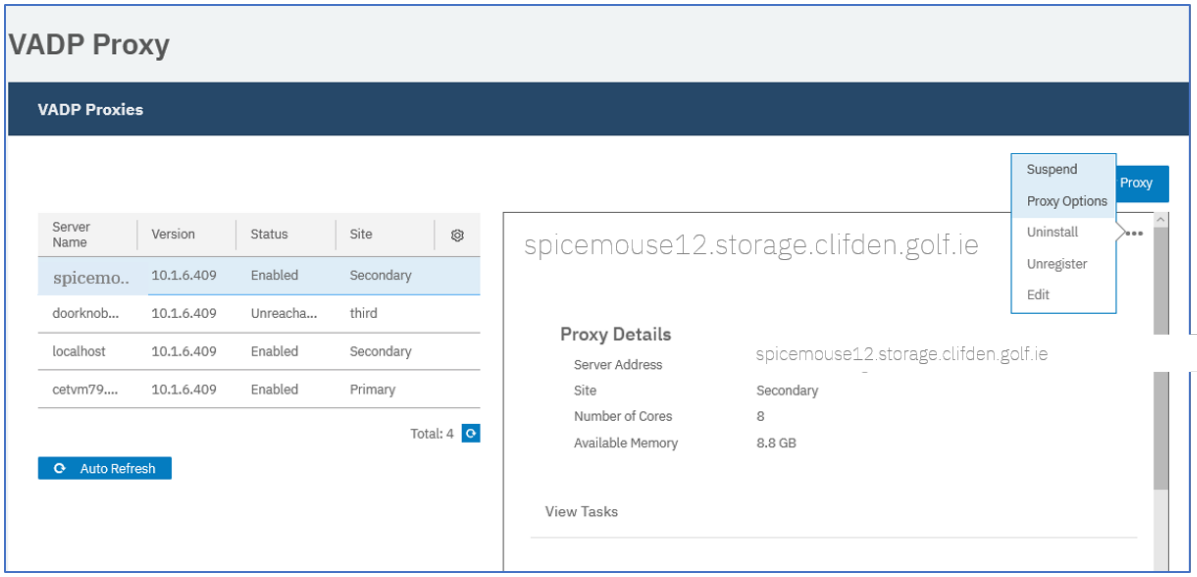
Communication

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.



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.

    For instructions, see “Creating jobs and job schedules” on page 620.

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

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

21. To resume jobs that failed after the vSnap server became unavailable, run a storage server inventory job. For instructions, see “Creating jobs and job schedules” on page 620.

Results

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

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

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

Chapter 6. Installing Container Backup Support

To protect persistent volumes of containers and cluster-scoped and namespace-scoped resources, you must install and configure IBM Spectrum Protect Plus Container Backup Support in a Kubernetes or Red Hat OpenShift Container Platform environment.

On Kubernetes and OpenShift clusters, you can install Container Backup Support by running an installation script at the command line. On OpenShift clusters, you can also install Container Backup Support on the OpenShift web console.

Installation prerequisites for Container Backup Support

Before you can install Container Backup Support on a Kubernetes or Red Hat OpenShift cluster, ensure that all system requirements and prerequisites are met.

For the system requirements for Container Backup Support, see "Container Backup Support requirements" on page 75. Ensure that the prerequisite software is up-to-date on your system and all security-related patches have been applied to the prerequisite software.

Ensure that IBM Spectrum Protect Plus vSnap server is registered with an IP address or fully qualified domain name (FQDN) and not just the hostname. For instructions, see "Registering a vSnap server as a backup storage provider" on page 158.

Then, to meet the prerequisites for Container Backup Support, complete the following actions:

- "Installing Helm 3 and renaming the binary file" on page 197
- "Installing Velero to protect resources" on page 198
- "For Kubernetes: Verifying whether Metrics Server is running" on page 204
- "Defining the application and persistent volume claim relationship" on page 205

Installing Helm 3 and renaming the binary file

The installation process for Container Backup Support uses a Helm 3 chart. The installation script that is provided with the installation package requires that the Helm 3 binary file is renamed to `helm3`.

About this task

Helm 3 is an application package manager that runs on Kubernetes or OpenShift. Helm is designed to simplify the definition, storage, and management of applications.

Procedure

1. Ensure that Helm 3 is installed on the host that is managing your cluster by issuing the following command:

   ```
   helm version --short
   ```

   If Helm 3 is installed, the output will display the Helm version.
2. If Helm 3 is installed, copy or rename the `helm` binary file in the `/usr/local/bin` directory to `helm3`.
3. If Helm 3 is not installed, you must install Helm 3 and rename the `helm` binary file to `helm3` because the Container Backup Support scripts use `helm3` as the binary name. The following steps include actions to rename the existing version of the Helm binary so that it can coexist with the Helm 3 binary:
   a) Verify whether an existing `helm` binary file (for example, for Helm 2) is installed in the `/usr/local/bin` directory. If so, rename the `helm` binary file to `helm2`.
   b) Issue the following commands to download Helm 3 and rename the `helm` binary file to `helm3`:

© Copyright IBM Corp. 2017, 2021

```
curl -fsSL -o get_helm.sh https://raw.githubusercontent.com/helm/helm/master/scripts/get-
helm-3
chmod 700 get_helm.sh
./get_helm.sh
sudo mv /usr/local/bin/helm /usr/local/bin/helm3
```

c) In the `/usr/local/bin` directory, rename the `helm2` binary back to `helm`.

Installing Velero to protect resources

If you plan to protect cluster-scoped and namespace-scoped resources in your container workloads, you must install the Velero tool.

For the version of Velero that is supported, see “Container Backup Support requirements” on page 75.

Space requirements

The MinIO Object Store serves as an S3 object store for snapshot backups. The MinIO Pod is integrated in the Container Backup Support (BaaS) installation package and is deployed to the BaaS namespace. This Pod claims a persistent volume with a size of 10 GB, and uses the default Storage Class based on the cluster configuration. If no default storage class exists, the **`minioStorageClass`** parameter in the `baas-values.yaml` file can be used to specify the storage class.

Protecting Kubernetes cluster-scoped and namespace-scoped resources

If you are planning to protect cluster-scoped and namespace-scoped resources in your container workloads, you must install and configure a dedicated version of Velero in your cluster. For instructions, see “For Kubernetes: Installing and configuring Velero” on page 198. If Velero is already installed for a different purpose, follow the instructions in “Installing a second instance of Velero ” on page 663.

Protecting Red Hat OpenShift cluster-scoped and namespace-scoped resources

If you are planning to protect cluster-scoped and namespace-scoped resources in your container workloads, you must install and configure a dedicated version of Velero in your cluster by using the OpenShift APIs for Data Protection (OADP) Operator. For instructions, see “For OpenShift: Installing and configuring Velero by using the OADP Operator ” on page 200. If Velero is already installed for a different purpose, follow the instructions in “Installing a second instance of Velero ” on page 663.

For Kubernetes: Installing and configuring Velero

To protect cluster-scoped resources and namespace-scoped resources in your container workloads, you must install and configure Velero in a dedicated namespace. The suggested default for this namespace is `spp-velero`.

Before you begin

If you installed an instance of Velero in the cluster for another purpose, you must install another instance of Velero for IBM Spectrum Protect Plus. Follow the instructions in “Installing a second instance of Velero ” on page 663.

About this task

Velero is required to protect cluster-scoped and namespace-scoped resources. If Velero is installed in a cluster in the default namespace, called `velero`, you can uninstall Velero and follow this procedure to install a new instance.

Procedure

1. Download the supported Velero installation package from the following VMware download site:

https://github.com/vmware-tanzu/velero/releases

To download Velero, run the following command in the `kubectl` command-line interface:

```
curl -fsSL -o velero-version-linux-amd64.tar.gz https://github.com/vmware-tanzu/velero/
releases/download/version/velero-version-linux-amd64.tar.gz
```

For example,

```
curl -fsSL -o velero-v1.5.2-linux-amd64.tar.gz https://github.com/vmware-tanzu/velero/
releases/download/v1.5.2/velero-v1.5.2-linux-amd64.tar.gz
```

2. Extract the .tar file by issuing the following command:

```
tar -xvf velero-version-linux-amd64.tar.gz
```

3. Change to the directory where the extracted .tar file is located and run the following command:

```
./velero install \
        --use-volume-snapshots=false \
        --no-default-backup-location \
        --no-secret \
        --plugins velero/velero-plugin-for-aws:v1.1.0 -n spp-velero
```

When the installation finishes, a message similar to the following message is displayed:

```
No secret file was specified, no Secret created.

No bucket and provider were specified, no default backup storage location created.

Velero is installed! ⛵ Use 'kubectl logs deployment/velero -n spp-velero' to view the status.
```

In this example, the Velero namespace parameter was named spp-velero by changing the **veleroNamespace** parameter in the `baas-values.yaml` file.

4. Ensure that Velero is installed and running by issuing the following command:

```
kubectl describe deployment velero -n spp-velero | grep Image
```

The following text block is an example of the output that is displayed:

```
Image:      velero/velero-plugin-for-aws:v1.1.0
Image:      velero/velero:v1.5.2
```

What to do next

If you no longer require Velero, uninstall it by issuing the following commands:

```
kubectl delete namespace/spp-velero clusterrolebinding/velero
kubectl delete crds -l component=velero
```

Tip:

After you installed Velero and Container Backup Support and before you can protect cluster-scoped and namespaces-scoped resources, you must run an inventory in the IBM Spectrum Protect Plus user interface so that the Velero instance is detected. For instructions, see "Detecting Kubernetes resources" on page 411.

During the inventory operation, the `BackupStorageLocation` that connects to the BaaS MinIO data mover is created automatically.

For OpenShift: Installing and configuring Velero by using the OADP Operator

To protect OpenShift cluster-scoped resources and namespace-scoped resources, you must use the OpenShift APIs for Data Protection (OADP) operator to install and configure the Velero tool in a dedicated namespace. The suggested name for the IBM Spectrum Protect Plus Velero namespace is `spp-velero`.

Before you begin

If you installed an instance of Velero in the cluster for another purpose, you must install another instance of Velero for IBM Spectrum Protect Plus. Follow the instructions in "Installing a second instance of Velero " on page 663.

Download the OADP operator from the Operator Hub.

Important: To use the OADP operator in Red Hat OpenShift Container Platform 4.5, you must complete extra steps to add the OADP operator to the Operator Hub. For instructions, see https://github.com/konveyor/oadp-operator/blob/master/docs/olm.md.

About this task

Tip: The term *namespace* is used to refer to Red Hat OpenShift *project*.

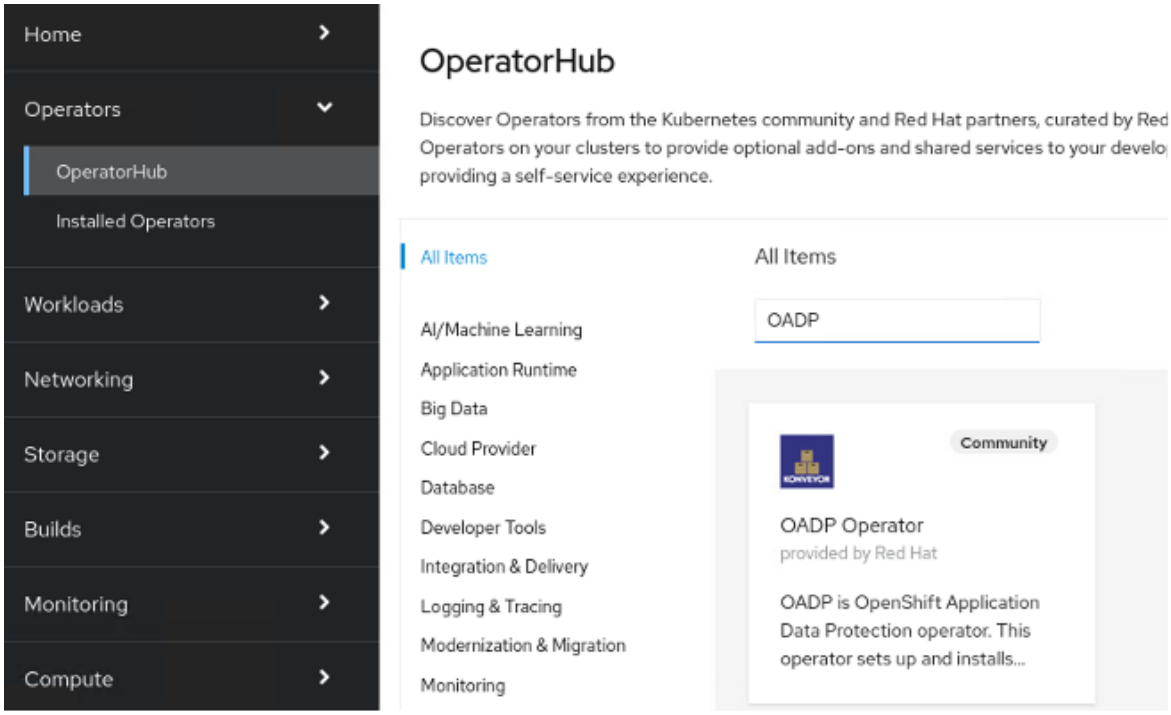
Procedure

1. Create an empty namespace called `spp-velero`.
2. Create a secret file with the following content:

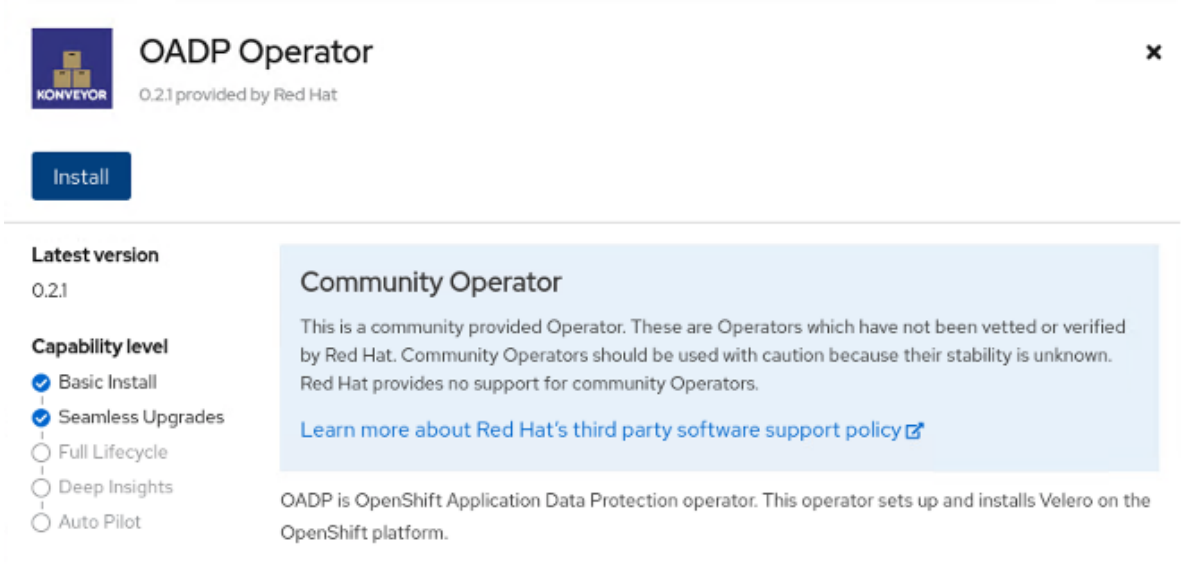
   ```
   [default]
   aws_access_key_id=
   aws_secret_access_key=
   ```

   The secret file does not need to have a user ID or a password, the fields can remain empty and unspecified.
3. Create a secret from the secret file that was created in the previous step by entering the following command:

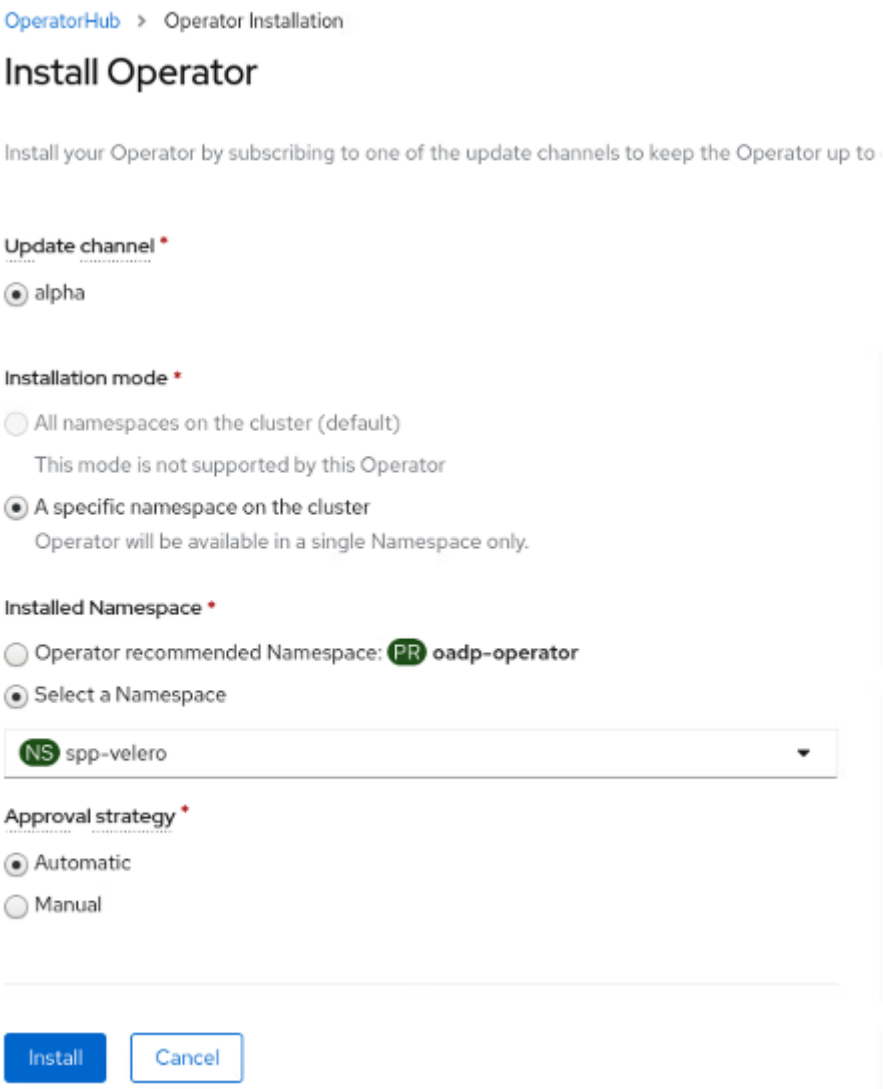
   ```
   oc create secret generic cloud-credentials
   --namespace spp-velero
   --from-file cloud=<path_to_secret_file>
   ```
4. Log on to the OpenShift web console as the cluster administrator.
5. In the navigation panel, click **Operators** > **OperatorHub**.
6. To install the OADP Operator, enter OADP in the search field. Click the **OADP Operator** card.



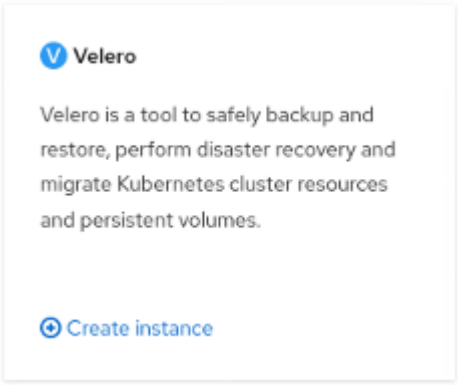
7. Click **Install**.



8. In the **Install Operator** page, change the default namespace name from `oadp-operator` to `spp-velero` and click **Install**.



9. In the **OADP Operator** page, look for the **Velero** card and click **Create Instance** to create a Velero custom resource (CR).

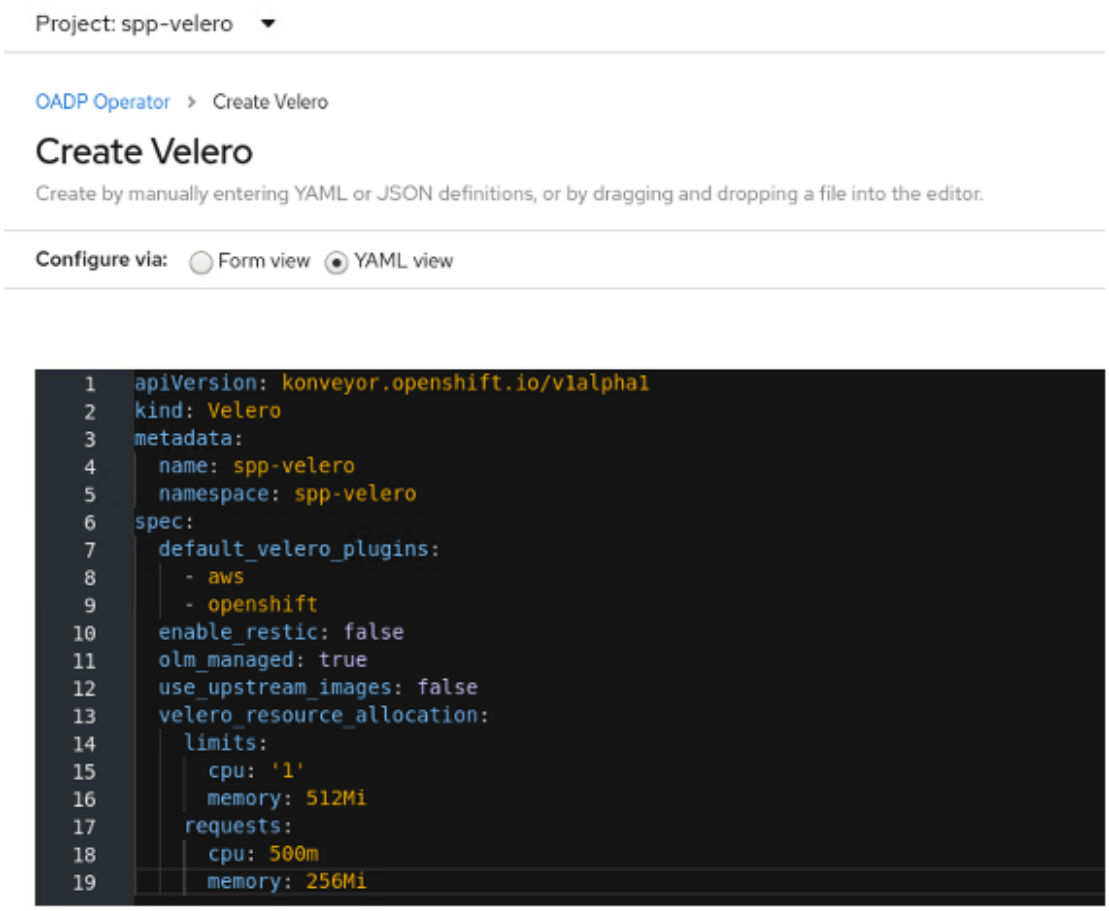


10. Edit the YAML file. Replace it with the following details and click **Create**.

```
apiVersion: konveyor.openshift.io/v1alpha1
kind: Velero
metadata:
  name: spp-velero
  namespace: spp-velero
spec:
  default_velero_plugins:
```

```
    - aws
    - openshift
  enable_restic: false
  olm_managed: true
  use_upstream_images: false
  velero_resource_allocation:
    limits:
      cpu: '1'
      memory: 512Mi
    requests:
      cpu: 500m
      memory: 256Mi
```

For example:



11. Ensure that Velero is installed and running by issuing the following command at the **oc** command line:

```
oc describe deployment velero -n spp-velero | grep Image
```

The following text block is an example of the output that is displayed:

```
Image:      quay.io/konveyor/openshift-velero-plugin:oadp-0.2.3
Image:      quay.io/konveyor/velero-plugin-for-aws:oadp-0.2.3
Image:      quay.io/konveyor/velero:oadp-0.2.3
Image:      quay.io/konveyor/velero:oadp-0.2.3
```

What to do next

If you no longer require Velero, uninstall it by deleting the Velero instance in the OADP operator page in the OpenShift web console.

Tip:

After you installed Velero and Container Backup Support and before you can protect cluster-scoped and namespaces-scoped resources, you must run an inventory in the IBM Spectrum Protect Plus user interface so that the Velero instance is detected. For instructions, see "Detecting OpenShift cluster resources" on page 433.

During the inventory operation, the `BackupStorageLocation` that connects to the BaaS MinIO data mover is created automatically.

For Kubernetes: Verifying whether Metrics Server is running

As an optional step, to help optimize product performance and scalability, ensure that Kubernetes Metrics Server v0.3.5 or later is installed and running properly on your Kubernetes cluster.

Before you begin

For instructions about deploying Metrics Server, review the `README.md` file at https://github.com/kubernetes-sigs/metrics-server. For general information about Kubernetes Metrics Server, see Resource metrics pipeline.

About this task

Metrics Server is used by the Container Backup Support scheduler to determine the resources that are used by concurrent data mover instances.

If Metrics Server does not return data, a limited number of data movers are used for backup operations. This limitation might negatively impact performance.

Procedure

Verify that Metrics Server is installed and returning metrics by completing the following steps:

1. Verify the installation by issuing the following command:

   ```
   kubectl get deploy,svc -n kube-system | egrep metrics-server
   ```

   If Metrics Server is installed, the output is similar to the following example:

   ```
   deployment.extensions/metrics-server   1/1     1            1           3d4h
   service/metrics-server   ClusterIP   198.51.100.0   <none>        443/TCP         3d4h
   ```

2. Verify that Metrics Server is returning data for all nodes by issuing the following command:

   ```
   kubectl get --raw "/apis/metrics.k8s.io/v1beta1/nodes"
   ```

   If Metrics Server is returning data for all nodes, the output is similar to the following example:

   ```
   {"kind":"NodeMetricsList","apiVersion":"metrics.k8s.io/v1beta1","metadata":{"selfLink":"/
   apis/metrics.k8s.io/v1beta1/nodes"},"items":[{"metadata":
   {"name":"cirrus12","selfLink":"/apis/metrics.k8s.io/v1beta1/nodes/cirrus12",
   "creationTimestamp":"2019-08-08T23:59:49Z"},"timestamp":"2019-08-08T23:59:08Z",
   "window":"30s","usage":{"cpu":"1738876098n","memory":"8406880Ki"}}]}
   ```

Results

The command might fail with empty output for the "items" key. This error is likely caused by installing Metrics Server with a self-signed certificate. To resolve this issue, install Metrics Server with a correctly signed certificate that is recognized by the cluster.

Defining the application and persistent volume claim relationship

You can optionally tie your stateful applications to their persistent volume claims (PVCs) by defining an owner-dependent relationship. By defining this relationship, you enable cascading actions for the applications.

For example, if an owner-dependent relationship is defined, scaling up and scaling down an application can cause the scheduled backups of its PVC to be paused and resumed. Similarly, deleting the application causes the deletion of the PVC, which in turn triggers the deletion of the backups.

After an application starts to use a PVC to store persistent data, you can reconfigure the PVC definition with its owner application.

For guidance, see the following sample configuration file for a PVC, which shows the owner-dependent relationship between an application and a PVC object. The PVC object includes the details of the owner deployment in the **ownerReferences** field:

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: demo-pvc
  ownerReferences:
  - apiVersion: apps/v1beta1
    blockOwnerDeletion: true
    kind: Deployment
    name: Dept10-deployment
    uid: 3b760e89-7da5-11e9-8c5a-0050568ba59c
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 1Gi
  storageClassName: csi-rbd
```

For Kubernetes or OpenShift: Installing Container Backup Support by using the command line

You can use the command line to install and deploy Container Backup Support on a Kubernetes or OpenShift cluster.

About this task

A Helm 3 chart is used to install and deploy Container Backup Support on your Kubernetes or OpenShift cluster. Scripts are provided in the installation package to deploy Container Backup Support on your cluster.

During the installation and deployment procedure, you must update the `baas-options.sh` and `baas-values.yaml` files with specifications for your environment, and then run an installation script. When you run the installation script, Helm 3 is used to deploy Container Backup Support in your environment.

You can install Container Backup Support by using one of the following methods:

Table 70. Methods for installing Container Backup Support	
Installation method	**Steps**
By downloading and installing the product package in an airgap environment	The installation package from IBM Passport Advantage Online is a larger but self-contained package. Internet access is not required at deployment time. Complete the following steps: 1. “Setting up the installation variables” on page 206 2. “Installing Container Backup Support in an airgap environment” on page 214
By fetching and installing the product package from IBM Helm Charts Repository and IBM Entitled Registry	The Helm package is smaller in size and therefore takes less time to download. Internet access is required to pull containers at deployment time. Complete the following steps: 1. “Setting up the installation variables” on page 206 2. “Installing Container Backup Support from IBM Helm Charts Repository and IBM Entitled Registry” on page 218

Restrictions:

- To upgrade from a previous version, you must uninstall the previous version and then install Container Backup Support Version 10.1.8 by using configuration values from the previous version. In Versions 10.1.6 and earlier, the configuration values are set up in the `baas-config.cfg` file. In V10.1.8, the configuration values are set up in the `baas-options.sh` and `baas-values.yaml` files.
- A rollback to a previous version is not supported. In other words, you cannot use previous versions of Container Backup Support to restore data that was backed up by Container Backup Support V10.1.8.
- Due to underlying changes in the `BaasReq` object, you cannot use Container Backup Support V10.1.8 to restore data that was backed up by V10.1.5.

Setting up the installation variables

A script is used as part of the installation process for Container Backup Support. Set up the environment and installation variables that are used by the installation script.

You must set up the variables in the following two files:

`baas-options.sh`
Contains the variables that are used to configure the prerequisites for Container Backup Support. This file is used to replace the sample `baas-options.sh` file that is provided in the installation package.

`baas-values.yaml`
Contains the values that are used to install Container Backup Support or to update an existing configuration. This file is used to replace the sample `baas-values.yaml` file that is provided in the installation package.

Information is provided on how to set up the installation variables for installing Container Backup Support in an airgap environment or from the IBM Helm Charts Repository that is linked to IBM Entitled Registry.

About the airgap environment: As used in this documentation, an airgap cluster is any Kubernetes or OpenShift cluster that does not have internet access, and therefore cannot pull container images from a Docker registry. The airgap installation package includes the container images for Container Backup Support. During the installation, the container images are loaded by the **`docker load`** command and then tagged and pushed to the Docker registry that is specified in the `baas-options.sh` file. As a result, the container images can be pulled during the Helm installation.

To set up the variables in the `baas-options.sh` file, see "Setting up installation variables in the baas-options.sh file" on page 207.

To set up the configuration parameters in the `baas-values.yaml` file, see "Setting up installation variables in the baas-values.yaml file" on page 209.

Setting up installation variables in the `baas-options.sh` file

Set up the variables in the `baas-options.sh` file to configure the prerequisites for Container Backup Support. Use this file to replace the sample `baas-options.sh` file that is provided in the installation package.

Before you begin, create a directory in your home folder (~) for storing the `baas-options.sh` and `baas-values.yaml` files. Issue the following commands:

```
mkdir install_vars_dir
cd install_vars_dir
```

where *install_vars_dir* is the name of the directory that you created.

Obtain an entitlement key

If you plan to pull Container Backup Support images from the IBM Entitled Registry, you must obtain a key from the IBM Container Library. To obtain an entitlement key:

1. Log in to the IBM Container software library with the IBMid and password that is associated with the entitled software.
2. Click **Get entitlement key**.
3. In the **Access your container software** page, click **Copy key** to copy the generated entitlement key.
4. Save the key to a secure location for later use.

Set up the variables in `baas-options.sh`

Update the values in the following text block and save the text block to a file called `baas-options.sh` in the *install_vars_dir* directory that you created:

```
export DOCKER_REGISTRY_ADDRESS='your_docker_registry'
export DOCKER_REGISTRY_USERNAME='your_docker_username'
export DOCKER_REGISTRY_PASSWORD='your_docker_password'
export DOCKER_REGISTRY_NAMESPACE='your_docker_registry_namespace'
export SPP_ADMIN_USERNAME='your_protectplus_containers_admin_username'
export SPP_ADMIN_PASSWORD='your_protectplus_containers_admin_password'
export DATAMOVER_USERNAME='create_a_datamover_username'
export DATAMOVER_PASSWORD='create_a_datamover_password'
export MINIO_USERNAME='create_a_minio_username'
export MINIO_PASSWORD='create_a_minio_password'
export BAAS_VERSION='protectplus_version'
```

The following table contains the descriptions for the environment variables in the `baas-options.sh` file. You must enclose the values with single quotation marks (`' '`).

Table 71. Installation variables in the `baas-options.sh` file

Environment variable	Description
DOCKER_REGISTRY_ADDRESS	The address of the Docker registry in your environment where container images are loaded. If you are pulling images from the IBM Entitled Registry, you must specify 'cp.icr.io/cp'. The value for DOCKER_REGISTRY_ADDRESS must match the value for the **`imageRegistry`** parameter in the `baas-values.yaml` file.

Table 71. Installation variables in the baas-options.sh file (continued)

Environment variable	Description
DOCKER_REGISTRY_USERNAME	The user account for the Docker registry where container images are loaded. If you are pulling images from the IBM Entitled Registry, you must specify 'cp'.
DOCKER_REGISTRY_PASSWORD	The user password for the Docker registry where the container images are loaded. To pull images from the IBM Entitled Registry, specify the entitlement key that you obtained from the IBM Container software library. You can avoid putting the password in the file by specifying an environment variable for any of the passwords. For example, ${DOCKERUSER_PW} or ${IBMCLOUD_API_KEY}.
DOCKER_REGISTRY_NAMESPACE	The namespace of the Docker registry where the container images are loaded. The namespace does not have to be created ahead of time. To pull images from the IBM Entitled Registry, you must specify 'sppc'. The value for DOCKER_REGISTRY_NAMESPACE must match the value for the **imageRegistryNamespace** parameter in the baas-values.yaml file.
SPP_ADMIN_USERNAME	The user ID of the IBM Spectrum Protect Plus containers administrator. The containers administrator is an IBM Spectrum Protect Plus administrator with the **Containers Admin** role.
SPP_ADMIN_PASSWORD	The IBM Spectrum Protect Plus password for the containers administrator. You can optionally specify an environment variable for the password. For example, ${PROTECTPLUS_ADMIN_PW}.
DATAMOVER_USERNAME	The user ID to create for use with the data mover. The value does not have to exist already. It is created for the installation. The data mover username must adhere to the rules for usernames and passwords for Red Hat Enterprise Linux (RHEL) 7 operating system. The rules are the same as the ones for creating a new user on RHEL 7. For example, the password and the username must not be the same.

Table 71. Installation variables in the baas-options.sh file (continued)	
Environment variable	**Description**
DATAMOVER_PASSWORD	The user password to create for use with the data mover. The value does not have to exist already. It is created for the installation. The data mover password must adhere to the rules for usernames and passwords for RHEL 7. The rules are the same as the ones for creating a new user on RHEL 7. For example: • The password must be at least 8 characters in length, and must contain letters and numbers. • No dictionary words are allowed in the password. • The password cannot be the same as the username.
MINIO_USERNAME	The username to create for the MinIO user. MinIO object storage is used to store backups of cluster and namespace resources. The value does not have to exist already. It is created for the installation.
MINIO_PASSWORD	The password to create for the MinIO user. The value does not have to exist already. It is created for the installation.
BAAS_VERSION	The version of IBM Spectrum Protect Plus that you are installing, for example, 10.1.8 or 10.1.8.*x*, where *x* specifies an interim fix (iFix).

Setting up installation variables in the baas-values.yaml file

Set up the variables that are used to install or update Container Backup Support in the `baas-values.yaml` file. Use this file to replace the sample `baas-values.yaml` file that is provided in the installation package.

Before you begin, create a directory in your home folder (~) for storing the `baas-options.sh` and `baas-values.yaml` files. Issue the following commands:

```
mkdir install_vars_dir
cd install_vars_dir
```

where *install_vars_dir* is the name of the directory that you created.

Set up the variables in baas-values.yaml

Update the values in the following text block and save the text block to a file called `baas-values.yaml` in the *install_vars_dir* directory that you created.

```
license: false | true
isOCP: false | true
clusterName: create_a_cluster_name
networkPolicy:
  clusterAPIServerips:
    - kubernetes_host_ip1
    - kubernetes_host_ip2
    - kubernetes_host_ip3
  clusterAPIServerport: your_cluster_api_port
  clusterCIDR: x.x.x.x/y
  isServerInstalledOnAnotherCluster: false
SPPfqdn: your_protectplus_server_DNS_address
SPPips: your_protectplus_server_ip
SPPport: your_protectplus_server_port
productLoglevel: INFO | WARNING | ERROR | DEBUG
imageRegistry: your_docker_registry
imageRegistryNamespace: your_docker_registry_namespace
minioStorageClass: name_of_storageclass_to_use_with_minio
veleroNamespace: spp-velero
```

Ensure that the spacing is maintained as specified in the YAML file. Tabs are not allowed.

Configuring baas-values.yaml when IBM Spectrum Protect Plus server runs in a container environment: If the IBM Spectrum Protect Plus server is installed on an OpenShift cluster, ensure that you set the values for the **isServerInstalledOnAnotherCluster**, **SPPfqdn**, and **SPPips** accordingly. For the specifications of these parameters, see Table 72 on page 210.

The following table contains the descriptions and default values for the configuration parameters in the `baas-values.yaml` file:

Table 72. Configuration parameters in the `baas-values.yaml` file

Parameter	Description	Default value
license	The product license for Container Backup Support. The English license file is located in the `LICENSES/LICENSE-en` directory, which is included in the installation package. Versions of the license in English and other languages are available in the "IBM Spectrum Protect Plus Capacity - Version 10.1.8" license agreements at License Information documents. Set the value to `true` to indicate that you have reviewed and agree to the license agreement.	false
isOCP	The type of cluster on which you are installing Container Backup Support. If you are installing the product on an OpenShift cluster, set the value to `true`. If you are installing the product on a Kubernetes cluster, set the value to `false`.	false
clusterName	The unique cluster name that is used to register the application host to the IBM Spectrum Protect Plus server. The cluster name can be any name of your choice, but it must be unique from the IBM Spectrum Protect Plus server.	None
clusterAPIServerips	The IP address for the cluster API server. To obtain the cluster API server address, issue the following command: For Kubernetes: `kubectl get endpoints -n default -o yaml kubernetes` For OpenShift: `oc get endpoints -n default -o yaml kubernetes` Use all of the provided addresses listed under the **addresses** field in the output, or add or remove IP addresses as needed. Specify multiple addresses as follows: `networkPolicy:` `  clusterAPIServerips:` `    - x.x.x.x` `    - y.y.y.y` `    - z.z.z.z`	*x.x.x.x*

Table 72. Configuration parameters in the baas-values.yaml file (continued)

Parameter	Description	Default value
clusterAPIServerport	The port address for the cluster API server. To obtain the cluster API server port, issue the following command: For Kubernetes: `kubectl get endpoints -n default -o yaml kubernetes` For OpenShift: `oc get endpoints -n default -o yaml kubernetes` Use the port number listed in the **port** field in the output.	6443
clusterCIDR	The Classless Inter-Domain Routing (CIDR) value for the cluster. To obtain the CIDR, issue the following command: For Kubernetes: `kubectl cluster-info dump | grep -m 1 cluster-cidr` For OpenShift: `oc get network -o yaml | grep -A1 clusterNetwork:` Use the displayed IP address as the cluster CIDR address. **Tip for Kubernetes:** If the command does not return the CIDR value, change the **grep** expression to look for the combination of "cluster" and "CIDR" and run the command again.	192.168.0.0/16
additionalAddresses	Enter a value for the **dm.nfs.additionalAddresses** field if the network environment contains routers, proxies, gateways, or Network Address Translation (NAT) devices between the cluster nodes and the IBM Spectrum Protect Plus vSnap server. This address is used for data transfer during backup and restore operations between the cluster and vSnap server. Enter the public facing addresses for the device. The vSnap server uses an IP filter to regulate data transfer. You can specify more than one address, which can be in IP address or CIDR format. The addresses must be on the allowlist. Otherwise, the addresses are rejected. You can specify multiple addresses as follows: `dm:` `  nfs:` `    additionalAddress:` `      - x.x.x.x` `      - y.y.y.y` `      - z.z.z.z`	*x.x.x.x*
isServerInstalledOnAnotherCluster	For OpenShift environments, set the value to `false` regardless of whether the IBM Spectrum Protect Plus server is installed on the same cluster as Container Backup Support or on a separate cluster. If you are installing Container Backup Support on a Kubernetes cluster, or if the IBM Spectrum Protect Plus server is installed as a virtual appliance, set the value to `false`.	false

Table 72. Configuration parameters in the baas-values.yaml file (continued)

Parameter	Description	Default value
SPPfqdn	The DNS address for the IBM Spectrum Protect Plus server. You can specify an IP address or a fully qualified domain name (FQDN). If the IBM Spectrum Protect Plus server is installed as a virtual appliance and no DNS server is available, specify the IP address that is used for the **SPPips** parameter. If the IBM Spectrum Protect Plus server is installed in an OpenShift container environment, retrieve the DNS address by issuing the following command: `oc get route --namespace spp_server_namespace` where *spp_server_namespace* specifies the namespace in which the IBM Spectrum Protect Plus server is installed. The DNS address to use is listed in the HOST/PORT column in the command output. For example: `NAME  HOST/PORT  PATH SERVICES` `spp-rte` **`my.plus.server.example`** `/  sppproxy`	None
SPPips	The IBM Spectrum Protect Plus server IP address. If the IBM Spectrum Protect Plus server is installed as a virtual appliance, specify an IP address. If the IBM Spectrum Protect Plus server is installed in an OpenShift container environment, obtain the IP address as follows: 1. Begin with the FQDN of the route that was used in the description for the **SPPfqdn** parameter. 2. Using the URL from the output of the **oc get route** query, run the **nslookup** command or any DNS tools to resolve the URL to the IP address of the route. You can use one of the following methods: • Run the following command by using the example from the **SPPfqdn** description: `nslookup my.plus.server.example` • For more reliable results, create the Kubernetes dnsutils pod and run a DNS lookup of the IP address from that pod. For more information, see Debugging DNS Resolution. For example, run the following command, where *my.plus.server.example* is the FQDN of the route: `kubectl exec -i -t dnsutils -- nslookup` `my.plus.server.example`	*x.x.x.x*
SPPport	The IBM Spectrum Protect Plus server port. You must set the port number to 443.	443

Table 72. Configuration parameters in the baas-values.yaml file (continued)		
Parameter	**Description**	**Default value**
productLoglevel	The trace levels for troubleshooting issues with the Container Backup Support transaction manager, controller, and scheduler components. The following trace levels are available: INFO, WARNING, DEBUG, and ERROR.	INFO
imageRegistry	The address of the Docker registry in your environment where the container images are loaded. If you are pulling images from the IBM Entitled Registry, you must specify cp.icr.io/cp. The value for the **imageRegistry** parameter must match the value for the DOCKER_REGISTRY_ADDRESS variable in the baas-options.sh file.	*docker-repo-hostname*: 5000
imageRegistryNamespace	The namespace of the Docker registry where the container images are loaded. The namespace does not have to be created ahead of time. To pull images from the IBM Entitled Registry, you must specify sppc. The value for the **imageRegistryNamespace** parameter must match the value for the DOCKER_REGISTRY_NAMESPACE variable in the baas-options.sh file.	baas
minioStorageClass	The name of the storage class to use for the MinIO server. The MinIO server is used to store the backups of cluster and namespace resources. If you do not specify a value for this parameter, the default storage class of your cluster is used. Ensure that a default storage class is defined. **Important:** To safeguard resource snapshot backups in the case where the BaaS is uninstalled or has been reinstalled, set the storage class with a **Reclaim Policy** with the Retain value specified. Backups that have been transferred to the vSnap server are not affected. Certain upgrade scenarios may also lead to losing the minIO PVC content if the **Reclaim Policy** is not set to Retain.	None
veleroNamespace	Specify the namespace of the Velero installation that is dedicated to IBM Spectrum Protect Plus Container Backup Support, for example, spp-velero. If you do not specify a value for this parameter, Velero integration is unavailable and you can use Container Backup Support to protect only PVCs.	None

Examples of baas-options.sh files

The following table shows examples of the baas-options.sh file for installations in different environments.

Table 73. Examples of baas-options.sh files

Kubernetes installation with a Docker registry in an airgap environment	OpenShift installation with the IBM Entitled Registry
`export DOCKER_REGISTRY_ADDRESS='192.0.2.28:5000'` `export DOCKER_REGISTRY_USERNAME='dockeruser'` `export DOCKER_REGISTRY_PASSWORD=${DOCKER_PW}` `export DOCKER_REGISTRY_NAMESPACE='baas'` `export SPP_ADMIN_USERNAME='container-admin'` `export SPP_ADMIN_PASSWORD=${SPP_ADMIN_PW}` `export DATAMOVER_USERNAME='spectrum'` `export DATAMOVER_PASSWORD='Pr0tect!'` `export MINIO_USERNAME='spp-user'` `export MINIO_PASSWORD='aust1np0w3rs'` `export BAAS_VERSION='10.1.8'`	`export DOCKER_REGISTRY_ADDRESS='cp.icr.io/cp'` `export DOCKER_REGISTRY_USERNAME='cp'` `export DOCKER_REGISTRY_PASSWORD=$` `{IBMCLOUD_API_KEY}` `export DOCKER_REGISTRY_NAMESPACE='sppc'` `export SPP_ADMIN_USERNAME='container-admin'` `export SPP_ADMIN_PASSWORD=${SPP_ADMIN_PW}` `export DATAMOVER_USERNAME='spectrum'` `export DATAMOVER_PASSWORD='Pr0tect!'` `export MINIO_USERNAME='spp-user'` `export MINIO_PASSWORD='aust1np0w3rs'` `export BAAS_VERSION='10.1.8'`

Examples of `baas-values.yaml` files

The following table shows examples of the `baas-values.yaml` file for installations in different environments.

Table 74. Examples of baas-values.yaml files

Kubernetes installation with a Docker registry in an airgap environment	OpenShift installation with the IBM Entitled Registry
`license: true` `isOCP: false` `clusterName: example-k8s-cluster` `networkPolicy:` `  clusterAPIServerips:` `  - 192.0.2.63` `  clusterAPIServerport: 6443` `  clusterCIDR: 192.168.0.0/16` `  isServerInstalledOnAnotherCluster: false` `SPPfqdn: my.ova.plus.server.example` `SPPips: 192.0.2.83` `SPPport: 443` `productLoglevel: INFO` `imageRegistry: 192.0.2.28:5000` `imageRegistryNamespace: baas` `minioStorageClass: csi-rbd` `veleroNamespace: spp-velero`	`license: true` `isOCP: true` `clusterName: example-ocp-cluster` `networkPolicy:` `  clusterAPIServerips:` `  - 198.51.100.1` `  - 198.51.100.2` `  - 198.51.100.3` `  clusterAPIServerport: 6443` `  clusterCIDR: 198.51.100.0/24` `  isServerInstalledOnAnotherCluster: false` `SPPfqdn: my.ocp.plus.server.example` `SPPips: 198.51.100.12` `SPPport: 443` `productLoglevel: INFO` `imageRegistry: cp.icr.io/cp` `imageRegistryNamespace: sppc` `minioStorageClass: csi-rbd` `veleroNamespace: spp-velero`

Related tasks

“Installing Container Backup Support in an airgap environment” on page 214

“Installing Container Backup Support from IBM Helm Charts Repository and IBM Entitled Registry” on page 218

You can install Container Backup Support by using the IBM Helm Charts Repository that is linked to the IBM Entitled Registry.

“For OpenShift: Installing Container Backup Support by using the OpenShift web console” on page 220

You can install Container Backup Support by using the OpenShift web console to take advantage of the benefits that are afforded by the web console, such as monitoring the deployments from the web console.

Installing Container Backup Support in an airgap environment

You can install Container Backup Support in an airgap cluster by using the installation package from IBM Passport Advantage Online.

As used in this documentation, an airgap cluster is any Kubernetes or OpenShift cluster that does not have internet access, and therefore cannot pull container images from a Docker registry. The airgap

installation package includes the container images for Container Backup Support. During the installation, the container images are loaded by the **docker load** command and then tagged and pushed to the Docker registry that is specified in the `baas-options.sh` file. As a result, the container images can be pulled during the Helm installation.

Before you begin

For the system requirements for Container Backup Support, see "Container Backup Support requirements" on page 75.

Ensure that prerequisites are met and preliminary tasks are completed:

- Ensure that you are logged in to the target cluster as a user with `cluster-admin` privileges.
- Ensure that you complete the installation prerequisites. For instructions, see "Installation prerequisites for Container Backup Support" on page 197.
- Ensure that you set up the installation variables in the `baas-options.sh` and `baas-values.yaml` files. For instructions, see "Setting up the installation variables" on page 206.
- On an OpenShift cluster, the `amq-streams-cluster-operator` pod is installed from OpenShift OperatorHub. In an airgap environment, ensure that you set up OperatorHub to operate in a restricted network. For instructions, see Using Operator Lifecycle Manager on restricted networks.

About this task

You must first download the Container Backup Support installation package from the IBM Passport Advantage Online website. Then, extract the package and use the script that is provided in the installation package to deploy Container Backup Support on your Kubernetes or OpenShift cluster.

By using the installation variables that you set up in the `baas-options.sh` and `baas-values.yaml` files, the provided script, `baas-install-ppa.sh`, automatically runs prerequisite tasks and installs Container Backup Support on your cluster.

The following tasks are performed automatically:

- Checking for prerequisites.
- Logging in to your Docker registry.
- Removing any existing Container Backup Support resources and images.
- Loading and pushing the Container Backup Support Docker images to your Docker registry.
- Creating the Kubernetes product namespace or OpenShift project (`baas`) and secret.
- Creating an image pull secret called `baas-registry-secret` for the namespace (or project) `baas`. This secret is automatically copied to other namespaces (or projects) of the PVCs that you want to protect.
- Creating custom resource definitions for Container Backup Support and Kafka.

Procedure

1. Download the SPP_V*sppversion*_for_Containers.tar.gz package from the IBM Passport Advantage Online to your home folder (~), where *sppversion* specifies the version of IBM Spectrum Protect Plus that you are installing, such as `10.1.8`.

   For information about downloading files, see technote 6415111.
2. Extract the installation package by issuing the following command:

   ```
   tar -xvf SPP_V*sppversion*_for_Containers.tar.gz
   ```

   where *sppversion* specifies the version of IBM Spectrum Protect Plus that you are installing, such as `10.1.8`.
3. Change to the `installer` directory by issuing the following command:

```
cd installer
```

The following packages are in the directory:

File name	Description
baas-*version*.tar.gz Where *version* specifies the version of IBM Spectrum Protect Plus that you are installing, such as 10.1.8.	The package that includes the container images for Container Backup Support. Do not manually expand this package. This package was generated from a **docker save** command and can be processed only by a **docker load** command that is issued by the baas installation scripts.
ibm-spectrum-protect-plus-prod-*chart_version*.tgz Where *chart_version* specifies the version of the Helm chart. For example, specify 1.2.0 for IBM Spectrum Protect Plus V10.1.8, or 1.2.*x* for V10.1.8.x, where *x* specifies an interim fix.	The package that contains the Helm 3 chart for Container Backup Support.

4. Expand the ibm-spectrum-protect-plus-prod-*chart_version*.tgz file by issuing the following command:

```
tar -xvf ibm-spectrum-protect-plus-prod-chart_version.tgz
```

where *chart_version* specifies the version of the Helm chart. For example, specify 1.2.0 for IBM Spectrum Protect Plus V10.1.8, or 1.2.*x* for V10.1.8.x, where *x* specifies an interim fix.

5. Copy the baas-options.sh and baas-values.yaml files that you created to the Helm chart installation directory:

```
cd ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install
cp ~/install_vars_dir/baas-options.sh .
cp ~/install_vars_dir/baas-values.yaml .
chmod +x *.sh
```

where *install_vars_dir* is the directory where you saved your custom baas-options.sh and baas-values.yaml files.

Restriction: Ensure that you do not add any large files to the installer/ibm-spectrum-protect-plus-prod directory. The size of the contents in this directory, including files and subdirectories, must not exceed the limit set by Helm (3145728 bytes).

6. Issue the following command to deploy Container Backup Support:

```
./baas-install-ppa.sh
```

Results

You can verify that Container Backup Support is installed by issuing the following command:

```
helm3 list -n baas
```

The output is similar to the following example:

```
NAME                            NAMESPACE REVISION UPDATED                                  STATUS
CHART                                 APP VERSION
ibm-spectrum-protect-plus-prod baas      1        2020-10-28 13:15:08.154754539 -0700 MST deployed ibm-
spectrum-protect-plus-prod-1.2.0 10.1.8
```

All of the Container Backup Support pods will load and change to the Running state after a few minutes.

When all pods are running, the deployment is completed. To verify that all pods are in the Running state and no components are missing, issue the following command:

```
kubectl get pods -n baas -w
```

For Kubernetes, the output is similar to the following example:

```
NAME                                          READY   STATUS      RESTARTS   AGE
baas-controller-5f75fc6c9-tmg5l               1/1     Running     0          6h15m
baas-entity-operator-c99f4c49b-p9v9c          3/3     Running     1          6h15m
baas-kafka-0                                  2/2     Running     0          6h15m
baas-minio-0                                  1/1     Running     3          6h15m
baas-scheduler-dfdcd9467-88hb5                1/1     Running     0          6h15m
baas-spp-agent-db6b98f85-svdxz                1/1     Running     0          6h15m
baas-strimzi-cluster-operator-7b5c4f9597-88xfn 1/1    Running     0          6h15m
baas-transaction-manager-f654f7f48-7mdxt      3/3     Running     0          6h15m
baas-zookeeper-0                              1/1     Running     0          6h15m
baas-zookeeper-1                              1/1     Running     0          6h15m
baas-zookeeper-2                              1/1     Running     0          6h15m
```

For OpenShift, the output is similar to the following example:

```
NAME                                                  READY   STATUS      RESTARTS   AGE
amq-streams-cluster-operator-v1.5.3-5b795f4c69-gdsrx  1/1     Running     0          24m
baas-controller-5f75fc6c9-tmg5l                       1/1     Running     0          24m
baas-entity-operator-c99f4c49b-p9v9c                  3/3     Running     1          24m
baas-kafka-0                                          2/2     Running     0          24m
baas-minio-0                                          1/1     Running     3          24m
baas-scheduler-dfdcd9467-88hb5                        1/1     Running     0          24m
baas-spp-agent-db6b98f85-svdxz                        1/1     Running     0          24m
baas-transaction-manager-f654f7f48-7mdxt              3/3     Running     0          24m
baas-zookeeper-0                                      1/1     Running     0          24m
baas-zookeeper-1                                      1/1     Running     0          24mm
baas-zookeeper-2                                      1/1     Running     0          24m
```

What to do next

After the deployment is completed, the application host for the Container Backup Support container is automatically registered upon startup of the cluster host in Kubernetes or OpenShift. However, if no clusters are displayed in the **Manage Protection** > **Containers** > **Kubernetes** page or the **Manage Protection** > **Containers** > **OpenShift** page in the IBM Spectrum Protect Plus user interface, automatic registration was unsuccessful. You must then manually register the cluster. For instructions, see “Registering a Kubernetes cluster” on page 409 or “Registering an OpenShift cluster” on page 431.

Updating Container Backup Support: You can modify an existing configuration of Container Backup Support or upgrade the Helm chart. For instructions, see “Updating Container Backup Support” on page 227.

Related tasks

“Installing Container Backup Support from IBM Helm Charts Repository and IBM Entitled Registry” on page 218
You can install Container Backup Support by using the IBM Helm Charts Repository that is linked to the IBM Entitled Registry.

Related information

“Kafka cluster” on page 402

Installing Container Backup Support from IBM Helm Charts Repository and IBM Entitled Registry

You can install Container Backup Support by using the IBM Helm Charts Repository that is linked to the IBM Entitled Registry.

Before you begin

For the system requirements for Container Backup Support, see "Container Backup Support requirements" on page 75.

Ensure that prerequisites are met and preliminary tasks are completed:

- Ensure that you are logged in to the target cluster as a user with `cluster-admin` privileges.
- Ensure that internet access is available to pull containers at deployment time.
- Ensure that you complete the installation prerequisites. For instructions, see "Installation prerequisites for Container Backup Support" on page 197.
- Ensure that you set up the installation variables in the `baas-options.sh` and `baas-values.yaml` files. For instructions, see "Setting up the installation variables" on page 206.

About this task

You must first prepare your environment by adding the IBM Helm 3 repository to your local repository list and fetching the Container Backup Support Helm package from the IBM Helm Charts Repository. Then, extract the package and use the script that is provided in the Container Backup Support Helm package to deploy Container Backup Support on your Kubernetes or OpenShift cluster.

By using the installation variables that you set up in the `baas-options.sh` and `baas-values.yaml` files, the provided script, `baas-install-entitled-registry.sh`, automatically runs prerequisite tasks and installs Container Backup Support on your cluster. The following tasks are performed automatically:

- Checking for prerequisites.
- Logging in to the IBM Entitled Registry.
- Removing any existing Container Backup Support resources and images.
- Creating the product Kubernetes namespace or OpenShift project (`baas`) and secret.
- Creating an image pull secret called `baas-registry-secret` for the namespace (or project) `baas`. This secret is automatically copied to other namespaces (or projects) of the PVCs that you want to protect.
- Creating custom resource definitions for Container Backup Support and Kafka.
- Pulling the Container Backup Support containers from the IBM Entitled Registry at deployment time.

Procedure

1. Complete this one-time preparation to add the IBM Helm Charts repository to the local repository list. Issue the following commands:

```
helm3 repo add ibm-helm https://raw.githubusercontent.com/IBM/charts/master/repo/ibm-helm
helm3 repo list
helm3 repo update
helm3 search repo spectrum
```

2. Fetch the Container Backup Support Helm package from the IBM Helm Charts Repository:

```
mkdir installer
cd installer
helm3 fetch ibm-helm/ibm-spectrum-protect-plus-prod --version "chart_version"
```

where *chart_version* specifies the version of the Helm chart. For example, specify 1.2.0 for IBM Spectrum Protect Plus V10.1.8, or 1.2.*x* for V10.1.8.*x*, where *x* specifies an interim fix.

3. Extract the Helm package:

```
tar -xvf ibm-spectrum-protect-plus-prod-chart_version.tgz
```

   where *chart_version* specifies the version of the Helm chart. For example, specify 1.2.0 for IBM Spectrum Protect Plus V10.1.8, or 1.2.*x* for V10.1.8.*x*, where *x* specifies an interim fix.

4. Copy the `baas-options.sh` and `baas-values.yaml` files that you created to the Helm chart installation directory:

```
cd ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install
cp ~/install_vars_dir/baas-options.sh .
cp ~/install_vars_dir/baas-values.yaml .
chmod +x *.sh
```

   where *install_vars_dir* is the directory where you saved your custom `baas-options.sh` and `baas-values.yaml` files.

   Restriction: Ensure that you do not add any large files to the `installer/ibm-spectrum-protect-plus-prod` directory. The size of the contents in this directory, including files and subdirectories, must not exceed the limit set by Helm (3145728 bytes).

5. Issue the following command to deploy Container Backup Support:

```
./baas-install-entitled-registry.sh
```

Results

You can verify that Container Backup Support is installed by issuing the following command:

```
helm3 list -n baas
```

The output is similar to the following example:

```
NAME                            NAMESPACE REVISION UPDATED                                  STATUS
CHART                                 APP VERSION
ibm-spectrum-protect-plus-prod baas      1        2020-10-28 13:15:08.154754539 -0700 MST deployed ibm-
spectrum-protect-plus-prod-1.2.0 10.1.8
```

All of the Container Backup Support pods will load and change to the `Running` state after a few minutes.

When all pods are running, the deployment is completed. To verify that all pods are in the `Running` state and no components are missing, issue the following command:

```
kubectl get pods -n baas -w
```

For Kubernetes, the output is similar to the following example:

```
NAME                                          READY   STATUS      RESTARTS   AGE
baas-controller-5f75fc6c9-tmg5l               1/1     Running     0          6h15m
baas-entity-operator-c99f4c49b-p9v9c          3/3     Running     1          6h15m
baas-kafka-0                                  2/2     Running     0          6h15m
baas-minio-0                                  1/1     Running     3          6h15m
baas-scheduler-dfdcd9467-88hb5                1/1     Running     0          6h15m
baas-spp-agent-db6b98f85-svdxz                1/1     Running     0          6h15m
baas-strimzi-cluster-operator-7b5c4f9597-88xfn 1/1     Running     0          6h15m
baas-transaction-manager-f654f7f48-7mdxt      3/3     Running     0          6h15m
baas-zookeeper-0                              1/1     Running     0          6h15m
baas-zookeeper-1                              1/1     Running     0          6h15m
baas-zookeeper-2                              1/1     Running     0          6h15m
```

For OpenShift, the output is similar to the following example:

```
NAME                                          READY   STATUS      RESTARTS   AGE
amq-streams-cluster-operator-v1.5.3-5b795f4c69-gdsrx   1/1     Running     0          24m
baas-controller-5f75fc6c9-tmg5l               1/1     Running     0          24m
baas-entity-operator-c99f4c49b-p9v9c          3/3     Running     1          24m
```

```
baas-kafka-0                                      2/2     Running     0          24m
baas-minio-0                                      1/1     Running     3          24m
baas-scheduler-dfdcd9467-88hb5                    1/1     Running     0          24m
baas-spp-agent-db6b98f85-svdxz                    1/1     Running     0          24m
baas-transaction-manager-f654f7f48-7mdxt          3/3     Running     0          24m
baas-zookeeper-0                                  1/1     Running     0          24m
baas-zookeeper-1                                  1/1     Running     0          24mm
baas-zookeeper-2                                  1/1     Running     0          24m
```

What to do next

After the deployment is completed, the application host for the Container Backup Support container is automatically registered upon startup of the cluster host in Kubernetes or OpenShift. However, if no clusters are displayed in the **Manage Protection** > **Containers** > **Kubernetes** page or the **Manage Protection** > **Containers** > **OpenShift** page in the IBM Spectrum Protect Plus user interface, automatic registration was unsuccessful. You must then manually register the cluster. For instructions, see "Registering a Kubernetes cluster" on page 409 or "Registering an OpenShift cluster" on page 431.

Updating Container Backup Support: You can modify an existing configuration of Container Backup Support or upgrade the Helm chart. For instructions, see "Updating Container Backup Support" on page 227.

Related tasks

"Installing Container Backup Support in an airgap environment" on page 214

For OpenShift: Installing Container Backup Support by using the OpenShift web console

You can install Container Backup Support by using the OpenShift web console to take advantage of the benefits that are afforded by the web console, such as monitoring the deployments from the web console.

Before you begin

For the system requirements for Container Backup Support, see "Container Backup Support requirements" on page 75.

Restriction: You can install Container Backup Support from the OpenShift web console only if you are running OpenShift Container Platform 4.6 or later. If your are using OpenShift Container Platform 4.5, you must install Container Backup Support by using the command line. For instructions, see "For Kubernetes or OpenShift: Installing Container Backup Support by using the command line" on page 205.

The instructions that are provided apply to pulling images from the IBM Helm Charts Repository that is linked to IBM Entitled Registry. If you are operating in an airgap environment, use the command line to install Container Backup Support. For instructions, see "Installing Container Backup Support in an airgap environment" on page 214.

Ensure that prerequisites are met and preliminary tasks are completed:

- Ensure that you are logged in to the target cluster as a user with `cluster-admin` privileges.
- Ensure that internet access is available to pull containers at deployment time.
- Ensure that you complete the installation prerequisites. For instructions, see "Installation prerequisites for Container Backup Support" on page 197.
- Ensure that you set up the installation variables in the `baas-options.sh` and `baas-values.yaml` files. For instructions, see "Setting up the installation variables" on page 206.

Procedure

1. "Adding the IBM Entitled Registry to your Helm repository" on page 221
2. "Creating a project for Container Backup Support" on page 222
3. "Creating image pull secrets" on page 222

4. “Creating the credentials secret” on page 223
5. “Creating custom resource definitions” on page 224
6. “Installing Container Backup Support from the OpenShift web console” on page 224

Adding the IBM Entitled Registry to your Helm repository

To prepare the OpenShift web console to pull Container Backup Support images from the IBM Helm Charts Repository that is linked to the IBM Entitled Registry, you must add the IBM Entitled Registry to your Helm repository.

About this task

Adding the IBM Helm Charts Repository that is linked to the IBM Entitled Registry is a one-time task.

Procedure

1. Define the location of the IBM Helm Charts Repository by saving the following text to a YAML file called `entitled-registry-repo.yaml`:

```
#--------------------------------------
# Filename: entitled-registry-repo.yaml
#--------------------------------------
apiVersion: helm.openshift.io/v1beta1
kind: HelmChartRepository
metadata:
  name: entitledregistry
spec:
  name: entitledregistry
  connectionConfig:
    url: https://raw.githubusercontent.com/IBM/charts/master/repo/ibm-helm
```

2. Apply the YAML file by issuing the following command:

```
oc apply -f entitled-registry-repo.yaml
```

   The following message is displayed:

```
helmchartrepository.helm.openshift.io/entitledregistry configured
```

   The Helm charts in the IBM Entitled Registry, including the Helm chart for IBM Spectrum Protect Plus Container Backup Support, are added to the **Developer Catalog** page of the OpenShift web console.

What to do next

To take advantage of the scripts that are provided in the installation package, you must fetch and extract the installation package by issuing the following commands from your home folder (~):

```
helm3 repo add ibm-helm https://raw.githubusercontent.com/IBM/charts/master/repo/ibm-helm
helm3 repo list
helm3 repo update
helm3 search repo spectrum
mkdir installer
cd installer
helm3 fetch ibm-helm/ibm-spectrum-protect-plus-prod --version "chart_version"
tar -xvf ibm-spectrum-protect-plus-prod-chart_version.tgz
```

where *chart_version* specifies the version of the Helm chart. For example, specify 1.2.0 for IBM Spectrum Protect Plus V10.1.8, or 1.2.*x* for V10.1.8.*x*, where *x* specifies an interim fix.

Restriction: Ensure that you do not add any large files to the `installer/ibm-spectrum-protect-plus-prod` directory. The size of the contents in this directory, including files and subdirectories, must not exceed the limit set by Helm (3145728 bytes).

Creating a project for Container Backup Support

You must create a project (namespace) for Container Backup Support called `baas` in the OpenShift web console.

Before you begin

Remove previous installations of Container Backup Support by taking the following actions:

1. Set up the installation variables in the `baas-options.sh` file as described in "Setting up the installation variables" on page 206.
2. Copy the `baas-options.sh` file to the `~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install` directory.
3. Uninstall the previous installation of Container Backup Support. For instructions, see "Uninstalling Container Backup Support" on page 231.

About this task

Creating the `baas` project is a one-time task.

Procedure

1. Log on to the OpenShift web console as the cluster administrator.
2. Select the **Developer** view and click **Project**.
3. On the **Project Details** page, click the **Create a project** link.
4. In the **Create Project** window, enter `baas` in the **Name** field.
5. Optional: Enter a display name ad a description for the project.
6. Click **Create**.

   The `baas` project is created.

Creating image pull secrets

Create an image pull secret in the `baas` project that you created for Container Backup Support.

About this task

The image pull secret provides the credentials that are required by OpenShift to pull images from the IBM Helm Charts Repository that is linked to the IBM Entitled Registry.

Procedure

1. Log on to the OpenShift web console as the cluster administrator.
2. Select the **Developer** view and click **Secrets**.
3. Ensure that the **baas** project is selected.

4. On the **Secrets** page, click **Create** > **Image Pull Secret**.
5. On the **Create Image Pull Secret** page, complete the following fields:

   Secret Name
   Enter `baas-registry-secret`.

   Authentication Type
   Select **Image Registry Credentials**.

Registry Server Address
Specify the Docker registry where images are pulled. For example, enter the value for the DOCKER_REGISTRY_ADDRESS variable that you set up in the `baas-options.sh` file. For the IBM Entitled Registry, specify `cp.icr.io/cp`.

Username
Enter the Docker registry username that is associated with the DOCKER_REGISTRY_USERNAME value in the `baas-options.sh` file. For the IBM Entitled Registry, specify `cp`.

Password
Enter the entitlement key for the IBM Entitled Registry. For instructions on obtaining the entitlement key, see Obtain an entitlement key.

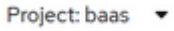
6. Click **Create**.

Creating the credentials secret

To store credentials for the `baas` project, you must create a secret named `baas-secret`.

Procedure

1. Log on to the OpenShift web console as the cluster administrator.
2. Select the **Developer** view and click **Secrets**.
3. Ensure that the **baas** project is selected.



4. Create the credentials secret by manually entering credentials or by using a script.

- Creating the credentials secret by invoking a script:
  a. On the **Secrets** page, click **Create** > **From YAML**. The **Create Secret** page is displayed.
  b. Go to the command line and issue the following commands to run the `baas-prereqs-create-baas-secret-yaml.sh` script:

```
cd ~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install
./baas-prereqs-create-baas-secret-yaml.sh
```

  The script extracts the values for the account credentials from the `baas-options.sh` file, converts the credentials into base64 encoding, and puts the credentials into a YAML format.

  The output is in the following format:

```
apiVersion: v1
kind: Secret
metadata:
  name: baas-secret
  namespace: baas
type: Opaque
data:
  baasadmin: xxxxxxxxxxxx
  baaspassword: xxxxxxxxxxx=
  datamoveruser: xxxxxxxxxxx=
  datamoverpassword: xxxxxxxxxxx=
  miniouser: xxxxxxxxxxx=
  miniopassword: xxxxxxxxxxxxxxx
```

  c. Copy the output and paste it to the YAML editor on the **Create Secret** page in the OpenShift web console.
  d. Click **Create**.
- Creating the credentials secret by manually entering key-value pairs for the credentials:
  a. On the **Secrets** page, click **Create** > **Key/Value Secret**.
  b. In the **Secret Name** field, enter `baas-secret`.
  c. Create the following credentials by specifying the key-value pairs. Click **Add Key/Value** each to add a key-value pair and click **Create** to create the `baas-secret` secret.

The following table lists the key-value pairs that you must provide:

Table 75. Key-value pairs for the baas-secret secret	
Key	**Value**
baasadmin	Specify the username for the IBM Spectrum Protect Plus containers administrator. Enter the value that is used for the SPP_ADMIN_USERNAME variable in the baas-options.sh file.
baaspassword	Specify the password for the IBM Spectrum Protect Plus containers administrator. Enter the value that is used for the SPP_ADMIN_PASSWORD variable in the baas-options.sh file.
datamoveruser	Specify the data mover username. Use the value that you specified for the DATAMOVER_USERNAME variable in the baas-options.sh file.
datamoverpassword	Specify the password for the data mover. Use the value that you specified for the DATAMOVER_PASSWORD variable in the baas-options.sh file.
miniouser	Specify the username for the minIO server. Use the value that you specified for the MINIO_USERNAME variable in the baas-options.sh file.
miniopassword	Specify the password for the minIO user. Use the value that you specified for the MINIO_PASSWORD variable in the baas-options.sh file.

Creating custom resource definitions

Before you can install Container Backup Support in the OpenShift web console, you must create custom resource definitions (CRDs) for Container Backup Support and Kafka.

About this task

The CRDs for Container Backup Support and Kafka are provided in the installation package, but you must apply them at the command line. The Kafka cluster, which is included in the installation, handles messaging operations in Container Backup Support.

Use the provided script, baas-crd-apply.sh, to create the CRDs.

Procedure

1. Change to the installation directory by issuing the following command:

```
cd ~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install
```

2. Apply the CRDs by issuing the following command:

```
./baas-crd-apply.sh
```

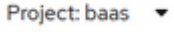
Installing Container Backup Support from the OpenShift web console

After you set up your environment to pull images from the IBM Entitled Registry, created the baas project, and created image pull and credentials secrets, you can install Container Backup Support from the OpenShift web console.

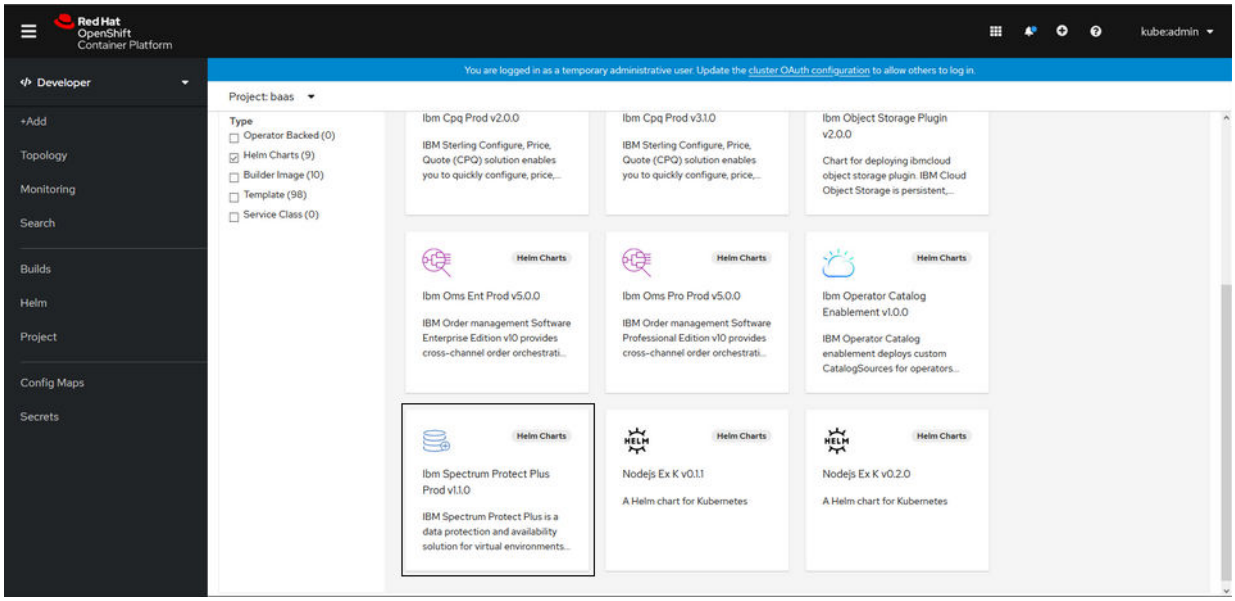
Procedure

1. Log on to the OpenShift web console as the cluster administrator.

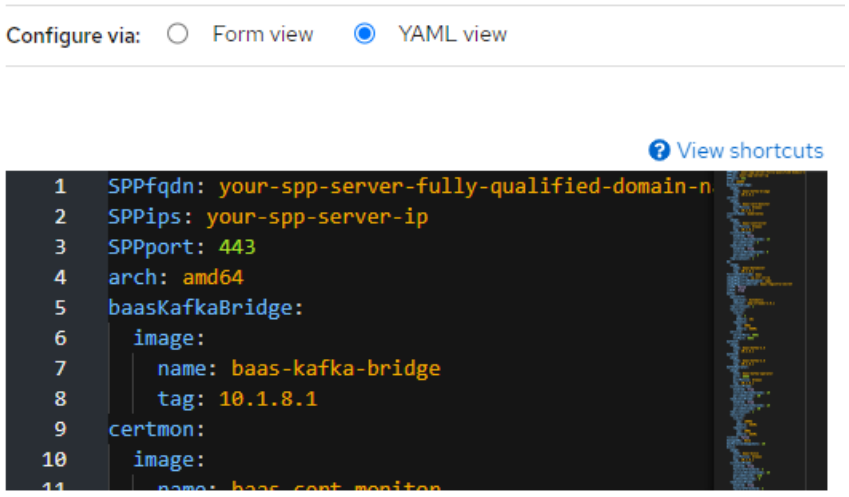
2. Select the **Developer** perspective and click **Add** > **Helm Chart**.
3. Ensure that the **baas** project is selected in the **Project** list.

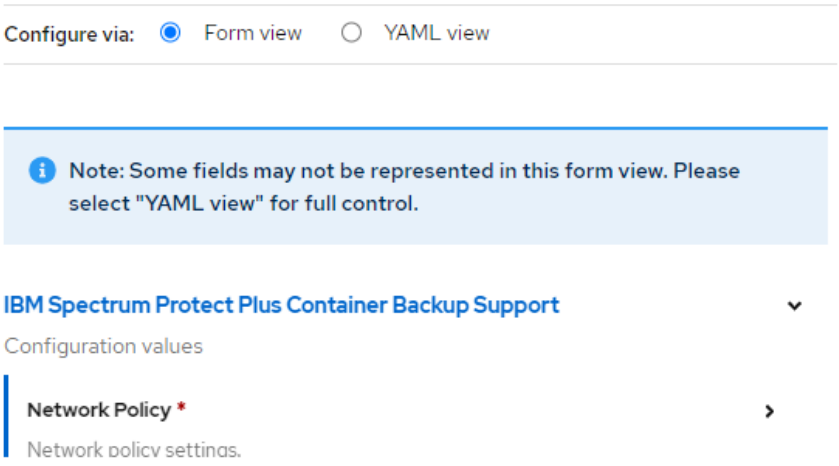
   

4. On the **Developer Catalog** page, click the **IBM Spectrum Protect Plus Prod v1.2.1** tile.



5. On the **IBM Spectrum Protect Plus Prod v1.2.1** description page, click **Install Helm Chart**.
6. On the **Install Helm Chart** page, provide the configuration values to use for the installation of Container Backup Support. You can use either the **YAML view** or **Form view** method:

Method	Steps
By using the YAML view	For best results, use the **YAML view** to paste the contents of your custom `baas-values.yaml` file. This method can be less error-prone than the **Form view** method. a. Click **YAML view**. b. Copy the contents of the `baas-values.yaml` file that you customized and paste the contents to the YAML editor. 

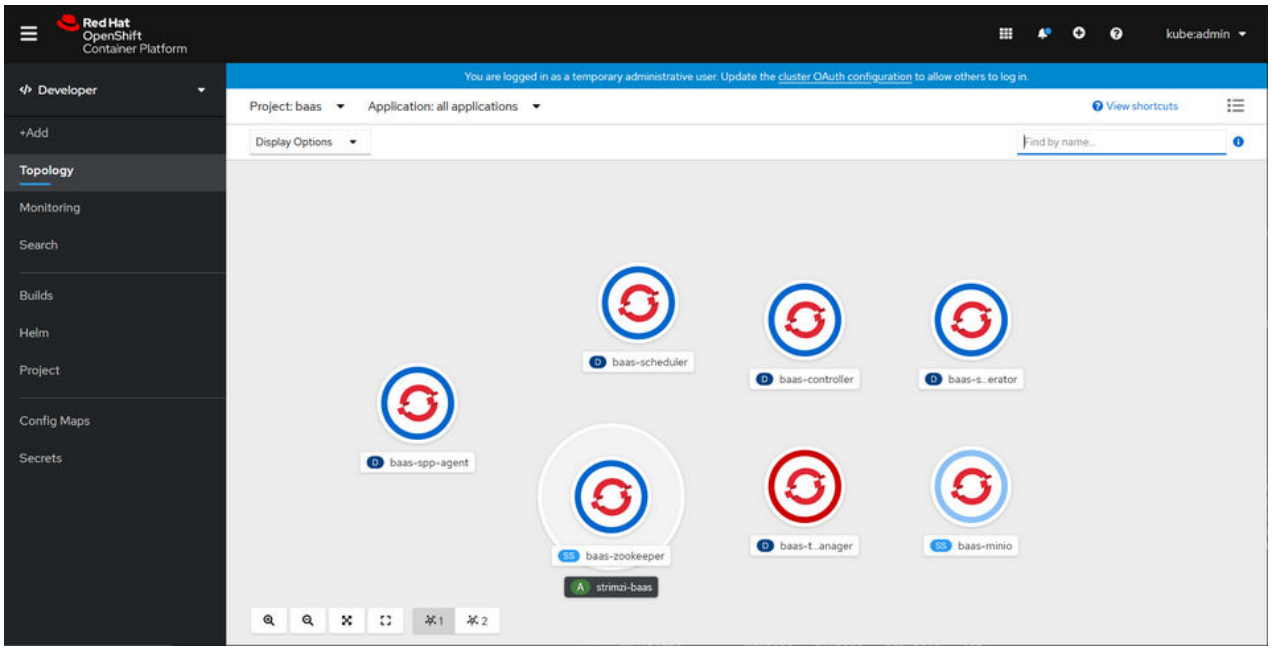
Method	Steps
	The following text block shows sample content from a `baas-values.yaml` file: `license: true` `isOCP: true` `clusterName: example-cluster` `networkPolicy:` `  clusterAPIServerips:` `  - 198.51.100.1` `  - 198.51.100.2` `  - 198.51.100.3` `  clusterAPIServerport: 6443` `  clusterCIDR: 203.0.113.0/24` `SPPips: 192.0.2.83` `SPPport: 443` `productLoglevel: INFO` `imageRegistry: cp.icr.io/cp` `imageRegistryNamespace: sppc` `minioStorageClass: example-csi-rbd` `veleroNamespace: spp-velero` **Tip:** You can replace all the contents in the YAML editor with these statements. Any values that are not specified here will come from the `values.yaml` file. You can also add any other variables from the `values.yaml` file to the YAML editor. For information about the parameters that are used, see Table 78 on page 234.
By using the Form view	Manually enter the configuration values that you set up in the `bass-values.yaml` file by completing the following steps: a. Click **Form view**. b. To view the configuration values, expand the **IBM Spectrum Protect Plus Container Backup Support** twistie. The configuration values that are used for Container Backup Support are listed in hierarchical order. The listed values are a subset of all the values in the `values.yam` file.  c. Review the configuration parameters and replace the listed values with the values that you customized in the `baas-values.yaml` file. Descriptions of the parameters are provided. To view a nested section, expand a twistie. To search for a parameter, expand all the twisties, press Ctrl+F and enter the search term.

Method	Steps
	For more information about the configuration values in the `baas-values.yaml` file, see Table 78 on page 234

7. Review your updates and click **Install**.

Results

You can review the status of the installation and manage Container Backup Support components by using the OpenShift web console. The following figure shows the topology view of the Container Backup Support containers as they are being installed.



What to do next

You can view the details of the Container Backup Support components on the **baas** project page.

After the deployment is completed, the application host for the Container Backup Support container is automatically registered upon startup of the cluster host in OpenShift. However, if no clusters are displayed in the **OpenShift** page in the IBM Spectrum Protect Plus user interface, automatic registration was unsuccessful. You must then manually register the cluster. For instructions, see "Registering an OpenShift cluster" on page 431.

Updating Container Backup Support: You can modify an existing configuration of Container Backup Support or upgrade the Helm chart. For instructions, see "Updating Container Backup Support" on page 227.

Updating Container Backup Support

You can update an existing configuration of Container Backup Support or upgrade the Helm chart for Container Backup Support.

Before you begin

Ensure that you set up the installation variables in the `baas-values.yaml` file. For instructions, see "Setting up installation variables in the baas-values.yaml file" on page 209.

About this task

For Kubernetes or OpenShift: To update Container Backup Support at the command line, modify the parameters in the `baas-values.yaml` file as required for your environment. Then, run the `baas-upgrade.sh` script to begin the update.

For OpenShift: To update Container Backup Support from the OpenShift web console, modify the parameters in the `baas-values.yaml` file as required for your environment. Then, log in to the OpenShift web console to complete the update.

Restriction: The `baas-upgrade.sh` script cannot upgrade the Helm chart from a Helm 2 chart, nor does it upgrade the product from one version to another. For example, the Container Backup Support Version 10.1.7.x software cannot be upgraded to V10.1.8 by using the Helm 3 chart. To upgrade from one version to another, you must uninstall the installed version, and then install V10.1.8 by using the configuration values from V10.1.7.

The following procedure contains examples for how to change the logging level for Container Backup Support at the command line and from the OpenShift web console.

Procedure

Use one of the following methods to change the logging level for Container Backup Support from INFO to DEBUG.

- For Kubernetes or OpenShift: Update Container Backup Support by using the command line. This method works in both airgap and online environments.
  a) Edit the `baas-values.yaml` file that you customized and modify the **`productLoglevel`** statement as follows:

     ```
     productLoglevel: DEBUG
     ```

     Save the file.
  b) Change to the installation directory (`~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install`).
  c) Copy your custom `baas-values.yaml` file to the installation directory:

     ```
     cp ~/install_vars_dir/baas-values.yaml .
     ```

     where *install_vars_dir* is the directory where you saved your custom `baas-values.yaml` files.
  d) Issue the following command to update Container Backup Support:

     ```
     ./baas-upgrade.sh
     ```

  e) After the update is completed, verify the values that you specified by issuing the following command:

     ```
     helm3 get values -n baas
     ```

     You can also show the revision history by issuing the following command:

     ```
     helm3 history ibm-spectrum-protect-plus-prod -n baas
     ```

     The output is similar to the following example:

```
REVISION UPDATED                   STATUS     CHART                                 APP VERSION DESCRIPTION
1        Mon Nov  9 09:05:11 2020 superseded ibm-spectrum-protect-plus-prod-1.1.0 10.1.8      Install
complete
2        Tue Nov 10 10:30:12 2020 deployed   ibm-spectrum-protect-plus-prod-1.1.0 10.1.8      Upgrade
complete
```

- For OpenShift: Update Container Backup Support by using the OpenShift web console. This method works only in an online environment.

a) Edit the `baas-values.yaml` that you customized and modify the **`productLogLevel`** statement as follows:

```
productLoglevel: DEBUG
```

Save the file.

b) Log on to the OpenShift web console as the cluster administrator.

c) From the **Developer** perspective in the OpenShift web console, click **Helm**.

d) On the **Helm Releases** page, ensure that the **baas** project is selected. Then, click the **ibm-spectrum-protect-plus-prod** release.

e) On the **ibm-spectrum-protect-plus-prod** Helm release details page, click **Details** > **Action** > **Upgrade**.

f) On the **Upgrade Helm Release** page, paste the contents of the updated `baas-values.yaml` file to the YAML editor.

g) Click **Upgrade**.

h) Verify that the update has completed by clicking **Revision History** on the **ibm-spectrum-protect-plus-prod** release details page.

i) After the update is completed, go to the command line. At the installation directory (`~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install`), run the following command to delete any obsolete data mover containers:

```
./baas-delete-dm.sh
```

Updating your credentials after installation

After Container Backup Support is installed in your environment, if you update your IBM Spectrum Protect Plus credentials due to a password rotation, account change, or any other reason, you must update the secret that stores credentials for the `baas` namespace (or project).

About this task

You must re-create the secret that is named `baas-secret` if you change one or more of the following credentials:

- IBM Spectrum Protect Plus containers administrator username or password
- Data mover username or password
- minIO username or password

You can re-create the secret by using the command line in the Kubernetes or OpenShift environment, or by using the OpenShift web console.

Procedure

- To re-create the `baas-secret` secret by using the Kubernetes or OpenShift command line, complete the following steps:

  a) Update one or more of the following variables in the `baas-options.sh` file:

  ```
  export SPP_ADMIN_USERNAME='your_protectplus_containers_admin_username'
  export SPP_ADMIN_PASSWORD='your_protectplus_containers_admin_password'
  export DATAMOVER_USERNAME='create_a_datamover_username'
  export DATAMOVER_PASSWORD='create_a_datamover_password'
  export MINIO_USERNAME='create_a_minio_username'
  export MINIO_PASSWORD='create_a_minio_password'
  ```

  For a description of the variables, see Table 77 on page 232.

b) At the installation directory (`~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install`), run the following script to re-create the secret:

```
./baas-prereqs-create-baas-secret.sh
```

c) Delete all of the `transaction-manager` pods and the `spp-agent` pods by issuing the following commands:

```
kubectl delete pods -l app.kubernetes.io/component=transaction-manager -n baas
kubectl delete pods -l app.kubernetes.io/component=spp-agent -n baas
```

After the pods are deleted, the pods will restart and capture the changes in the `baas-secret` secret.

- To re-create the `baas-secret` secret by using the OpenShift web console, complete the following steps:

a) Log in to the OpenShift web console as the cluster administrator.

b) Select the **Developer** perspective and click **Secrets**.

c) Click **Project** > **baas**.

d) Update one or more of the following key-values pairs for the `baas-secret` secret.

Table 76. Key-value pairs for the `baas-secret` secret

Key	Value
`baasadmin`	Specify the username for the IBM Spectrum Protect Plus containers administrator. Enter the value that is used for the SPP_ADMIN_USERNAME variable in the `baas-options.sh` file.
`baaspassword`	Specify the password for the IBM Spectrum Protect Plus containers administrator. Enter the value that is used for the SPP_ADMIN_PASSWORD variable in the `baas-options.sh` file.
`datamoveruser`	Specify the data mover username. Use the value that you specified for the DATAMOVER_USERNAME variable in the `baas-options.sh` file.
`datamoverpassword`	Specify the password for the data mover. Use the value that you specified for the DATAMOVER_PASSWORD variable in the `baas-options.sh` file.
`miniouser`	Specify the username for the minIO server. Use the value that you specified for the MINIO_USERNAME variable in the `baas-options.sh` file.
`miniopassword`	Specify the password for the minIO user. Use the value that you specified for the MINIO_PASSWORD variable in the `baas-options.sh` file.

e) From the **Pods** view of the **baas** project, manually delete all of the `baas-transaction-manager` pods and the `baas-spp-agent` pods.

After the pods are deleted from the web console, the pods will restart and capture the changes in the `baas-secret` secret.

Uninstalling Container Backup Support

You can completely uninstall Container Backup Support so that all components, configuration settings, and backups are removed from the Kubernetes or OpenShift environment.

Before you begin

Take the following actions:

- Stop all scheduled backup operations. For instructions, see Discontinuing SLA backups for a PVC or Modifying parameters in a YAML file.
- Wait for all running backup and restore jobs to finish.

Procedure

To completely uninstall Container Backup Support from the cluster that you are logged in to, complete the following steps on the command line:

1. Destroy all snapshot and copy backups with a **destroy** request. For instructions, see "Deleting container backups" on page 489.
2. Delete any persistent volume claims (PVCs) that were used for copy backups.

   Tip: To determine which PVCs were used for copy backups, look for the names of the PVCs that were backed up.
3. Uninstall Container Backup Support by issuing the following commands:

   ```
   cd ~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install
   helm3 uninstall ibm-spectrum-protect-plus-prod --namespace baas
   kubectl delete namespace baas
   ./baas-docker-remove-images.sh
   ```
4. Remove the custom resource definitions (CRDs) for Container Backup Support and Kafka by taking the following actions:
   a) Determine whether Kafka is used by other applications outside of Container Backup Support.
   b) If Kafka is used exclusively by Container Backup Support, issue the following commands:

      ```
      cd ~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install
      kubectl delete -f ../crds
      ```
   c) If Kafka is also used by other applications in your cluster, delete only the Container Backup Support CRDs by issuing the following commands:

      ```
      cd ~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install
      kubectl delete -f ../crds/baas.io_baasreqs_crd.yaml
      ```
5. Optional: To verify the progress of the uninstallation, enter the following command:

   ```
   kubectl get pods -n baas
   ```
6. Optional: For OpenShift, if the `amq-streams-cluster-operator` pod is still running after the uninstallation is completed, you must manually uninstall it. To uninstall the `amq-streams-cluster-operator` pod, you must delete all ClusterServiceVersion (CSV) objects in the installation namespace. For example, issue the following command:

   ```
   oc delete csv --namespace baas --all
   ```
7. Unregister the Kubernetes or OpenShift cluster by using the IBM Spectrum Protect Plus user interface:
   a) In the navigation panel, take one of the following actions:
      - For Kubernetes: Click **Manage Protection** > **Containers** > **Kubernetes**.
      - For OpenShift: Click **Manage Protection** > **Containers** > **OpenShift**.
   b) Click **Manage clusters**.

c) In the list of host addresses, click the deletion icon next the cluster that you want to unregister.

d) In the **Confirm** window, enter the displayed confirmation code, and click **Unregister**.

8. Remove the account identity that is used to register the Kubernetes or OpenShift cluster:

a) In the navigation panel, click **Accounts** > **Identity**.

b) Click the deletion icon that is associated with the cluster.

c) Click **Yes** to delete the identity.

9. Optional: Review the installation and configuration information and revert any prerequisite steps.

What to do next

If Container Backup Support was not uninstalled cleanly, see "Container Backup Support did not uninstall cleanly" in “Troubleshooting quick reference” on page 667.

Configuration parameters for Container Backup Support

The configuration parameters of the Container Backup Support Helm chart are provided.

The values for the parameters are specified in the following files:

`baas-options.sh`
Contains the variables that are used to configure the prerequisites for Container Backup Support. This file is used to replace the sample `baas-options.sh` file that is provided in the installation package.

`baas-values.yaml`
Contains the values that are used to install Container Backup Support or to update an existing configuration. This file is used to replace the sample `baas-values.yaml` file that is provided in the installation package.

For more information, see “Setting up the installation variables” on page 206.

The following table contains the descriptions for the environment variables in the `baas-options.sh` file. You must enclose the values with single quotation marks (`' '`).

Table 77. Installation variables in the `baas-options.sh` file

Environment variable	Description
DOCKER_REGISTRY_ADDRESS	The address of the Docker registry in your environment where container images are loaded. If you are pulling images from the IBM Entitled Registry, you must specify '`cp.icr.io/cp`'. The value for DOCKER_REGISTRY_ADDRESS must match the value for the **`imageRegistry`** parameter in the `baas-values.yaml` file.
DOCKER_REGISTRY_USERNAME	The user account for the Docker registry where container images are loaded. If you are pulling images from the IBM Entitled Registry, you must specify '`cp`'.

Table 77. Installation variables in the baas-options.sh file (continued)

Environment variable	Description
DOCKER_REGISTRY_PASSWORD	The user password for the Docker registry where the container images are loaded. To pull images from the IBM Entitled Registry, specify the entitlement key that you obtained from the IBM Container software library. You can avoid putting the password in the file by specifying an environment variable for any of the passwords. For example, ${DOCKERUSER_PW} or ${IBMCLOUD_API_KEY}.
DOCKER_REGISTRY_NAMESPACE	The namespace of the Docker registry where the container images are loaded. The namespace does not have to be created ahead of time. To pull images from the IBM Entitled Registry, you must specify 'sppc'. The value for DOCKER_REGISTRY_NAMESPACE must match the value for the **imageRegistryNamespace** parameter in the baas-values.yaml file.
SPP_ADMIN_USERNAME	The user ID of the IBM Spectrum Protect Plus containers administrator. The containers administrator is an IBM Spectrum Protect Plus administrator with the **Containers Admin** role.
SPP_ADMIN_PASSWORD	The IBM Spectrum Protect Plus password for the containers administrator. You can optionally specify an environment variable for the password. For example, ${PROTECTPLUS_ADMIN_PW}.
DATAMOVER_USERNAME	The user ID to create for use with the data mover. The value does not have to exist already. It is created for the installation. The data mover username must adhere to the rules for usernames and passwords for Red Hat Enterprise Linux (RHEL) 7 operating system. The rules are the same as the ones for creating a new user on RHEL 7. For example, the password and the username must not be the same.
DATAMOVER_PASSWORD	The user password to create for use with the data mover. The value does not have to exist already. It is created for the installation. The data mover password must adhere to the rules for usernames and passwords for RHEL 7. The rules are the same as the ones for creating a new user on RHEL 7. For example: • The password must be at least 8 characters in length, and must contain letters and numbers. • No dictionary words are allowed in the password. • The password cannot be the same as the username.
MINIO_USERNAME	The username to create for the MinIO user. MinIO object storage is used to store backups of cluster and namespace resources. The value does not have to exist already. It is created for the installation.

Table 77. Installation variables in the baas-options.sh file (continued)	
Environment variable	**Description**
MINIO_PASSWORD	The password to create for the MinIO user. The value does not have to exist already. It is created for the installation.
BAAS_VERSION	The version of IBM Spectrum Protect Plus that you are installing, for example, 10.1.8 or 10.1.8.*x*, where *x* specifies an interim fix (iFix).

The following table contains the descriptions and default values for the configuration parameters in the `baas-values.yaml` file:

Table 78. Configuration parameters in the baas-values.yaml file		
Parameter	**Description**	**Default value**
license	The product license for Container Backup Support. The English license file is located in the `LICENSES/LICENSE-en` directory, which is included in the installation package. Versions of the license in English and other languages are available in the "IBM Spectrum Protect Plus Capacity - Version 10.1.8" license agreements at License Information documents. Set the value to `true` to indicate that you have reviewed and agree to the license agreement.	false
isOCP	The type of cluster on which you are installing Container Backup Support. If you are installing the product on an OpenShift cluster, set the value to `true`. If you are installing the product on a Kubernetes cluster, set the value to `false`.	false
clusterName	The unique cluster name that is used to register the application host to the IBM Spectrum Protect Plus server. The cluster name can be any name of your choice, but it must be unique from the IBM Spectrum Protect Plus server.	None
clusterAPIServerips	The IP address for the cluster API server. To obtain the cluster API server address, issue the following command: For Kubernetes: `kubectl get endpoints -n default -o yaml kubernetes` For OpenShift: `oc get endpoints -n default -o yaml kubernetes` Use all of the provided addresses listed under the **addresses** field in the output, or add or remove IP addresses as needed. Specify multiple addresses as follows: `networkPolicy:` `  clusterAPIServerips:` `    - x.x.x.x` `    - y.y.y.y` `    - z.z.z.z`	*x.x.x.x*

Table 78. Configuration parameters in the baas-values.yaml file (continued)

Parameter	Description	Default value
clusterAPIServerport	The port address for the cluster API server. To obtain the cluster API server port, issue the following command: For Kubernetes: `kubectl get endpoints -n default -o yaml kubernetes` For OpenShift: `oc get endpoints -n default -o yaml kubernetes` Use the port number listed in the **port** field in the output.	6443
clusterCIDR	The Classless Inter-Domain Routing (CIDR) value for the cluster. To obtain the CIDR, issue the following command: For Kubernetes: `kubectl cluster-info dump | grep -m 1 cluster-cidr` For OpenShift: `oc get network -o yaml | grep -A1 clusterNetwork:` Use the displayed IP address as the cluster CIDR address. **Tip for Kubernetes:** If the command does not return the CIDR value, change the **grep** expression to look for the combination of "cluster" and "CIDR" and run the command again.	192.168.0.0/16
additionalAddresses	Enter a value for the **dm.nfs.additionalAddresses** field if the network environment contains routers, proxies, gateways, or Network Address Translation (NAT) devices between the cluster nodes and the IBM Spectrum Protect Plus vSnap server. This address is used for data transfer during backup and restore operations between the cluster and vSnap server. Enter the public facing addresses for the device. The vSnap server uses an IP filter to regulate data transfer. You can specify more than one address, which can be in IP address or CIDR format. The addresses must be on the allowlist. Otherwise, the addresses are rejected. You can specify multiple addresses as follows: `dm:` `  nfs:` `    additionalAddress:` `      - x.x.x.x` `      - y.y.y.y` `      - z.z.z.z`	*x.x.x.x*
isServerInstalledOnAnotherCluster	For OpenShift environments, set the value to `false` regardless of whether the IBM Spectrum Protect Plus server is installed on the same cluster as Container Backup Support or on a separate cluster. If you are installing Container Backup Support on a Kubernetes cluster, or if the IBM Spectrum Protect Plus server is installed as a virtual appliance, set the value to `false`.	false

Table 78. Configuration parameters in the baas-values.yaml file (continued)

Parameter	Description	Default value
SPPfqdn	The DNS address for the IBM Spectrum Protect Plus server. You can specify an IP address or a fully qualified domain name (FQDN). If the IBM Spectrum Protect Plus server is installed as a virtual appliance and no DNS server is available, specify the IP address that is used for the **SPPips** parameter. If the IBM Spectrum Protect Plus server is installed in an OpenShift container environment, retrieve the DNS address by issuing the following command: `oc get route --namespace spp_server_namespace` where *spp_server_namespace* specifies the namespace in which the IBM Spectrum Protect Plus server is installed. The DNS address to use is listed in the HOST/PORT column in the command output. For example: `NAME  HOST/PORT  PATH SERVICES` `spp-rte` **`my.plus.server.example`** `/  sppproxy`	None
SPPips	The IBM Spectrum Protect Plus server IP address. If the IBM Spectrum Protect Plus server is installed as a virtual appliance, specify an IP address. If the IBM Spectrum Protect Plus server is installed in an OpenShift container environment, obtain the IP address as follows: 1. Begin with the FQDN of the route that was used in the description for the **SPPfqdn** parameter. 2. Using the URL from the output of the **oc get route** query, run the **nslookup** command or any DNS tools to resolve the URL to the IP address of the route. You can use one of the following methods: • Run the following command by using the example from the **SPPfqdn** description: `nslookup my.plus.server.example` • For more reliable results, create the Kubernetes dnsutils pod and run a DNS lookup of the IP address from that pod. For more information, see Debugging DNS Resolution. For example, run the following command, where *my.plus.server.example* is the FQDN of the route: `kubectl exec -i -t dnsutils -- nslookup` `my.plus.server.example`	*x.x.x.x*
SPPport	The IBM Spectrum Protect Plus server port. You must set the port number to 443.	443

Table 78. Configuration parameters in the `baas-values.yaml` file (continued)

Parameter	Description	Default value
`productLoglevel`	The trace levels for troubleshooting issues with the Container Backup Support transaction manager, controller, and scheduler components. The following trace levels are available: INFO, WARNING, DEBUG, and ERROR.	`INFO`
`imageRegistry`	The address of the Docker registry in your environment where the container images are loaded. If you are pulling images from the IBM Entitled Registry, you must specify `cp.icr.io/cp`. The value for the **`imageRegistry`** parameter must match the value for the DOCKER_REGISTRY_ADDRESS variable in the `baas-options.sh` file.	*docker-repo-hostname*: 5000
`imageRegistryNamespace`	The namespace of the Docker registry where the container images are loaded. The namespace does not have to be created ahead of time. To pull images from the IBM Entitled Registry, you must specify `sppc`. The value for the **`imageRegistryNamespace`** parameter must match the value for the DOCKER_REGISTRY_NAMESPACE variable in the `baas-options.sh` file.	`baas`
`minioStorageClass`	The name of the storage class to use for the MinIO server. The MinIO server is used to store the backups of cluster and namespace resources. If you do not specify a value for this parameter, the default storage class of your cluster is used. Ensure that a default storage class is defined. **Important:** To safeguard resource snapshot backups in the case where the BaaS is uninstalled or has been reinstalled, set the storage class with a **`Reclaim Policy`** with the `Retain` value specified. Backups that have been transferred to the vSnap server are not affected. Certain upgrade scenarios may also lead to losing the minIO PVC content if the **`Reclaim Policy`** is not set to `Retain`.	None
`veleroNamespace`	Specify the namespace of the Velero installation that is dedicated to IBM Spectrum Protect Plus Container Backup Support, for example, `spp-velero`. If you do not specify a value for this parameter, Velero integration is unavailable and you can use Container Backup Support to protect only PVCs.	None

Chapter 7. 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.

For information about available update files and how to obtain them from an IBM download site, see technote 6415111.

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

Review the following restrictions and tips:

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

Before you update components, it is important that you back up your IBM Spectrum Protect Plus environment as described in “Backing up the IBM Spectrum Protect Plus catalog ” on page 615.

Updating IBM Spectrum Protect Plus in a virtual appliance environment

Update IBM Spectrum Protect Plus install as a virtual appliance by using the administrative console.

Before you begin

After the IBM Spectrum Protect Plus virtual appliance is updated, it cannot roll back to a previous version without a virtual machine snapshot. Create a virtual machine snapshot of the IBM Spectrum Protect Plus appliance before you update to a new version of IBM Spectrum Protect Plus. If you later want to roll back IBM Spectrum Protect Plus to an earlier version, you must have a virtual machine snapshot. After the upgrade is completed successfully, remove the virtual machine snapshot.



Attention: Do not create virtual machine snapshots of external vSnap servers. After vSnap servers are updated, the servers cannot roll back to a previous version.

Managing updates

An IBM Spectrum Protect Plus environment includes the IBM Spectrum Protect Plus server, one or more vSnap servers, and, optionally, one or more VADP proxies. To help ensure that IBM Spectrum Protect Plus operates normally, all components in the environment must be at the same version level. Review the instructions to carefully plan and complete the update process.

Before you begin

Complete the following steps:

1. Plan a maintenance and verification period for the update process. You can estimate the required time based on the number of components in the environment that must be updated.

© Copyright IBM Corp. 2017, 2021

The process of upgrading an IBM Spectrum Protect Plus environment depends on the number of components in the environment and network speeds of the locations involved. The following table contains the three IBM Spectrum Protect Plus components and the average time, in minutes, that it takes to apply the update and successfully restart the system.

Table 79. IBM Spectrum Protect Plus components and upgrade times

Component	Time to update	Time to restart	Total
IBM Spectrum Protect Plus server	10	15	25
vSnap server	15	10 - 30	25 - 45
VADP proxy server	15	Not required.	15

2. Gather version information for the components in your environment and determine the version levels for the update process. Determine if the vSnap servers must be updated as part of the upgrade process.

3. Adjust the start times of scheduled inventory or maintenance jobs so that they will run after the maintenance and verification period is concluded.

4. End any restore or reuse jobs, including object storage restore jobs. If necessary, schedule these jobs after the maintenance and verification period is completed.

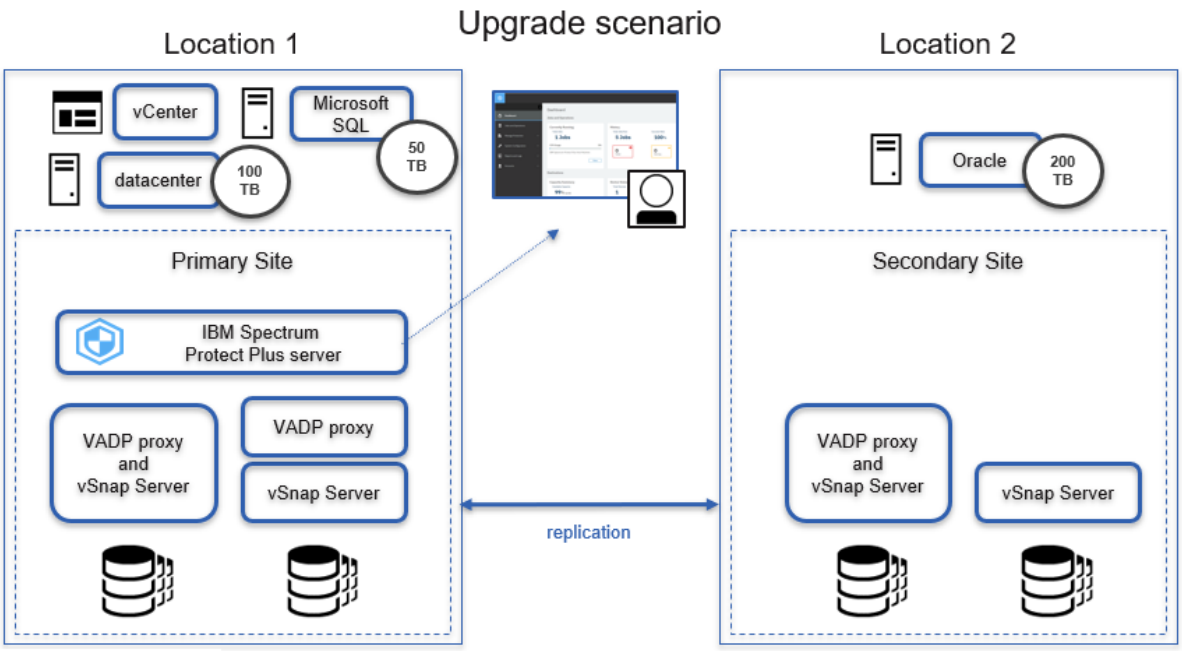
5. Pause any remaining jobs so that they do not run during the maintenance and verification period.

About this task

The procedure is based on an example environment, which includes the following components:

- 1 IBM Spectrum Protect Plus server
- 4 vSnap servers, with all four servers having replication relationships
- 2 VADP proxies co-installed with two of the vSnap servers
- 1 stand-alone VADP proxy

In the following figure, the components are displayed at their respective sites, Location 1 and Location 2:



Procedure

1. To prepare the system environment for the update process, complete the following steps:
   a) In the navigation panel, click **Manage Protection** > **Policy Overview** and then click the **Add SLA Policy** button.
   b) In the **New SLA Policy** pane, enter a policy name and click the radio button that includes the word **Catalog**. Click **Save**.
   c) Select the **Disable Schedule** checkbox and specify an appropriate retention period. From the **Target Site** list, select the site that will contain the catalog backup.
   d) Optionally, specify other options for the backup job. Click **Save**.
   e) In the navigation panel, click **Manage Protection** > **IBM Spectrum Protect Plus** > **Backup**.
   f) In the **SLA Policy** pane, select the policy that you created. Click **Save**.
   g) The policy is displayed in the **SLA Policy Status** pane. If it does not appear automatically, click the refresh button.
   h) To initiate the catalog backup, click **Actions** and then click **Start**.
   i) Verify completion of the catalog backup job. In the navigation panel, click **Jobs and Operations** to verify that the catalog backup job was completed successfully.
   j) Pause all scheduled jobs. In the navigation panel, click **Jobs and Operations** and click the **Schedule** tab. Click **Pause All Schedules**. The status for all scheduled jobs will change to **Held**.
   k) Pause all scheduled jobs. In the navigation panel, click **Jobs and Operations** and click on the **Schedule** tab. Click **Pause All Schedules**. The status for all scheduled jobs changes to `Held`.
   l) To verify that no jobs are running, click the **Running Jobs** tab. If jobs are running, allow the jobs to complete processing.
2. To prepare for updating vSnap servers, review the IBM Spectrum Protect Plus Blueprints at https://www.ibm.com/support/pages/node/1119489. Each vSnap server in your environment must be updated to the same IBM Spectrum Protect Plus version level. To update the vSnap servers, complete the following steps:
   a) Follow the steps for updating the operating system for vSnap servers, as described in "Updating the operating system for a virtual vSnap server" on page 249.

      Important: You must rename the downloaded ISO file as described in the procedure and move the file to the `/tmp` directory on the vSnap server if you wish to update the operating system.
   b) Complete the steps for updating a vSnap server, as described in "Updating a vSnap server" on page 250.

      Tip: After you update a vSnap server, it can take 15 minutes longer than in previous versions to restart the vSnap server. For more information, see https://www.ibm.com/support/pages/node/3531159.
3. Update the IBM Spectrum Protect Plus server by completing the following steps:
   a) Optional: If the IBM Spectrum Protect Plus server is deployed virtually, take a snapshot of the appliance in the appropriate hypervisor interface.
   b) Update the IBM Spectrum Protect Plus server. Follow steps 1 through 6 in the "Updating the IBM Spectrum Protect Plus server" on page 242 topic. Do not release the schedule or any jobs that are held as indicated in the last two steps.
   c) Log back in to the IBM Spectrum Protect Plus server.
4. Update VADP proxies. After you update the IBM Spectrum Protect Plus server, VADP proxies are updated automatically. However, the proxies might not be updated immediately.

   To update the VADP proxies immediately, follow the steps in the "Updating VADP proxies" on page 251 topic.
5. Verify that all components were updated successfully by completing the following steps:

a) Using the `serveradmin` account, log on to the IBM Spectrum Protect Plus administrative console. Follow the steps in "Logging on to the administrative console" on page 312.

b) Click **Product Information**. In the table, verify that the following items have the same version level: `spp-release`, `vsnap`, `vsnap-dist`, `vadp`, and `vadp-dist`.

c) Log out of the IBM Spectrum Protect Plus administrative console.

d) Load the IBM Spectrum Protect Plus splash screen by opening a supported browser and entering the following URL:

```
https://hostname/
```

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

e) Verify that the version and build on the splash screen match the `spp-release` that was displayed in the **Product Information** section of the administrative console.

f) To verify that a maintenance job can be completed successfully in the updated environment, in the navigation panel, click **Jobs and Operations** > **Schedule**. Click the options icon ⋮ next to Maintenance Job and select **Start**. Monitor the job progress through the **Jobs and Operations** pane.

6. Release scheduled jobs and, optionally, remove the snapshot. Complete the following steps:

a) Release all schedules. In the navigation panel, click **Jobs and Operations** > **Schedule**. Click **Release All Schedules**.

b) Optional: If you took a snapshot of the IBM Spectrum Protect Plus virtual appliance, you can delete the snapshot of the IBM Spectrum Protect Plus server by using the hypervisor interface. Follow the instructions in the hypervisor documentation.

What to do next

If necessary, restart any jobs that were stopped or paused during the maintenance and verification period.

Updating the IBM Spectrum Protect Plus server

Use the IBM Spectrum Protect Plus administrative console to update the IBM Spectrum Protect Plus server in a virtual appliance environment. Updating IBM Spectrum Protect Plus can be run offline or online.

Before you begin

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

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

1. Ensure that your IBM Spectrum Protect Plus environment is backed up. For instructions, see "Backing up the IBM Spectrum Protect Plus catalog " on page 615.
2. For offline updates, download the IBM Spectrum Protect Plus update package that is named *<part_number>*`.iso` to a directory on the computer that is running the browser for the administrative console. The update file is installed first.
3. Ensure that no jobs are running before starting the update procedure. Pause the schedule for any jobs that have a status of IDLE or COMPLETED. For instructions, see "Pausing and resuming jobs" on page 625.
4. Take a snapshot of the IBM Spectrum Protect Plus virtual appliance so that in the event of an update failure, the IBM Spectrum Protect Plus VM can be reverted.

For a list of download images, including the required operating system update for the virtual appliance, see technote 6415111.

About this task

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

Procedure

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

1. From a supported web browser, access the administrative console by entering the following address:

   ```
   https://hostname:8090/
   ```

   where *hostname* is the IP address of the virtual machine where the application is deployed.
2. In the login window, select one of the following authentication types in the **Authentication Type** list:

Authentication Type	Login information
IBM Spectrum Protect Plus	To log on as the IBM Spectrum Protect Plus superuser, enter the username and password. The IBM Spectrum Protect Plus superuser is the user who is assigned the SUPERUSER role.
System (recommended)	To log on as a system user, enter the `serveradmin` password. The default password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon.

3. Click **Updates and Hotfix Management** to open the updates management page.

   If you have access to the FTP site, public.dhe.ibm.com, the administrator console checks for available updates automatically and lists them.
4. Click **Run Update** to install the available updates.

   - When the updates are installed successfully, go to Step 6.
   - If you are planning to install an update from an ISO file, click **Click Here** to run the offline updates. Go to Step 5.

   Note: If you want to run online updates but can see only the offline mode, check your internet connectivity and reattempt to access the FTP site, public.dhe.ibm.com.
5. Choose the update that you want to run, as follows:

   - Online mode: Updates are listed automatically in the repository when they are made available. Click **Run Update**.
   - Offline mode: Click **Choose file** to browse for the downloaded file. The file has an `iso` or `rpm` extension like this example, `<filename>.iso`. Click **Upload Update Image (or) Hotfix**. You can select only one update file at a time.

     Important: There must be at least 4.2 GB of disk space available in the `/tmp` directory of the IBM Spectrum Protect Plus server.

   If the update fails during the prerequisite check, you must either correct the error found and then reboot the IBM Spectrum Protect Plus appliance before re-attempting the update or revert the IBM Spectrum Protect Plus appliance, correct the prerequisite issue, and then re-attempt the update. When the update is completed, IBM Spectrum Protect Plus automatically restarts.

   Important: After the IBM Spectrum Protect Plus update is completed, you must update the external vSnap and VADP proxy servers in your environment. For instructions see "Updating vSnap servers" on page 248 and "Updating VADP proxies" on page 251.
6. Clear the browser cache.

   HTML content from previous versions of IBM Spectrum Protect Plus might be stored in the cache.
7. Start the updated version of IBM Spectrum Protect Plus.

8. In the navigation panel, click **Jobs and Operations**, and then click the **Schedule** tab.
   Find the jobs that you paused.
9. Click the actions menu icon ⋮ for the job, and then click **Release Schedule**.

Related tasks

"Updating vSnap servers" on page 248
vSnap servers, both virtually deployed or physically installed, must occasionally be updated.

Updating IBM Spectrum Protect Plus in a container environment

Update IBM Spectrum Protect Plus components in the OpenShift web console or the IBM Spectrum Protect Plus user interface.

Updating IBM Spectrum Protect Plus by using the OpenShift web console

You can update IBM Spectrum Protect Plus on an OpenShift cluster by first updating the IBM Spectrum Protect Plus operator and then updating the IBM Spectrum Protect Plus server instance.

Before you begin

Ensure that Red Hat OpenShift Container Platform 4.5 or later is running on your cloud environment, and you have cluster administrator privileges for your OpenShift cluster.

Updating the IBM Spectrum Protect Plus operator

You must update the IBM Spectrum Protect Plus operator before you can update the IBM Spectrum Protect Plus instance on OpenShift Container Platform.

About this task

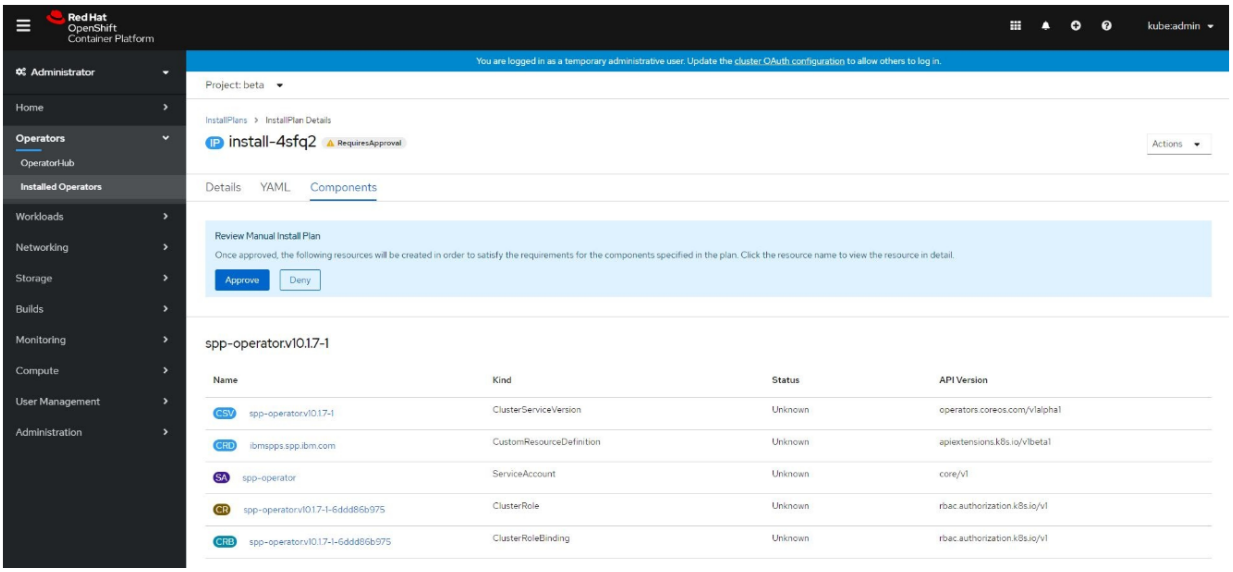
When you installed the IBM Spectrum Protect Plus operator, you were given the choice to configure a manual or automatic approval strategy for updating the operator. If you chose the manual approval strategy, when an operator update is available, an update request is created by OpenShift Operator Lifecycle Manager (OLM). You must manually approve the update request.

If you chose the automatic approval strategy, no interventions are required from you when an update is available.

Procedure

To update the IBM Spectrum Protect Plus operator by using the manual approval strategy, complete the following steps:

1. Log on to the OpenShift web console as the cluster administrator.
2. In the navigation panel, click **Operators** > **Installed Operators**.
3. On the **Installed Operators** page, ensure that the project that you created for IBM Spectrum Protect Plus is selected in the **Project** list.
4. Select the IBM Spectrum Protect Plus operator.
5. On the **Operator Details** page, click the **Subscription** tab and locate **Install Plan**.
6. Click the link that is displayed under **Install Plan** and click **Preview Install Plan**.
7. Review the install plan and click **Approve** to update to the specified version of the IBM Spectrum Protect Plus operator.



Results

When the operator is successfully updated, the upgrade status of the operator subscription is shown as **Up to date**. You can further verify that the upgraded operator is installed by selecting **Operators** > **Installed Operators** to verify that the ClusterServiceVersion (CSV) of the IBM Spectrum Protect Plus operator is installed successfully.

What to do next

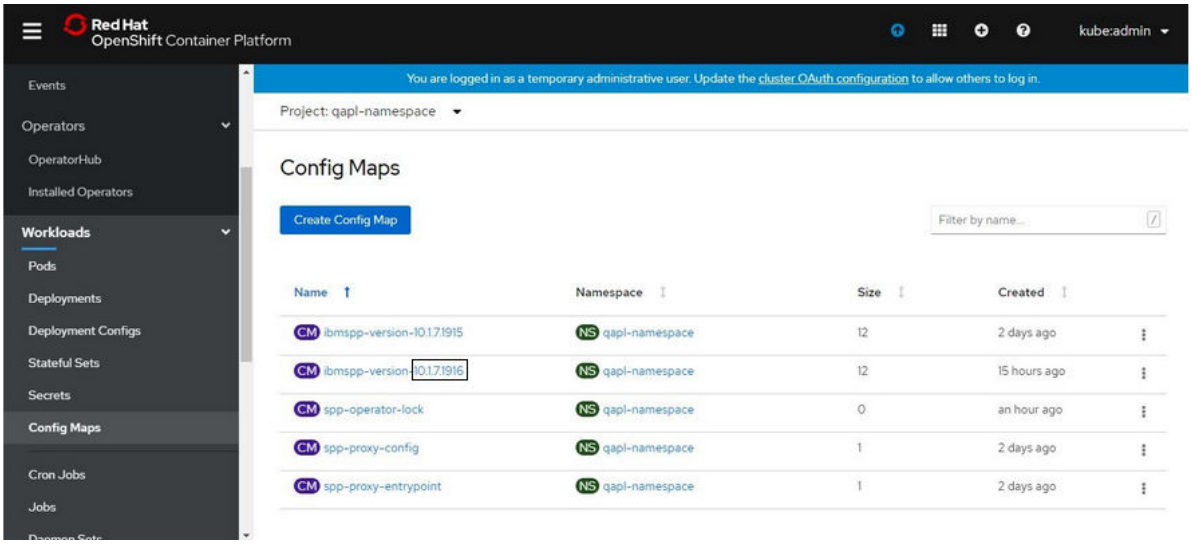
Update the IBM Spectrum Protect Plus instance in the OpenShift web console. For instructions, see “Updating an IBM Spectrum Protect Plus instance” on page 245.

Updating an IBM Spectrum Protect Plus instance

After the IBM Spectrum Protect Plus operator is updated, you can update the instance of IBM Spectrum Protect Plus. In this way, you can take advantage of the latest features of IBM Spectrum Protect Plus in the container environment.

Procedure

1. Log on to the OpenShift web console as the cluster administrator.
2. Check whether an updated version of IBM Spectrum Protect Plus is available by taking the following actions:
   a) In the navigation panel, click **Workloads** > **Config Maps**.
   b) Ensure that the project that you created for IBM Spectrum Protect Plus is selected.
   c) On the **Config Maps** page, note the version of IBM Spectrum Protect Plus that you want to upgrade to. The available IBM Spectrum Protect Plus versions are identified by the `ibmspp-version-` prefix.



3. In the navigation panel, click **Operators** > **Installed Operators**.
4. On the **Installed Operators** page, ensure that the project that you created for IBM Spectrum Protect Plus is selected.
5. Select the IBM Spectrum Protect Plus operator and click the **IBM Spectrum Protect Plus** tab.
6. Select the IBM Spectrum Protect Plus instance and click the **YAML** tab.
7. In the displayed YAML file, find the **`version`** field and replace the existing version number with the new version number, including any periods.
   For example:

   ```
   version: 10.1.7.1916
   ```

8. Update the instance of IBM Spectrum Protect Plus by clicking **Save**.

What to do next

You can verify that the IBM Spectrum Protect Plus instance was created updated by completing the following steps:

1. In the navigation panel of the OpenShift web console, click **Operators** > **Installed Operators**.
2. Click IBM Spectrum Protect Plus operator from the list of installed operators.
3. On the **Operator Details** page, click the IBM Spectrum Protect Plus tab. The list of running instances is displayed.
4. Click the name of an instance to show its status.
5. Scroll to the **Conditions** section of the page and review the status. The instance is updated successfully when the `True` message is displayed in the **Status** column and the `Successful` message is displayed in the **Reason** column.

Updating IBM Spectrum Protect Plus by using the user interface

Use the IBM Spectrum Protect Plus user interface to update the product online.

Before you begin

Complete the following tasks before you start the update process:

- Update the IBM Spectrum Protect Plus OpenShift operator by using one of the following methods:
  - Complete the steps in “Updating the IBM Spectrum Protect Plus operator” on page 244 to update the operator manually.

  - Set the operator approval status to **Automatic**, which updates the operator automatically when a new version is available.
- Log in to the IBM Spectrum Protect Plus as the superuser. The IBM Spectrum Protect Plus superuser is the user who is assigned the SUPERUSER role.
- Ensure that your IBM Spectrum Protect Plus environment is backed up before you run updates. For more information about backing up your environment, see “Backing up the IBM Spectrum Protect Plus catalog ” on page 615.
- Ensure that no jobs are running during the update procedure. Pause the schedule for any jobs that have a status of IDLE or COMPLETED.
- Clear the cache for the browser that you use to open the IBM Spectrum Protect Plus user interface. HTML content from previous versions of IBM Spectrum Protect Plus might be stored in the cache.

Procedure

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

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

   The installed versions are shown.
2. Click **Check for updates**.
3. Click the version that you want to update to, and then click **Proceed to update**.
4. Review the summary information, and then click **Update**.
5. Click **Yes** to confirm the update.

   A window opens showing the update state for each IBM Spectrum Protect Plus component.
6. When the status of all components is **Running**, click **Reload** and then click **Reload** again in the confirmation dialog box.

   The IBM Spectrum Protect Plus login page is opened with a message that the server is being brought up. When this message is no longer shown, you can log in to IBM Spectrum Protect Plus.
7. Log in to IBM Spectrum Protect Plus.
8. In the navigation panel, click **Jobs and Operations**, and then click the **Schedule** tab.

   Find the jobs that you paused.
9. Click the actions menu icon for the job, and then click **Release Schedule**.

Updating IBM Spectrum Protect Plus in a cloud environment

When IBM Spectrum Protect Plus is installed in a cloud environment, the steps to update the product components are similar to the steps for updating components in an on-premises environment.

About this task

The steps to update the components depend on the cloud environment that you are using: all on cloud or hybrid.

Hybrid environment

If you are updating IBM Spectrum Protect Plus in a hybrid cloud environment, the IBM Spectrum Protect Plus server is installed on-premises and the vSnap server is installed on the cloud.

To update the IBM Spectrum Protect Plus server on premises, follow the instructions in “Updating the IBM Spectrum Protect Plus server” on page 242.

To update the vSnap server on the cloud, use a virtual private network (VPN) or other connection such as a Bastion host to access the host where the application is installed. To update the vSnap server, follow the instructions in “Updating vSnap servers” on page 248.

All-on-cloud environment

If you are updating IBM Spectrum Protect Plus in an all-on-cloud environment, the IBM Spectrum Protect Plus server and the vSnap server are installed on the cloud.

To update these applications, use a VPN or other connection such as a Bastion host to access the host where the applications are installed, and then follow the instructions in "Updating the IBM Spectrum Protect Plus server" on page 242 and "Updating vSnap servers" on page 248.

Related concepts

"IBM Spectrum Protect Plus on IBM Cloud" on page 17
IBM Spectrum Protect Plus is available as a software offering in the IBM Cloud catalog or as an IBM Cloud for VMware Solutions service.

"IBM Spectrum Protect Plus on the AWS cloud platform" on page 18
IBM Spectrum Protect Plus on the Amazon Web Services (AWS) cloud platform is a data protection solution for users who want to protect databases that are running on AWS. In addition, users can protect virtual machines that are managed by VMware Cloud (VMC) on AWS while having the IBM Spectrum Protect Plus server installed on VMC and the vSnap server installed on an AWS Virtual Private Cloud (VPC).

"IBM Spectrum Protect Plus on the Microsoft Azure cloud platform" on page 19
IBM Spectrum Protect Plus on the Microsoft Azure cloud platform is a data protection solution for users who want to protect one or more databases that are running on Azure.

Updating 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.

You might also be required to update the operating system for the vSnap servers prior to updating the servers. For operating system requirements, see "Component requirements " on page 27.

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 250

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

Updating the operating system for a virtual vSnap server

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

Before you begin

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

Before you begin the update process, ensure that you have backed up your IBM Spectrum Protect Plus environment as described in "Backing up the IBM Spectrum Protect Plus catalog " on page 615. For information on obtaining the ISO file, see "Updating the IBM Spectrum Protect Plus server" on page 242.

Restriction: The ISO should not be used if updating a physical Red Hat Enterprise Linux server. It should only be used on OVA deployments.

Procedure

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

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

   Important: It is critical to rename the downloaded ISO file as described in this step and move it to the /tmp directory on the vSnap server if you wish to update the operating system.
2. Next, follow the instructions found in the "Updating a vSnap server" on page 250 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 250

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

Updating a vSnap server

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

Before you begin

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

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

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

1. Ensure that you have backed up your IBM Spectrum Protect Plus environment as described in “Backing up the IBM Spectrum Protect Plus catalog ” on page 615.
2. Download the vSnap update file. The image name that contains *vsnap-dist-el7-<version-build>* is for CentOS and Red Hat Enterprise Linux 7.x and must be used for upgrades. The image name has an associated **Passport Advantage Part Number**. Download the corresponding run file, *<part_number>*.run, for the image name from Passport Advantage Online. For information about downloading files, see technote 6415111.
3. Copy the upgrade file to a temporary location on the vSnap server.

Procedure

To update a vSnap server, complete the following steps:

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

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

3. Run the installer by issuing the following command:

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

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

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

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

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

   The vSnap packages are installed.
4. After the vSnap packages are installed, start the updated version of the vSnap server.
5. In the navigation panel, click **Jobs and Operations**, and then click the **Schedule** tab.

   Find the jobs that you paused.

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

Additional steps for updating virtual machines in Hyper-V Replica environments

Beginning with IBM Spectrum Protect Plus Version 10.1.5, you can protect virtual machines (VMs) that are enabled to use the Hyper-V Replica feature.

IBM Spectrum Protect Plus processes the data on the source and replicated instances of the VMs separately. For example, if a VM named `VM1` is on the Hyper-V host named `Host1` and the VM is replicated to `Host2`, IBM Spectrum Protect Plus assigns the IDs `VM1@Host1` and `VM1@Host2` to the VMs. You can then select one or both of the VMs for data protection.

Considerations for VMs that are defined in existing SLA policies

If you update IBM Spectrum Protect Plus, you might have to take additional steps to ensure that data protection continues for VMs that are currently included in your service level agreement (SLA) policies.

An SLA policy can *implicitly* or *explicitly* include a replicated VM. You might be required to update the SLA policy when you update to IBM Spectrum Protect Plus V10.1.5 or later.

An example of an SLA policy that implicitly includes a replicated VM is a scenario in which the policy protects all VMs on `Host1`, which contains the VM `VM1`. `VM1` is replicated to `Host2`. In this scenario, a change to the SLA policy is not required after you update IBM Spectrum Protect Plus. The SLA policy creates a full backup of the instance of `VM1` on `Host2` and creates a new full backup of the instance of `VM1` on `Host1`. Existing backups of `VM1` on `Host1` that were created before the update will expire based on the SLA policy retention settings.

An example of an SLA policy that explicitly includes a replicated VM is a scenario in which the policy protects `VM1` on `Host1`, and `VM1` is replicated to `Host2`. In this scenario, you must re-add the instance of the VM on each host to the SLA policy after you update IBM Spectrum Protect Plus.

Updating VADP proxies

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

Before you begin

Before you begin, ensure that you have backed up your IBM Spectrum Protect Plus environment as described in “Backing up the IBM Spectrum Protect Plus catalog ” on page 615.

Note: Only VADP proxies registered with IBM Spectrum Protect Plus will be updated. If the VADP proxy is not registered with IBM Spectrum Protect Plus, the VADP component will not be updated.

Procedure

If a VADP proxy update is available for external proxies during a restart of the IBM Spectrum Protect Plus virtual appliance, the update will be automatically applied to any VADP proxy associated with an identity. To associate a VADP proxy with an identity, navigate to **System Configuration** > **VADP Proxy**. Click the ellipses icon ••• and select **Edit**. Select **Use existing user** and choose a previously entered identity in **Select user** for the VADP proxy server.

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

1. Navigate to the **System Configuration** > **VADP Proxy** page in IBM Spectrum Protect Plus.
2. The **VADP Proxy** page displays each proxy server. If a newer version of the VADP proxy software is available, an update icon displays in the **Status** field.

3. Ensure that there are no active jobs that use the proxy, and then click the update icon .

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

When attempting to install a VADP proxy to a supported, stand-alone Linux deployment or when the account used to register the VADP proxy through IBM Spectrum Protect Plus is a non-root user, special instructions must be followed. Specifically, the username and password for the account used to register the VADP proxy must also exist on the machine to which the proxy is being installed or updated and have a matching `sudoers` configuration file. The `sudoers` configuration must allow the user to run commands without a password.

1. Log in to the VADP server as the `root` user.
2. In the case of a stand-alone VADP proxy deployed on a supported version of Linux, create a new user and assign a password. This is the account that will subsequently be used to register the VADP proxy through the IBM Spectrum Protect Plus user interface. In this example, the variable *vadpuser* is the username used to register the VADP proxy.

   ```
   # useradd vadpuser
   # passwd vadpuser_password
   ```

   In the case of updating a VADP proxy, verify that the user that was used to register the VADP server exists.

   ```
   # cat /etc/passwd | grep vadpuser
   ```

3. When installing a stand-alone VADP proxy on a supported version of Linux, you may need to install the `nfs-utils` package if it is not already installed. Answer yes ("y") to the prompts.

   ```
   # yum install nfs-utils
   ```

4. Next, create a `sudoers` configuration file in the `/etc/sudoers.d/` directory. Write the "`Defaults !requiretty`..." text to the file and save it by pressing CTRL+D on the keyboard when done.

   ```
   # cat /etc/sudoers.d/vadpuser
   Defaults !requiretty
        vadpuser ALL=NOPASSWD: /tmp/cdm_guestapps_vadpuser/runcommand.sh
   <<Press CTRL+D>>
   ```

5. Finally, set the appropriate permissions on the file.

   ```
   # chmod 0440 vadpuser
   ```

What to do next

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

Action	How to
Run the VMware backup job.	See “Backing up VMware data” on page 340. The proxies are indicated in the job log by a log message similar to the following text: `Run remote vmdkbackup of MicroService: http://<proxy` *`nodename`*`, IP:`*`proxy_IP_address`*

Related tasks

“Creating VADP proxies” on page 346

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

Related reference

“Editing firewall ports” on page 122
Use the provided examples as a reference for opening firewall ports on remote VADP proxy servers or application servers. You must restrict port traffic to only the required network or adapters.

Applying early availability updates

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

About this task

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

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

Chapter 8. Getting off to a quick start

To start using IBM Spectrum Protect Plus, you must define resources that you want to protect and create service level agreement (SLA) policies, also known as backup policies, for those resources. This getting started section provides these and other steps required to set up and start using IBM Spectrum Protect Plus to back up data. Other tasks such as copying and restoring data are discussed in detail in other areas of the documentation.

Ensure that you followed the instructions in the IBM Spectrum Protect Plus Blueprints to determine how to size, build, and place the components in your IBM Spectrum Protect Plus environment and that the tasks listed in the "Deployment storyboard for IBM Spectrum Protect Plus" on page 1 are complete.

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

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

Action	Owner	Description
Start IBM Spectrum Protect Plus	Infrastructure administrator and application administrator	The infrastructure administrator starts the application for the first time by using the default `admin` user account with the password `password`. The administrator is prompted to reset the username and password for this account. The administrator cannot reset the user name to `admin`, `root`, or `test`. After the initial startup, the application administrator can start the application by using this user account, which is referred to as the IBM Spectrum Protect Plus superuser account, or another account that the infrastructure administrator creates.
"Manage sites" on page 258	Infrastructure administrator	A site is used to group vSnap servers based on a physical or logical location to help quickly identify and interact with backup data. A site is assigned to a vSnap server when the server is added to IBM Spectrum Protect Plus. The default sites are named Primary and Secondary, but a custom site can also be created.

© Copyright IBM Corp. 2017, 2021

Action	Owner	Description
Create backup policies	Infrastructure administrator	Backup policies define the parameters that are applied to backup jobs. These parameters include the frequency and retention of backups and the options to replicate data from one vSnap server to another and to copy backup data to secondary backup storage for longer-term protection. Backup policies also define the target site to for backing up data. A site can contain one or more vSnap servers. Backup policies are called SLA policies in IBM Spectrum Protect Plus.
Create a user account for the application administrator	Infrastructure administrator	User accounts determine the resources and functions that are available to the user.
Add resources to protect	Application administrator	Resources are entities that you want to protect. After a resource is registered, an inventory of the resource is captured and added to the IBM Spectrum Protect Plus inventory.
Add resources to a job definition	Application administrator	Job definitions associate the resources that you want to protect with one or more SLA policies. The options and schedules that are defined in the SLA policies are used for backup jobs for the resources.
“Back up the IBM Spectrum Protect Plus catalog” on page 268	Application administrator	The catalog consists of data such as application configuration settings, SLA policies, backup storage settings, and information about registered resources, restore points, and jobs.
Start a backup job	Application administrator	Backup jobs are started as defined in the SLA policy that is associated with the job definition. You can also manually start a job.
Run a report	Application administrator	IBM Spectrum Protect Plus provides predefined reports that you can run with default parameters or modify to create custom reports.

Start IBM Spectrum Protect Plus

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

Procedure

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

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

   ```
   https://hostname
   ```

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

   Hostname for a container installation

   The hostname must be in the following format:

   ```
   instancename-spp.routerCanonicalHostname
   ```

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

   Hostname for a virtual appliance installation

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

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

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

   This user account is the superuser account and is assigned the SUPERUSER role. This role is assigned to only one IBM Spectrum Protect Plus user. The SUPERUSER role provides the user with access to all IBM Spectrum Protect Plus functions. For more information about the superuser account, see “Managing the superuser account” on page 655.

3. Click **Sign In**.

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

   The following rules are enforced when creating a new password:

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

Manage sites

A site is used to group vSnap servers based on a physical or logical location to help quickly identify and interact with backup data. A site is assigned to a vSnap server when the server is added to IBM Spectrum Protect Plus.

About this task

Review the available sites by clicking **System Configuration** > **Site** in the navigation panel and decide whether you want to add new sites or edit the existing ones for your vSnap servers.

Note: You can change the site name and other options for the default Primary and Secondary sites.

Procedure

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

1. In the navigation panel, click **System Configuration** > **Site**.
2. To add new sites or edit existing sites, take the appropriate action:

Action	How to
Add a new site.	a. Click **Add Site**. b. Enter a site name. c. Optional: Select options to manage backup and replication operations as described in “Adding a site” on page 291. **Important:** Modify the **VMware VM allocation** settings only at the direction of IBM Support. d. Click **Save**.
Edit a site.	a. Click **Edit Site**. b. Click the edit icon that is associated with a site. c. Optional: Select options to manage backup and replication operations as described in “Editing a site” on page 292. **Important:** Modify the **VMware VM allocation** settings only at the direction of IBM Support. d. Click **Save**.

Related concepts

“Product components” on page 6
The IBM Spectrum Protect Plus solution is provided as a container or virtual appliance that includes storage and data movement components.

“Managing sites” on page 291

A *site* is an IBM Spectrum Protect Plus policy construct that is used to manage the placement of data in an environment.

Create backup policies

Backup policies, which are also referred to as service level agreement (SLA) policies, define parameters that are applied to backup jobs. These parameters specify the frequency of backup jobs, the retention period for backed up data, and other directives for backup operations.

Before you begin

IBM Spectrum Protect Plus includes default SLA policies as described in Chapter 10, "Managing SLA policies for backup operations," on page 321. You can use the default policies as they are or modify the policies. You can also create custom SLA policies.

For example purposes, create the following SLA policies:

An SLA policy for the resources that you want to protect
This SLA policy is applied to backup jobs for the resources that you want to protect.

An SLA policy for the IBM Spectrum Protect Plus catalog
This SLA policy is applied to backup jobs for the IBM Spectrum Protect Plus catalog. The catalog consists of data such as application configuration settings, SLA policies, backup storage settings, and information about registered resources, restore points, and jobs.

A catalog backup is helpful in case of a natural disaster or other unexpected event. By accessing data in the catalog backup, you can help to restore normal system operations.

To optimize backup jobs, always create separate SLA policies for backing up resources and for backing up the IBM Spectrum Protect Plus catalog. To balance the system workload, ensure that the SLA policies do not define run times that would cause the jobs to run concurrently.

About this task

This task does not include instructions for enabling replication for vSnap servers or for copying data to secondary backup storage, which are optional features. For information about how to set up these features in the SLA policy, see "Creating an SLA policy for hypervisors, databases, and file systems" on page 324.

The following figures show separate SLA policies for resource and catalog backups. The resource SLA policy named `Copper` runs every Monday at midnight with a retention period of 1 month. The catalog SLA policy named `Catalog_Backup` runs every day at midnight with a retention period of 15 days.

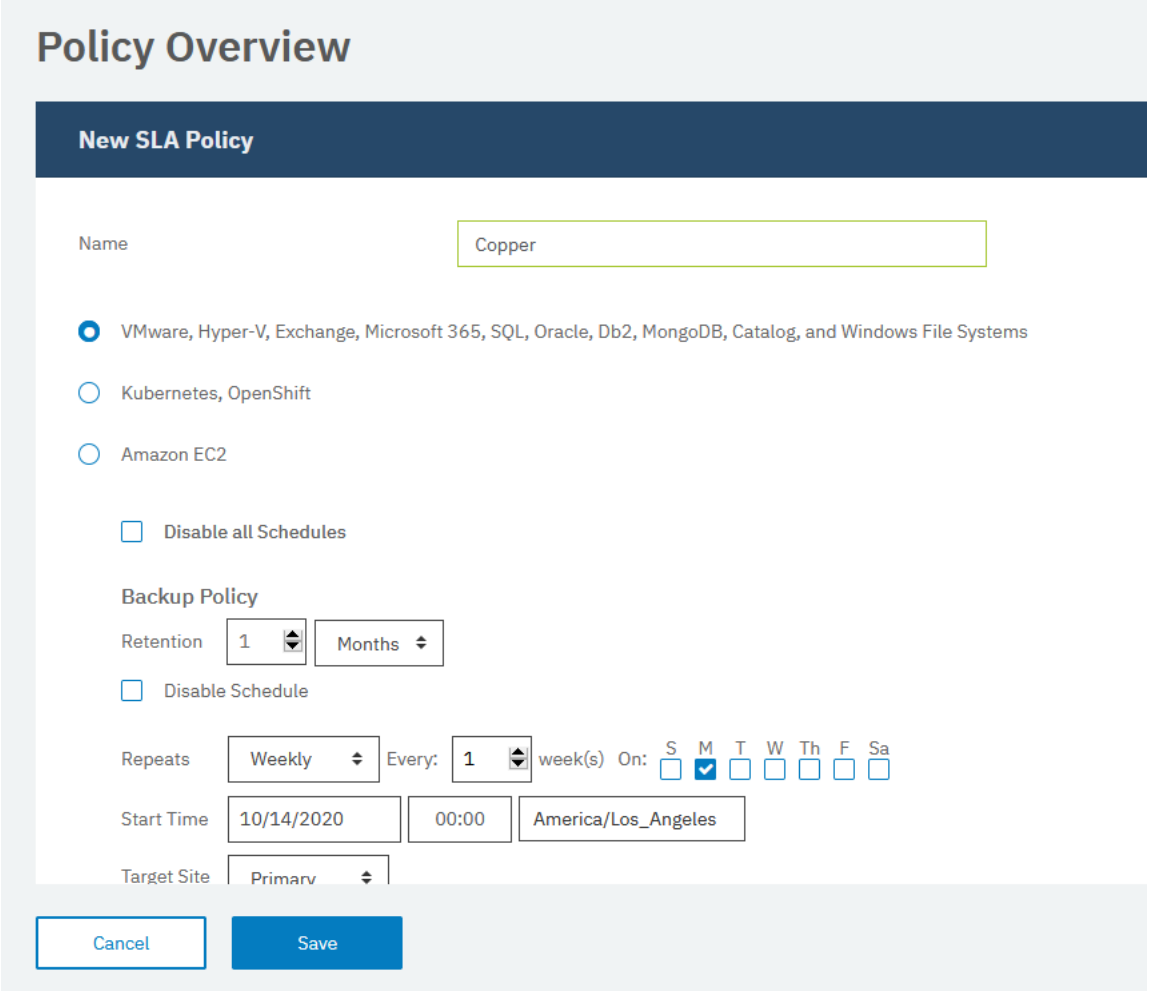


Figure 13. Creating an SLA policy for resource backup jobs

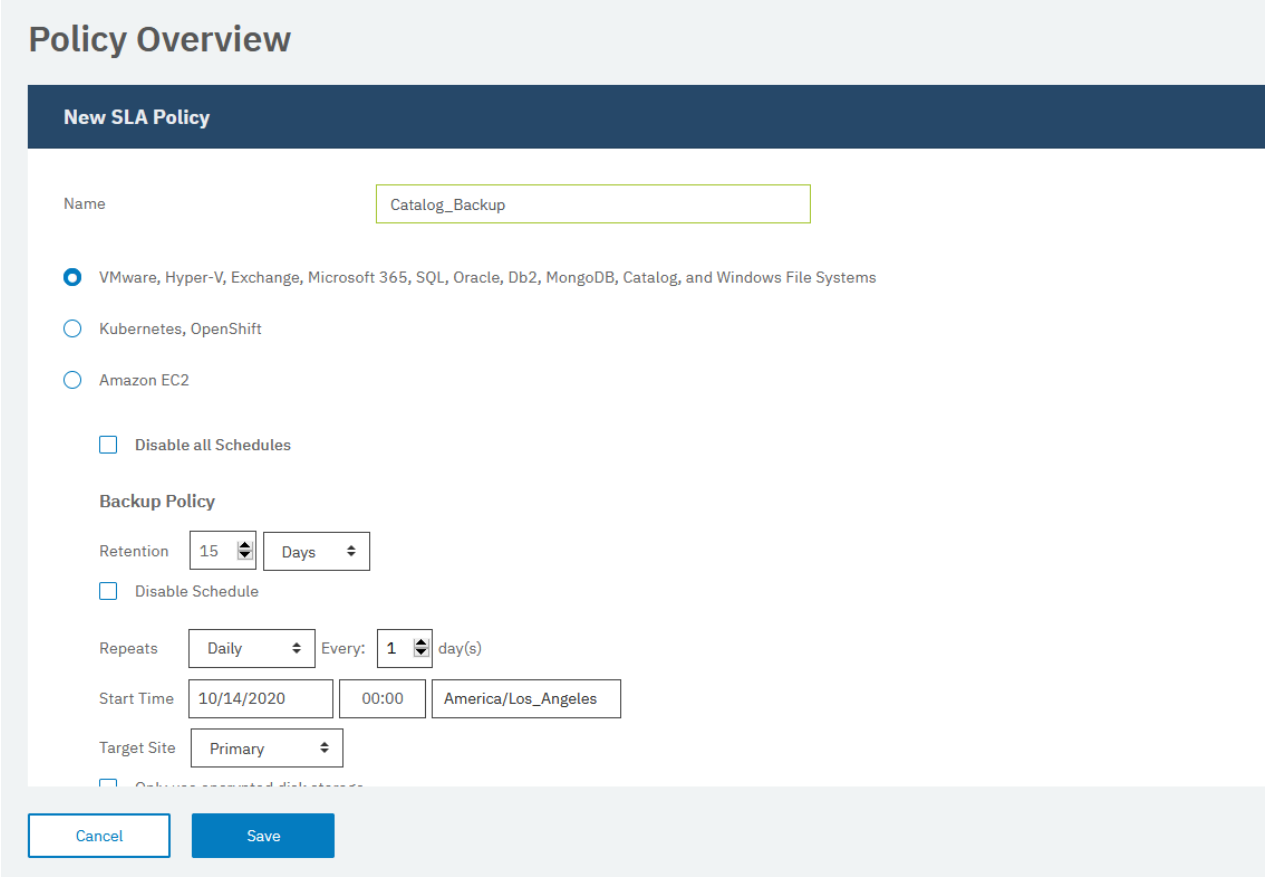


Figure 14. Creating an SLA policy for catalog backup jobs

Procedure

To create an SLA policy, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Policy Overview**.
2. Click **Add SLA Policy**.

   The **New SLA Policy** pane is displayed.
3. In the **Name** field, enter a name that provides a meaningful description of the SLA policy.
4. Click **VMware, Hyper-V, Exchange, Microsoft 365, SQL, Oracle, Db2, MongoDB, Catalog, and Windows File Systems** to set the backup options.

   Use this option to create an SLA policy for all resource types other than containers and Amazon EC2. To create SLA policies for containers and EC2, follow the instructions in "Creating an SLA policy for containers" on page 329 and "Creating an SLA policy for Amazon EC2 instances" on page 328.
5. In the **Backup Policy** section, set the following options for backup operations. These operations occur on the vSnap servers that are defined in the **System Configuration** > **Backup Storage** > **Disk** window.

   Retention
   Specify the retention period for the backup snapshots.

   Disable Schedule
   Select this checkbox to create the main policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand.

   Repeats
   Enter a frequency for backup jobs by selecting **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

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

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site
Select the target backup site for backing up data.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, the IBM Spectrum Protect Plus server manages data placement across the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a vSnap server, are not shown.

Only use encrypted disk storage
Select this checkbox to back up data to encrypted vSnap servers if your environment includes encrypted and unencrypted servers.

Restriction: If this option is selected and no encrypted vSnap servers are available, the associated job fails.

6. Click **Save**. The SLA policy can now be applied to backup job definitions as shown in "Add resources to a job definition" on page 266.

Related concepts

"Replicate backup-storage data " on page 14
When you enable replication of backup data, data from one vSnap server is asynchronously replicated to another vSnap server. For example, you can replicate backup data from a vSnap server on a primary site to a vSnap server on a secondary site.

"Copy snapshots to secondary backup storage" on page 15
The vSnap server is the primary backup location for snapshots. All IBM Spectrum Protect Plus environments have at least one vSnap server. Optionally, you can copy snapshots from a vSnap server to secondary backup storage.

"Managing SLA policies for backup operations" on page 321
Service level agreement (SLA) policies, also known as backup policies, define parameters for backup jobs. These parameters include the frequency and retention period of backups and the option to replicate or copy backup data. You can use predefined SLA policies, or customize them to meet your needs.

Create a user account for the application administrator

Create a user account for an administrator who can run backup and restore operations for the resources in your environment.

Before you begin

For example purposes, the following steps show how to create an account for an individual user who is responsible for protecting VMware data. This account uses an existing user role and resource group.

To create an account for an LDAP group, see "Creating a user account for an LDAP group" on page 653.

To create custom user roles and resource groups, see "Creating a resource group" on page 644 and "Creating a role" on page 650

Procedure

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

1. In the navigation panel, click **Accounts** > **User**.
2. Click **Add User**. The **Add User** pane is displayed.
3. Click **Select the type of user or group you want to add** > **Individual new user**.

4. Enter a name and password for the application administrator.
5. In the **Assign Role** section, select **VM Admin**.
   The permissions are shown in the **Permission Groups** section.

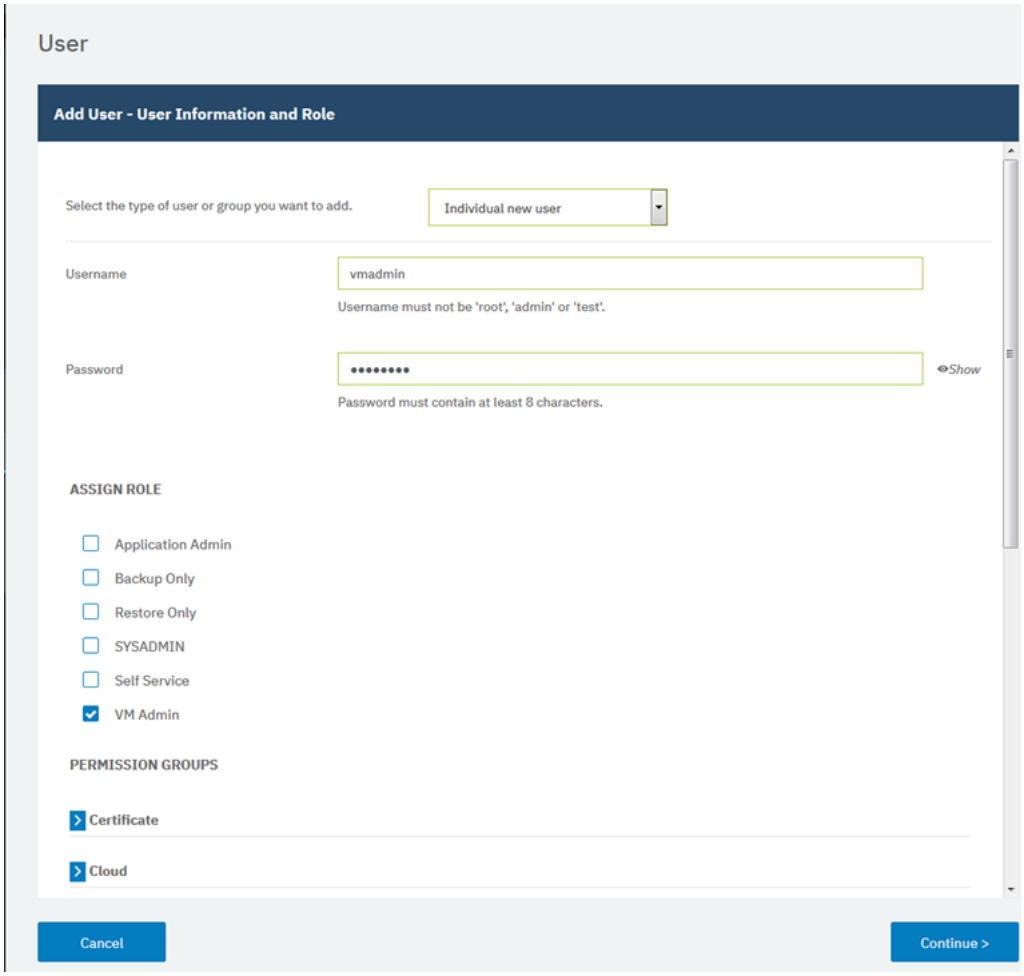


Figure 15. Creating a user account and assigning a role

6. Click **Continue**.
7. In the **Add Users - Assign Resources** section, select the **All Resources** resource group, and then click **Add resources**.
   The resource group is added to the **Selected Resources** section.

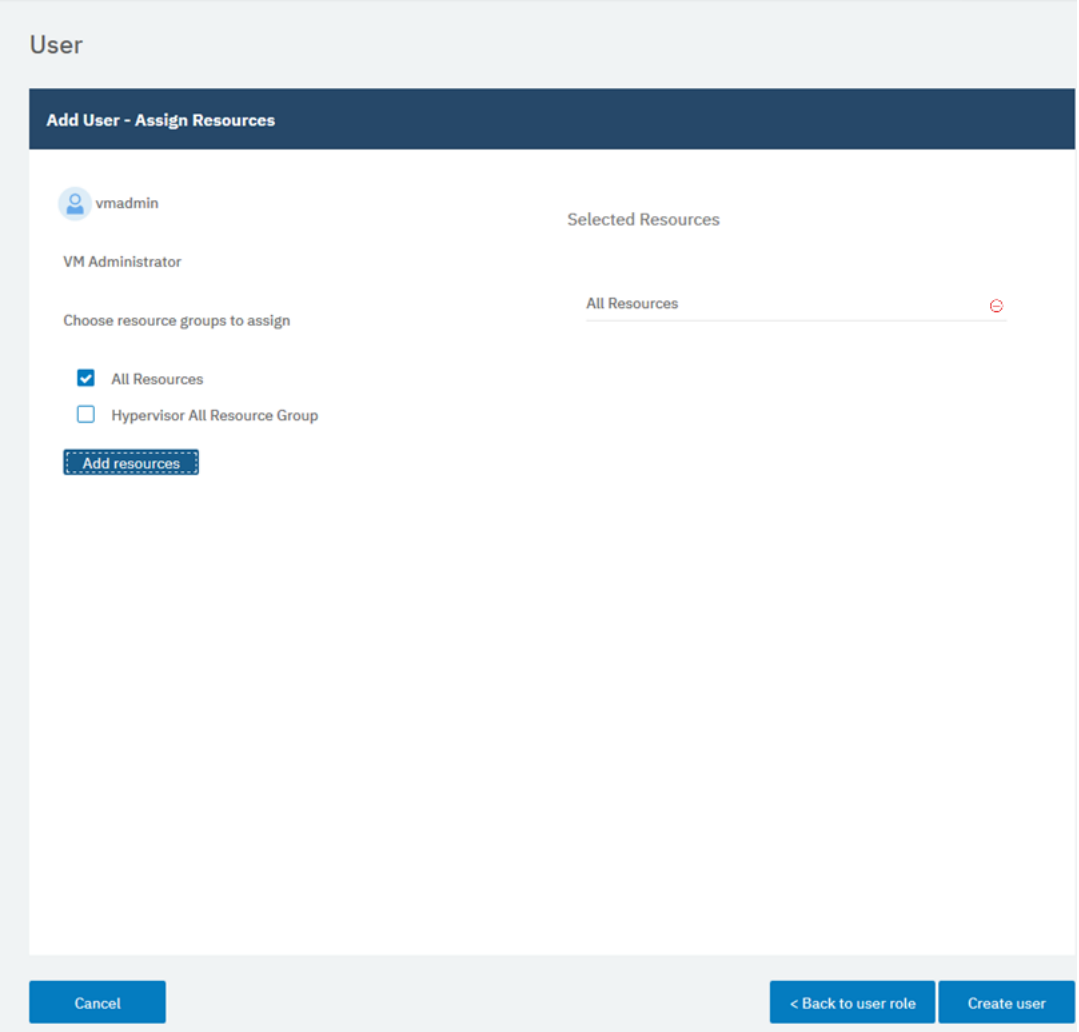


Figure 16. Selecting a resource group for the user account

8. Click **Create user**.

Related concepts

"Managing user access" on page 643
By using role-based access control, you can set the resources and permissions available to IBM Spectrum Protect Plus user accounts.

Back up resources

Add the resources that you want to protect to IBM Spectrum Protect Plus and create a backup job for those resources.

Add resources to protect

Resources are entities that you want to protect. After a resource is registered, an inventory of the resource is captured and added to the IBM Spectrum Protect Plus inventory, enabling you to complete backup and restore jobs, as well as to run reports.

About this task

For example purposes, this task describes how to add a VMware instance. To add other types of resources, see the instructions by resource type in the following sections:

- Chapter 11, "Protecting virtualized systems," on page 335

- Chapter 12, “Protecting file systems,” on page 383
- Chapter 13, “Protecting containers,” on page 401
- Chapter 14, “Protecting data on cloud systems,” on page 491
- Chapter 15, “Protecting databases,” on page 497

Procedure

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

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **VMware**.
2. Click **Manage vCenter**, and then click **Add vCenter**.
3. Populate the fields in the **vCenter Properties** section:

   Hostname/IP
   Enter the resolvable IP address or a resolvable path and machine name.

   Use existing user
   Enable to select a previously entered user name and password for the vCenter Server instance.

   Username
   Enter your user name for the vCenter Server instance.

   Password
   Enter your password for the vCenter Server instance.

   Port
   Enter the communications port of the vCenter Server instance. Select the **Use TLS** checkbox to enable an encrypted Transport Layer Security (TLS) connection. The typical default port is 80 for non TLS connections or 443 for TLS connections.
4. In the **Options** section, configure the following option:

   Maximum number of VMs to process concurrently per ESX server and per SLA
   Set the maximum number of concurrent VM snapshots to process on the ESX server. The default setting is 3.

   The following figure shows populated fields.

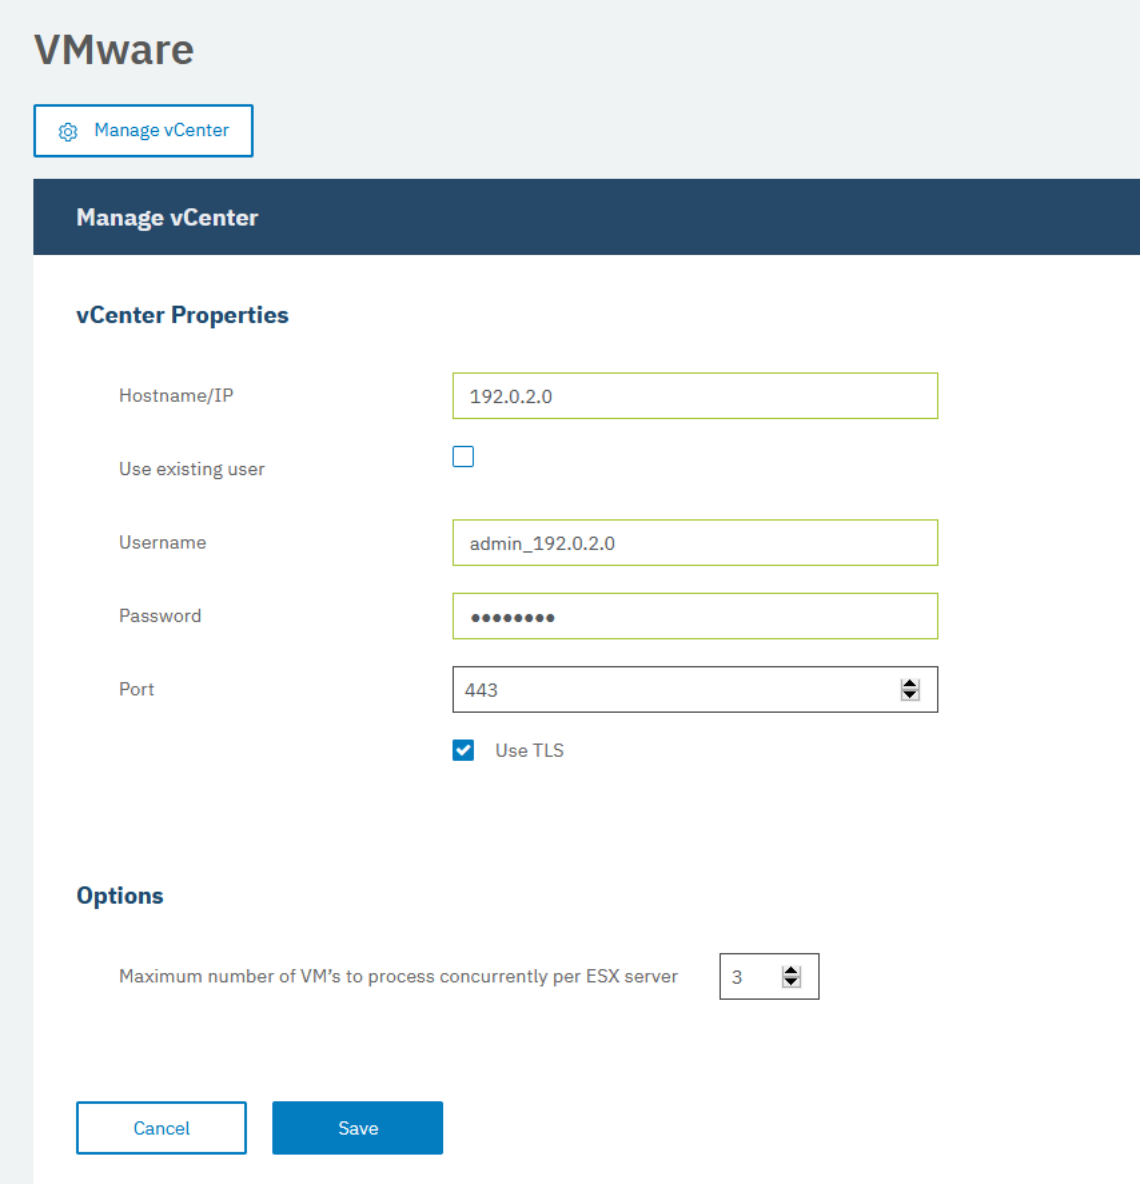


Figure 17. Adding a vCenter Server instance

5. Click **Save**.

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

Add resources to a job definition

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

About this task

For example purposes, this task describes how to select an SLA policy for resources that are in a VMware vCenter.

Procedure

To select an SLA policy, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **VMware**.
2. Select the resources that you want to back up. You can select all resources in a vCenter or drill down to select specific resources.

   Use the search function to search for available resources and toggle the displayed resources by using the **View** filter. Available options are **VMs and Templates**, **VMs**, **Datastore**, **Tags and Categories**, and **Hosts and Clusters**. Tags, which are applied in vSphere, make it possible assign metadata to virtual machines.

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

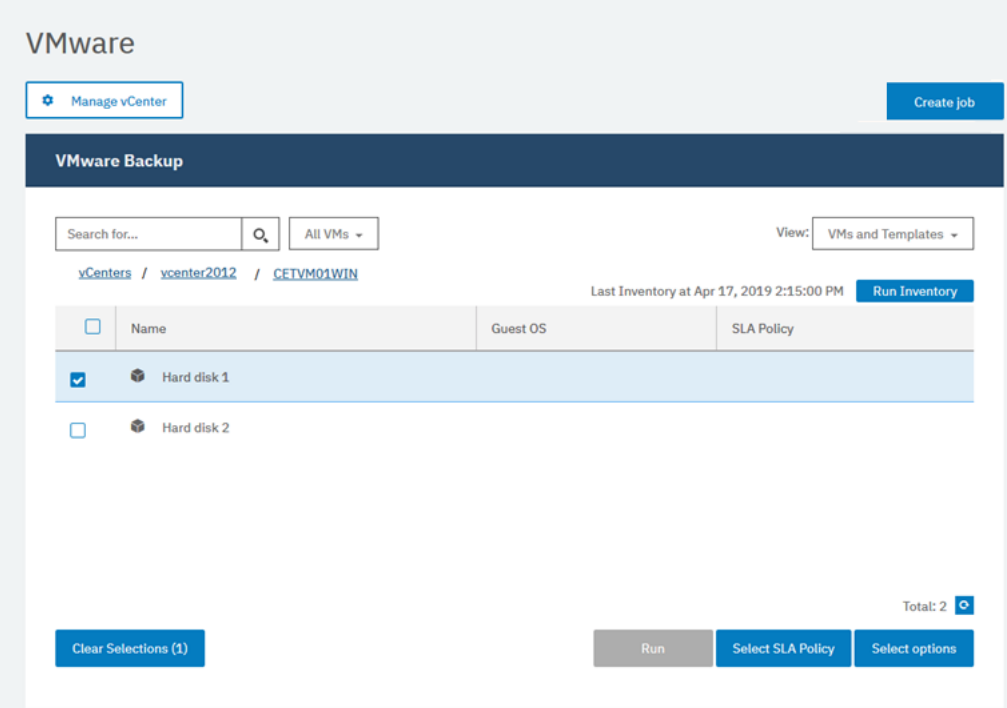


   Figure 18. Selecting resources for backup

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

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

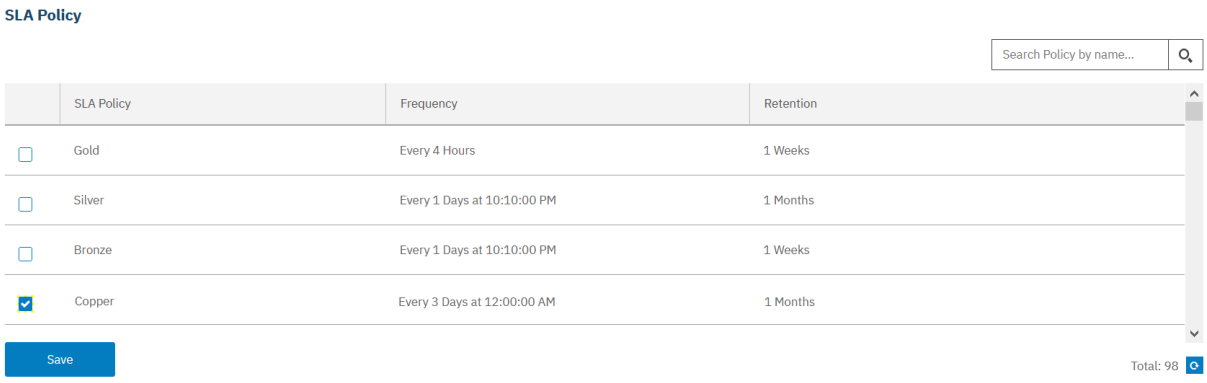


   Figure 19. Selecting an SLA policy

4. To create the job definition by using default options, click **Save**.

   The job name is auto generated and is constructed of the resource type followed by the SLA policy that is used for the job. For this example job, the name vmware_Copper is created.

5. Optional: To configure additional options, click **Select Options** and follow the instructions in "Backing up VMware data" on page 340.
6. Click **Save**.

After the job definition is saved, available virtual machine disks (VMDKs) in a virtual machine are discovered and are shown when **VMs and Templates** is selected in the **View** filter. By default, these VMDKs are assigned to the same SLA policy as the virtual machine. Optionally, to define a more granular policy by excluding individual VMDKs, follow the instructions in "Excluding VMDKs from the SLA policy for a job" on page 344.

Results

The job runs as defined by the SLA policies that you selected, or you can manually run the job by following the steps in "Start a backup job" on page 268.

Related concepts

"Protecting IBM Spectrum Protect Plus" on page 615

Protect the IBM Spectrum Protect Plus application by backing up the catalog. The catalog consists of data such as application configuration settings, SLA policies, registered resources, restore points, backup storage settings, and job information.

Back up the IBM Spectrum Protect Plus catalog

In addition to backing up the resources that you want to protect, you must also back up the IBM Spectrum Protect Plus catalog. In case of a natural disaster or other unexpected event, the catalog backup is required to restore system operations.

Procedure

The catalog consists of data such as application configuration settings, SLA policies, backup storage settings, and information about registered resources, restore points, and jobs.

1. In the navigation panel, click **Manage Protection** > **IBM Spectrum Protect Plus** > **Backup**.
2. Select the SLA policy that you created for the catalog backup job in "Create backup policies" on page 259.
3. Click **Save** to create the job definition.

Results

The job runs as defined by the SLA policy that you selected, or you can manually run the job by following the steps in "Start a backup job" on page 268.

Start a backup job

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

Procedure

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

1. In the navigation panel, click **Jobs and Operations**, and open the **Schedule** tab.
   If your job is not a scheduled job, but is an on-demand job, click the **Job History** tab.
2. Choose the job that you want to run and click the **Start** action as shown in the following figure:

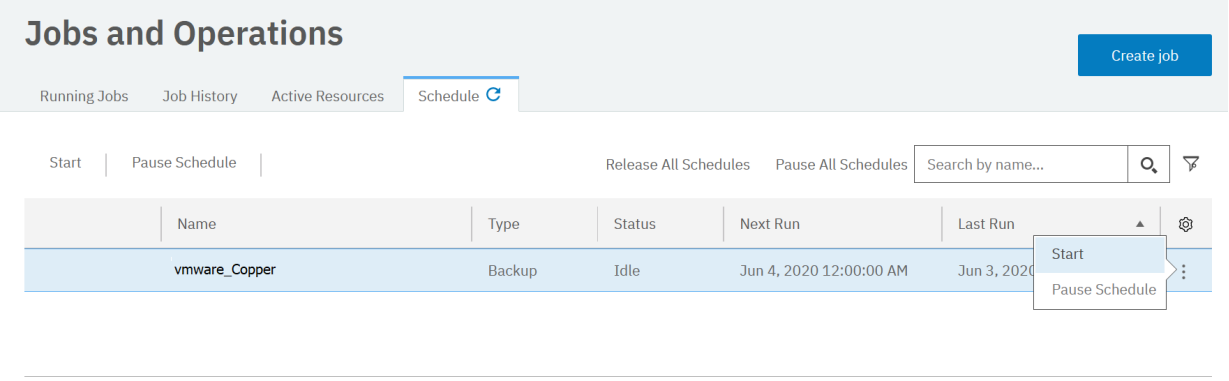


Figure 20. Starting a job

What to do next

To view the job log, click the job in the **Running Jobs** tab.

The log screen shows the following details:

- Status: shows whether the message is an error, warning, or informational message.
- Time: shows the time stamp of the message.
- ID: shows the unique identifier for the message if applicable.
- Description: shows what the message is.

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

Related concepts

"Managing jobs and operations" on page 619
You can manage and monitor jobs in the **Jobs and Operations** window. You can also configure scripts to run before or after jobs.

Run a report

Run reports with default parameters or custom parameters.

Procedure

To run a report, complete the following steps:

1. In the navigation panel, click **Reports and Logs** > **Reports**.
2. Click the **Reports** tab.

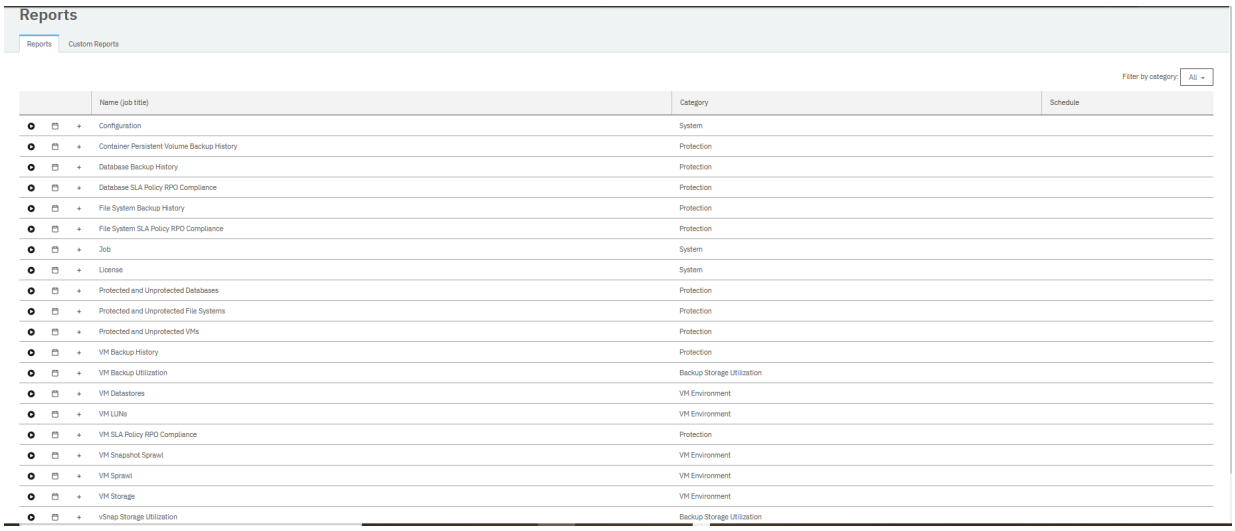


Figure 21. Selecting a report to run

3. Run the report by clicking the **Run Report** (▶) icon beside the report.
   - To run the report with custom parameters, set the parameters in the **Run Report** window, and click **Run**. Parameters are unique to each report.
   - To run the report with default parameters, click **Run**.

Related concepts

“Managing reports and logs” on page 631
IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Chapter 9. Configuring the system environment

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

Maintenance tasks include reviewing the configuration of the IBM Spectrum Protect Plus virtual appliance, collecting log files for troubleshooting, and managing Transport Layer Security (TLS) certificates.

All system management tasks apply to IBM Spectrum Protect Plus whether it is installed as a set of containers or as a virtual appliance, except for the tasks in "Configuring IBM Spectrum Protect Plus installed as a virtual appliance" on page 312. The tasks in this section apply only to IBM Spectrum Protect Plus virtual appliance installations.

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. To add keys and certificates in advance, complete the steps in the following topics:

- "Adding an access key" on page 297
- "Adding a certificate" on page 298

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 15
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.

© Copyright IBM Corp. 2017, 2021

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: • "Adding Amazon S3 Object Storage" on page 272 • "Adding IBM Cloud Object Storage as a backup storage provider" on page 273 • "Adding Microsoft Azure cloud storage as a backup storage provider" on page 275 • "Adding S3 compatible object storage" on page 276

Adding Amazon S3 Object Storage

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

Before you begin

Configure the key that is required for the cloud object. For instructions, see "Adding an access key" on page 297.

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

Procedure

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

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

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

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

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

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

   Key name
   Enter a meaningful name to help to identify the key.

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

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

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

5. Click **Get Buckets** to connect IBM Spectrum Protect Plus to AWS to retrieve the list of available buckets.
6. Select the bucket that you plan to use as the copy target.

   The **Standard object storage bucket** and **Archive object storage bucket** fields are displayed.
7. In the **Standard object storage bucket** field, select a bucket to serve as the copy target.
8. Optional: In the **Archive object storage bucket** field, select a cloud storage resource to serve as the archive target.

   Archiving data creates a full data copy and can provide longer-term protection, cost, and security benefits.

   For more information about archiving data, see the information about copying data to cloud archive storage in "Copy snapshots to secondary backup storage" on page 15.
9. Select **Deep Archive** to register Amazon S3 Glacier Deep Archive Buckets for long-term archiving.
10. Click **Register** to complete the operation.

    The cloud storage is added to the cloud servers table.

What to do next

After you add the S3 storage, complete the following action:

Action	How to
Associate the cloud storage with the SLA policy that is used for the backup job.	To create an SLA policy, see "Creating an SLA policy for hypervisors, databases, and file systems" on page 324. To modify an existing SLA policy, see "Editing an SLA policy" on page 333.

Adding IBM Cloud Object Storage as a backup storage provider

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

Before you begin

Configure the key and certificate that are required for the cloud object. For instructions, see "Adding an access key" on page 297 and "Adding a certificate" on page 298.

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

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

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

For on-premises IBM Cloud Object Storage (COS), the user associated with the access key and secret key that is registered in IBM Spectrum Protect Plus must be assigned as an Owner of the vault in the IBM COS Manager interface. It is not sufficient for the user to have only Read/Write access to the vault.

When adding IBM Cloud Object Storage (COS), the method for obtaining the access and secret key will depend on the deployment model. If on-premise, keys can be obtained from the IBM COS Manager Console. For IBM COS IaaS, keys are created when a service account is created and can be obtained from the softlayer portal. If using IBM COS (COS as a Service), the access and secret key are not

created by default; when a service account is created, check the **Include HMAC Credential** box, and add `{"HMAC":true}` to the **Add Inline Configuration Parameters** text area.

Procedure

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

1. In the navigation 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 15.
8. Click **Register**.

   The cloud storage is added to the cloud servers table.

What to do next

After you add the IBM Cloud Object Storage, complete the following action:

Action	How to
Associate the cloud storage with the SLA policy that is used for the backup job.	To create an SLA policy, see "Creating an SLA policy for hypervisors, databases, and file systems" on page 324. To modify an existing SLA policy, see "Editing an SLA policy" on page 333.

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 panel, 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 15.
8. Click **Register**.

The cloud storage is added to the cloud servers table.

What to do next

After you add the Microsoft Azure storage, complete the following action:

Action	How to
Associate the cloud storage with the SLA policy that is used for the backup job.	To create an SLA policy, see "Creating an SLA policy for hypervisors, databases, and file systems" on page 324. To modify an existing SLA policy, see "Editing an SLA policy" on page 333.

Adding S3 compatible object storage

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

Before you begin

Tip:

For information about compatible object storage providers, see technote 108714.

Configure the key that is required for the cloud object. For instructions, see "Adding an access key" on page 297.

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 15.
8. Click **Register**.
   The cloud storage is added to the cloud servers table.

What to do next

After you add the S3 compatible storage, complete the following action:

Action	How to
Associate the cloud storage with the SLA policy that is used for the backup job.	To create an SLA policy, see "Creating an SLA policy for hypervisors, databases, and file systems" on page 324. To modify an existing SLA policy, see "Editing an SLA policy" on page 333.

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.

User scenario	Purpose	Steps
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. Some of the data may be duplicated on the tape storage pool and cache storage pool.	To copy data to tape, you must create a tape storage pool first and then you must create a disk storage pool which is where the cold-data-cache storage pool will reside on the IBM Spectrum Protect server. Adding the object agent is a mandatory step. Follow steps 1-4 listed, here.
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 280. 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 281.
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 283.
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 286.
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 287.

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 289.
2. When that is done, you can create an SLA policy that defines the IBM Spectrum Protect server as the backup storage target. For more information to help you choose which type of policy you need, see “Configuration for copying or archiving data to IBM Spectrum Protect” on page 278

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

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

About this task

Before you start, ensure that you have sized your cold cache storage needs by using the sizing tool and the Blueprints. For information about how to do this, see the Blueprints. For more useful links and videos, see “Deployment storyboard for IBM Spectrum Protect Plus” on page 1.

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

Procedure

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

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

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

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

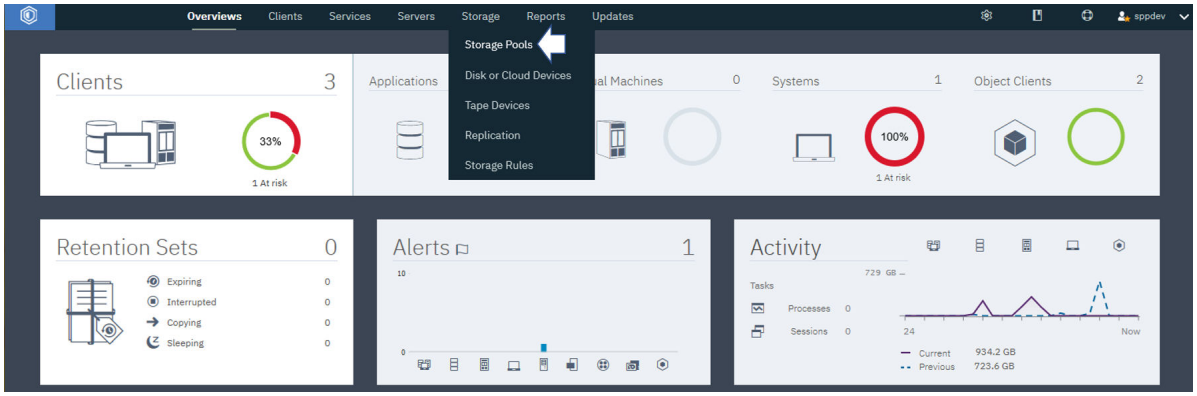
1. On the IBM Spectrum Protect server, configure a primary storage pool that represents a tape device or VTL. This primary storage pool is the destination for the object data that you want to copy.

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

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

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

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



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

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

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

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

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

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

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

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

What to do next

When you complete the configuration of the cold data cache storage pool, create the object domain. For instructions about how to do that, see “Step 2: Configuring an object policy domain” on page 281.

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 283.

Step 3: Setting up standard object storage

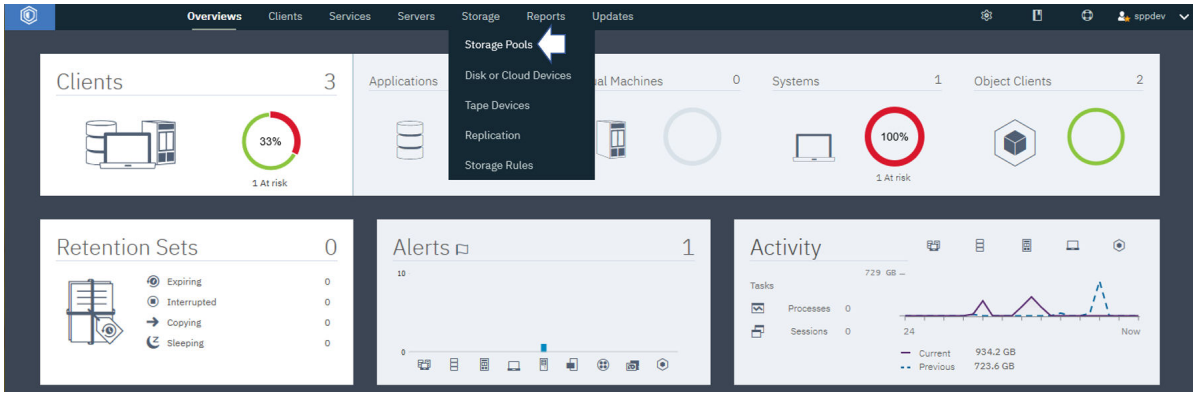
To set up standard object storage for copying data from IBM Spectrum Protect Plus to the IBM Spectrum Protect server, log in to the Operations Center and follow the procedure to set up storage pools. Complete the process by following the steps to create an object agent service by using the Operations 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 280.

Procedure

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



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

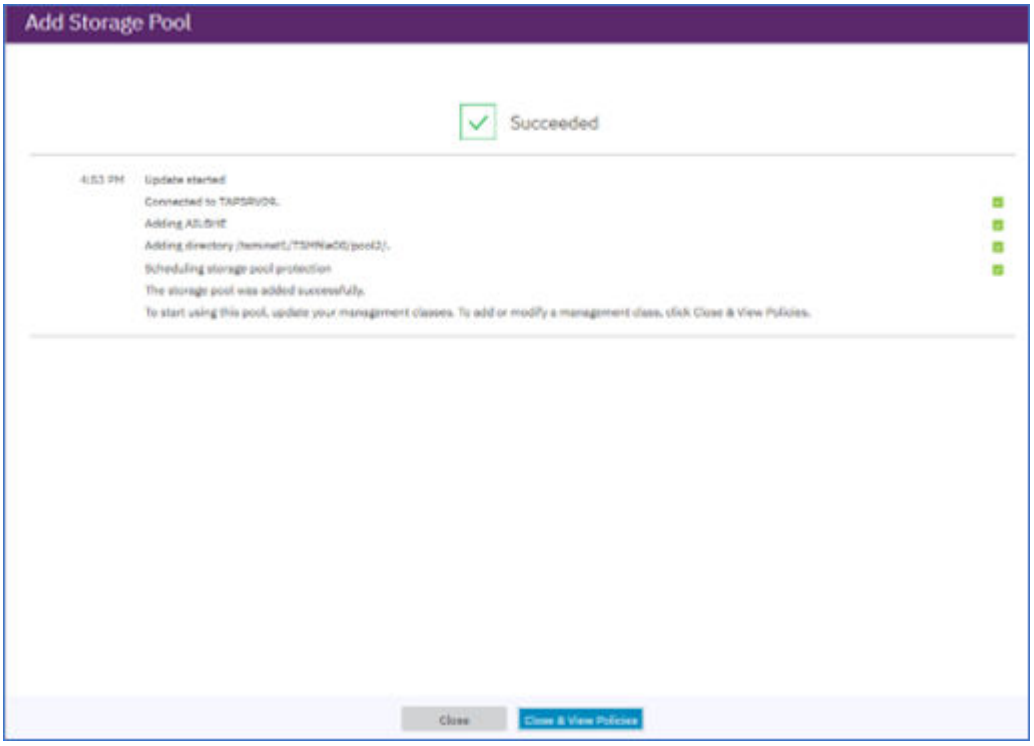
c) Complete the steps in the **Add Storage Pool** wizard.

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

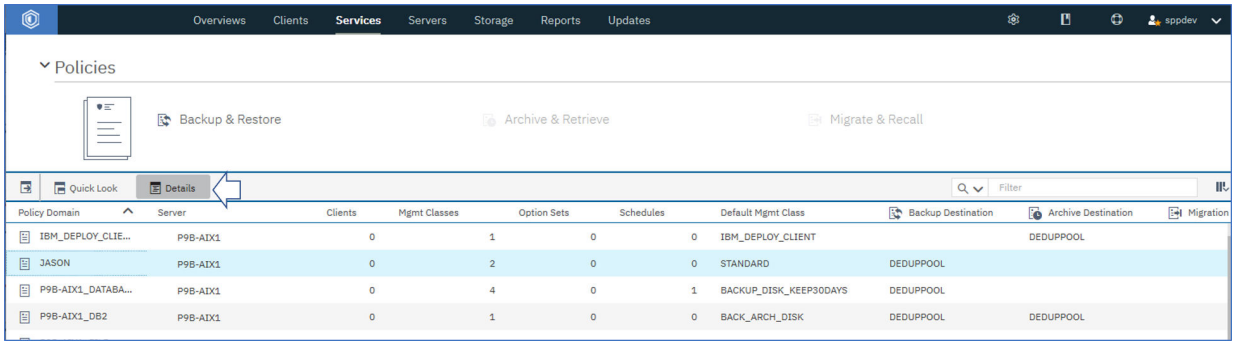
d) Review the **Protect Pool** summary, and click **Next**.

e) Specify an overflow pool is that is required.

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

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

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



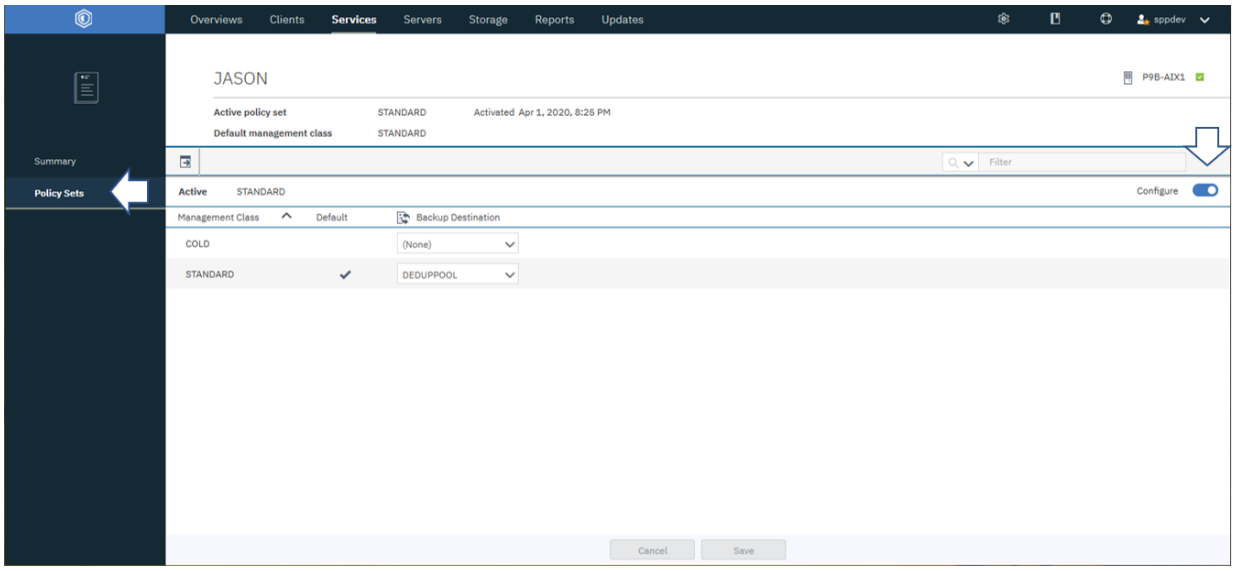
- You can edit an existing domain policy by following these steps:

a) Update one or more management classes to use the new pool by editing the **Backup Destination** field of the table.

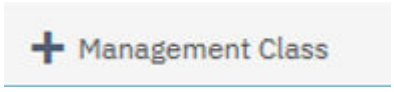
b) Click **Save**.

- 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 281.

3. On the **Details** page, click **Policy Sets**. Click the **Configure** toggle to make the policy sets editable.



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



to point to the new storage pool.

5. 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.

6. 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.
7. 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.
8. 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 287

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.

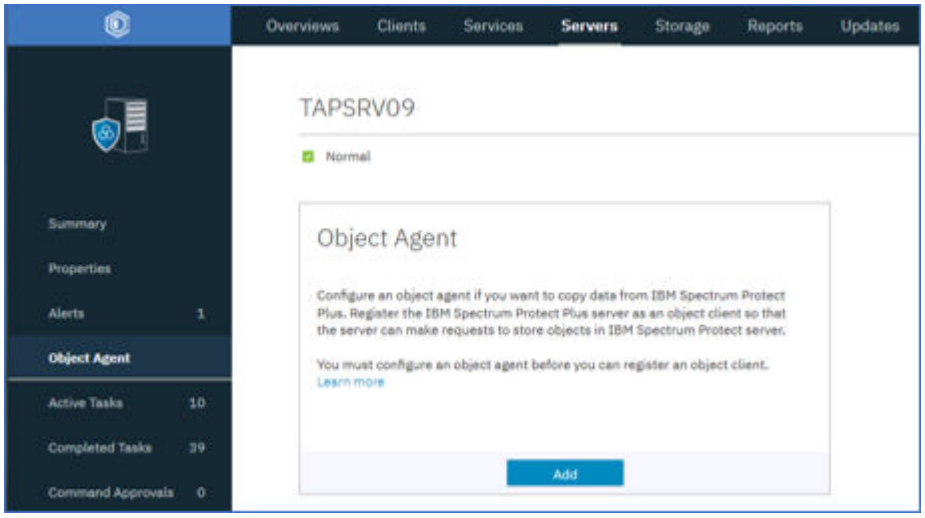
1. Ensure that you are logged in to the IBM Spectrum Protect server with an instance user ID.
2. Ensure that you have set up storage pools either for standard storage or for copying to tape. For instructions, see “Step 1: Creating a tape storage pool and a cold-data-cache storage pool for copying data to tape” on page 280 or “Step 3: Setting up standard object storage” on page 283.
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)

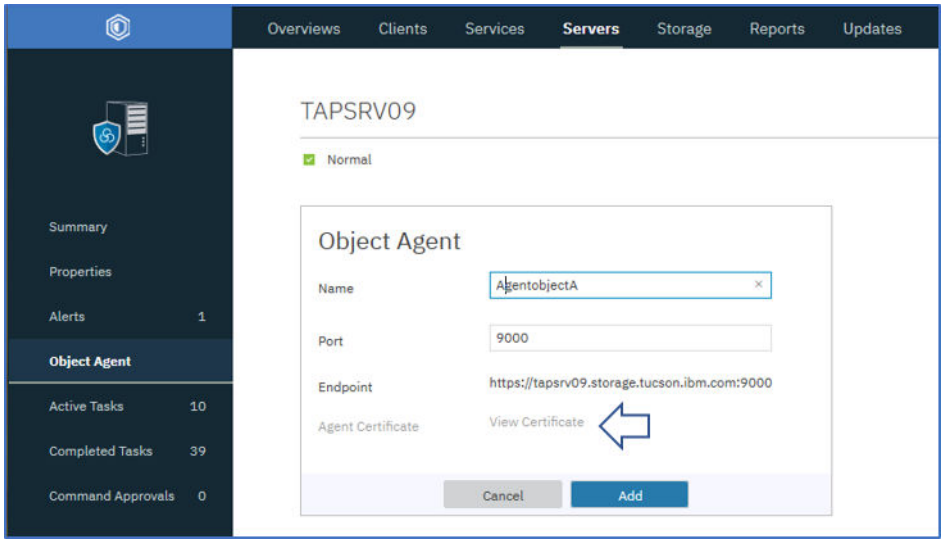
Procedure

1. On the Operations Center menu bar, click **Servers** Servers.
2. Select a server and click **Details**.
3. From the navigation panel, 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 280 or “Step 3: Setting up standard object storage” on page 283.
3. Ensure that an object domain and an object agent are created before you start.

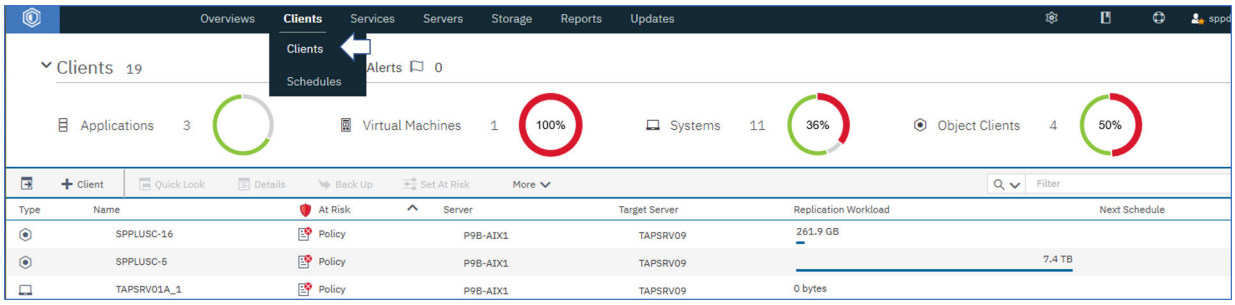
Tip: If you create an object client before you create the corresponding object agent, the **Add Client** wizard forces the creation of the object agent.

About this task

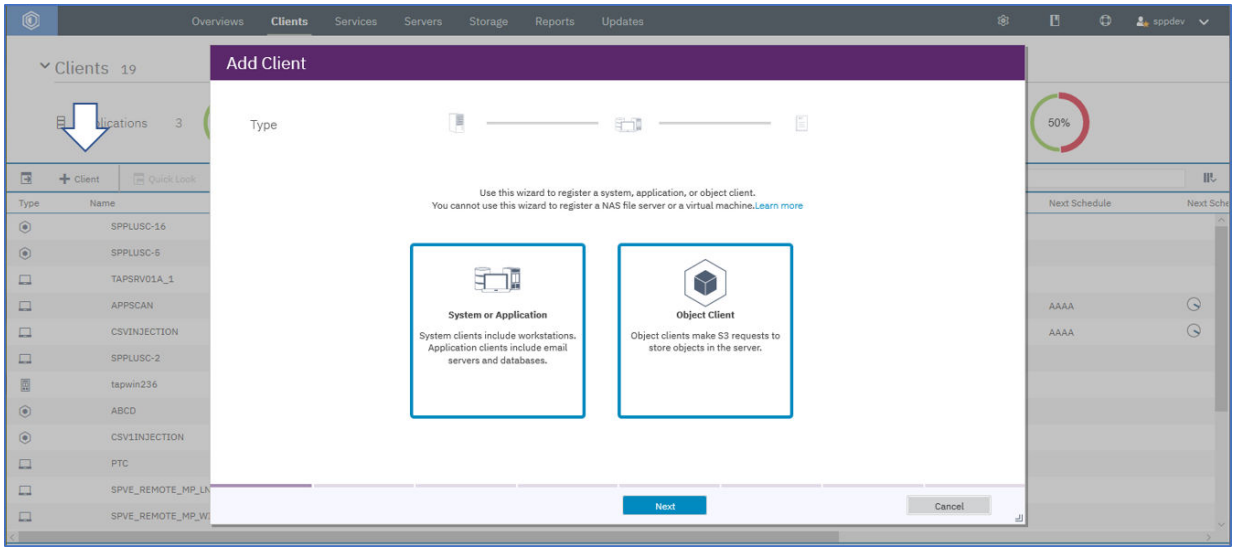
This procedure is based on an environment where the IBM Spectrum Protect server is installed on an IBM AIX operating system AIX Version 7.2 TL 1 and SP 4 or later, running on an IBM POWER8 or later server.

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 281.
   - 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 289. Once that is completed, you can create SLA policy jobs to copy data to the IBM Spectrum Protect server for standard storage or for archive to tape.

Registering a repository server as a backup storage provider

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

Before you begin

Configure the key and certificate that are required for the repository server. For instructions, see "Adding an access key" on page 297 and "Adding a certificate" on page 298.

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.

What to do next

After you add a repository server, complete the following action:

Action	How to
Associate the repository server with the SLA policy that is used for the backup job.	To create an SLA policy, see "Creating an SLA policy for hypervisors, databases, and file systems" on page 324. To modify an existing SLA policy, see "Editing an SLA policy" on page 333.

Related concepts

"Configuration for copying or archiving data to IBM Spectrum Protect" on page 278
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.

Managing sites

A *site* is an IBM Spectrum Protect Plus policy construct that is used to manage the placement of data in an environment.

A site can be physical, such as a data center, or logical, such as a department or organization. IBM Spectrum Protect Plus components are assigned to sites to localize and optimize data paths. An IBM Spectrum Protect Plus deployment always has at least one site per physical location.

By default, the IBM Spectrum Protect Plus environment has a Primary site and a Secondary site.

Adding a site

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

Procedure

To add a site, complete the following steps:

1. In the navigation panel, click **System Configuration** > **Site**.
2. Click **Add Site**.

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

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

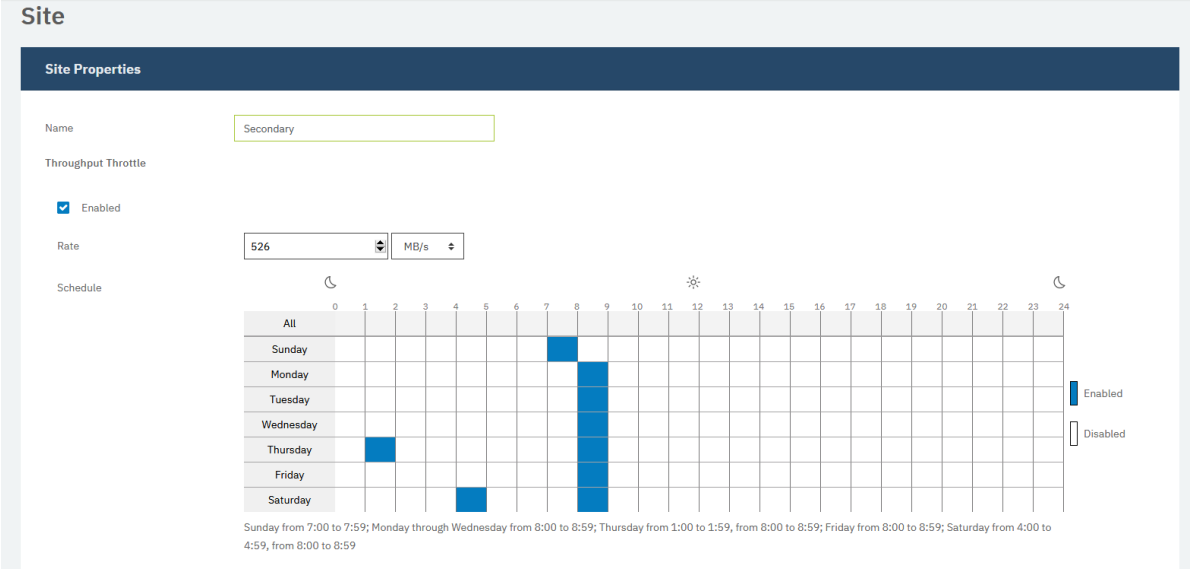


Figure 22. Enabling different rates of throttling for different times to improve throughput

c) In the weekly schedule table, select daily times for throttling, or select specific days and times for throttling. The time that is specified should be based on the local time of the one or more vSnap servers that are assigned to the site.

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

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

5. Optional: To optimize the storage on the vSnap servers in your environment for the back up of new VMware data, use the slider in the **VMware VM allocation** section. This slider determines how new VMware backup data is initially allocated to the vSnap servers: by free space, by VM count, or a factor of both. The slider is set to allocate backup data by VM count by default. Do not change this value unless instructed by IBM Support.
6. Click **Save** to commit the changes and close the pane.

Results

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

Editing a site

Revise site information to reflect changes in your IBM Spectrum Protect Plus environment.

Procedure

To edit a site, complete the following steps:

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

i) Change the numerical rate of throughput by clicking the up or down arrows.

ii) Select a unit for the throughput. The choices include **bytes/s**, **KB/s**, **MB/s**, and **GB/s**.

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

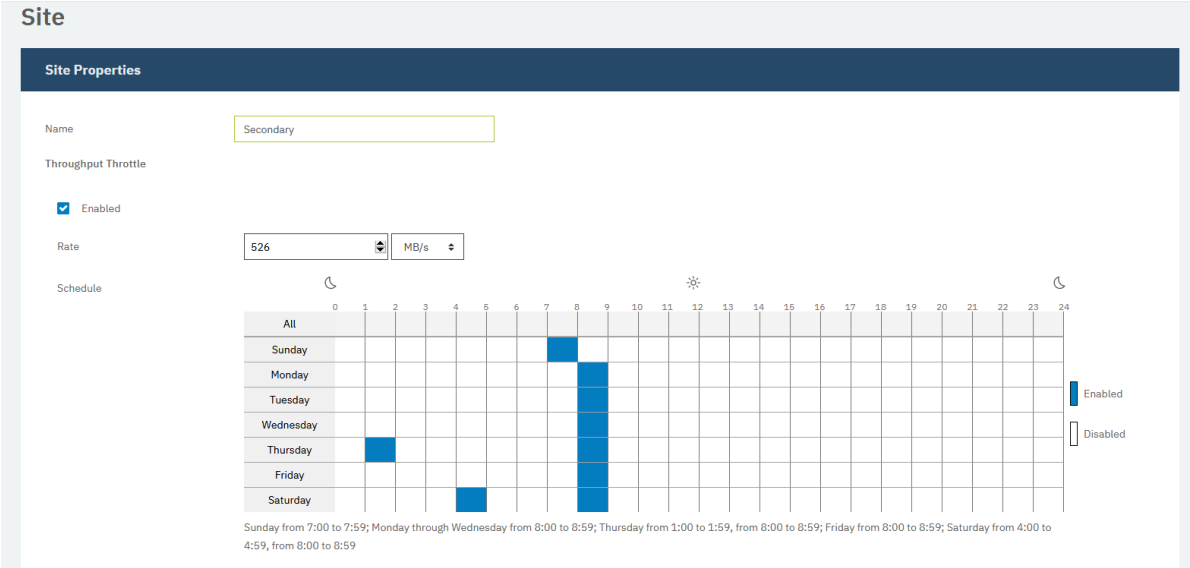


Figure 23. Enabling different rates of throttling for different times to improve throughput

c) In the weekly schedule table, select daily times for throttling, or select specific days and times for throttling. The time that is specified should be based on the local time of the one or more vSnap servers that are assigned to the site.

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

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

5. Optional: To optimize the storage on the vSnap servers in your environment for the back up of new VMware data, use the slider in the **VMware VM allocation** section. This slider determines how new VMware backup data is initially allocated to the vSnap servers: by free space, by VM count, or a factor of both. The slider is set to allocate backup data by VM count by default. Do not change this value unless instructed by IBM Support.
6. Click **Save** to commit the changes and close the pane.

Deleting a site

Delete a site when it becomes obsolete. Ensure that you reassign your backup storage to different sites before deleting the site.

Procedure

To delete a site, complete the following steps:

1. In the navigation panel, click **System Configuration** > **Site**.
2. Click the delete icon ✕ that is associated with a site.
3. Click **Yes** to delete the site.

Managing LDAP and SMTP servers

You can add a Lightweight Directory Access Protocol (LDAP) and Simple Mail Transfer Protocol (SMTP) server for use in the IBM Spectrum Protect Plus for use in user account and report features.

Related tasks

"Creating a user account for an LDAP group" on page 653
With IBM Spectrum Protect Plus, you can use a Lightweight Directory Access Protocol (LDAP) server to manage users. When you create an LDAP user account, you can add the user account to a user group.

"Scheduling a report" on page 640
You can schedule reports in IBM Spectrum Protect Plus to run at specific times.

Adding an LDAP server

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

About this task

You can add a Microsoft Active Directory or OpenLDAP server. Note that OpenLDAP does not support the sAMAaccountName user filter that is commonly used with Active Directory. Additionally, the **memberOf** option must be enabled on the OpenLDAP server.

Procedure

To register an LDAP server, complete the following steps:

1. In the navigation panel, click **System Configuration** > **LDAP/SMTP Server**.
2. In the **LDAP Servers** pane, click **Add LDAP Server**.
3. Populate the following fields in the **LDAP Servers** pane:

Host Address

The IP address of the host or logical name of the LDAP server.

Port

The port on which the LDAP server is listening. The typical default port is 389 for non TLS connections or 636 for TLS connections.

TLS

Enable the TLS option to establish a secure connection to the LDAP server.

Use existing user

Enable to select a previously entered user name and password for the LDAP server.

Bind Name

The bind distinguished name that is used for authenticating the connection to the LDAP server. IBM Spectrum Protect Plus supports simple bind.

Password

The password that is associated with the Bind Distinguished Name.

Base DN

The location where users and groups can be found.

User Filter

A filter to select only those users in the Base DN that match certain criteria. An example of a valid default user filter is `cn={0}`.

Tips:

- To enable authentication by using the **sAMAccountName** Windows user naming attribute, set the filter to `samaccountname={0}`. When this filter is set, users log in to IBM Spectrum Protect Plus by using only a user name. A domain is not included.
- To enable authentication using the user principal name (UPN) naming attribute, set the filter to `userprincipalname={0}`. When this filter is set, users log in to IBM Spectrum Protect Plus by using the username@domain format.
- To enable authentication by using an email address that is associated with LDAP, set the filter to `mail={0}`.

The **User Filter** setting also controls the type of user name that appears in the IBM Spectrum Protect Plus display of users.

User RDN

The relative distinguished path for the user. Specify the path where user records can be found. An example of a valid default RDN is `cn=Users`.

Group RDN

The relative distinguished path for the group. If the group is at a different level than the user path, specify the path where group records can be found.

4. Click **Save**.

Results

IBM Spectrum Protect Plus completes the following actions:

1. Confirms that a network connection is made.
2. Adds the LDAP server to the database.

After the SMTP server is added, the **Add LDAP Server** button is no longer available.

What to do next

If a message is returned indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connections.

Related tasks

"Creating a user account for an LDAP group" on page 653
With IBM Spectrum Protect Plus, you can use a Lightweight Directory Access Protocol (LDAP) server to manage users. When you create an LDAP user account, you can add the user account to a user group.

Adding an SMTP server

You must add an SMTP server to send scheduled reports to email recipients. Only one SMTP server can be associated with a IBM Spectrum Protect Plus virtual appliance.

Procedure

To add an SMTP server, complete the following steps:

1. In the navigation panel, click **System Configuration** > **LDAP/SMTP Server**.
2. In the **SMTP Servers** pane, click **Add SMTP Server**.
3. Populate the following fields in the **SMTP Servers** pane:

Host Address

The IP address of the host, or the path and host name of the SMTP server.

Port

The communications port of the server that you are adding. The typical default port is 25 for non-TLS connections or 443 for TLS connections.

Username

The name that is used to access the SMTP server.

Password

The password that is associated with the user name.

Timeout

The email timeout value in milliseconds.

From Address

The address that is associated with email communications from IBM Spectrum Protect Plus.

Subject Prefix

The prefix to add to the email subject lines sent from IBM Spectrum Protect Plus.

4. Click **Save**.

Results

IBM Spectrum Protect Plus completes the following actions:

1. Confirms that a network connection is made.
2. Adds the server to the database.

If a message is returned indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connections.

To test the SMTP connection, click the **Test SMTP Server** button, then enter an e-mail address. Click **Send**. A test e-mail message is sent to the e-mail address to verify the connection.

After the SMTP server is added, the **Add SMTP Server** button is no longer available.

What to do next

Related tasks

"Scheduling a report" on page 640

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

Editing settings for an LDAP or SMTP server

Edit the settings for an LDAP or SMTP server to reflect changes in your IBM Spectrum Protect Plus environment.

Procedure

To edit the settings for an LDAP or SMTP server, complete the following steps:

1. From the navigation menu, click **System Configuration** > **LDAP/SMTP Server**.
2. Click the edit icon that is associated with the server.
   The edit pane is displayed.
3. Revise the settings for the server, and then click **Save**.

Deleting an LDAP or SMTP server

Delete an LDAP or SMTP server when it becomes obsolete. Ensure that the server is not in use by IBM Spectrum Protect Plus before deleting the server.

Procedure

To delete an LDAP or SMTP server, complete the following steps:

1. From the navigation menu, click **System Configuration** > **LDAP/SMTP Server**.
2. Click the delete icon ✕ that is associated with the server.
3. Click **Yes** to delete server.

Managing keys and certificates for connection to IBM Spectrum Protect Plus components

Keys and certificates are used in the IBM Spectrum Protect Plus environment to provide secure connections to IBM Spectrum Protect Plus components.

Keys, and in some environments certificates, are required to enable IBM Spectrum Protect Plus to connect to the following components:

Secondary backup storage
The cloud resources and repository servers that provide secondary backup storage require credentials to serve as copy destinations. Access keys and secret keys are provided by your cloud resource or repository server interface. These keys serve as the username and password of your copy destinations and allow them to be accessed by IBM Spectrum Protect Plus. Some copy destinations also require TLS certificates for additional data security. The TLS certificate can be a certificate that is issued by a certificate authority (CA).

Linux-based resources
You can add a Secure Shell (SSH) key to provide credentials for Linux-based resources on virtual machines managed by vCenter and Hyper-V, as well as Oracle, Db2, and MongoDB application servers. SSH keys help to provide a secure connection between IBM Spectrum Protect Plus and target resources for file indexing and restore operations.

When you add a key or certificate to IBM Spectrum Protect Plus, the list of available keys and certificates is updated so that you can select a key or certificate as needed in the user interface.

For information about using an TLS certificate for secure connections to the IBM Spectrum Protect Plus user interface, see “Uploading a TLS certificate” on page 302.

Adding an access key

Add an access key to provide cloud resource or repository server credentials.

Procedure

To add a key, complete the following steps:

1. Create your access key and secret key through the interface of the cloud resource or repository server. Make note of the access key and secret key.
2. In the navigation menu, click **System Configuration** > **Keys and Certificates**.
3. From the **Access Keys** section, click **Add Access Key**.
4. Complete the fields in the **Key Properties** pane:

   Name

   Enter a meaningful name to help identify the access key.

   Access Key

Enter the access key of the cloud resource or repository server. For Microsoft Azure, enter the storage account name.

Secret Key

Enter the secret key of the cloud resource or repository server. For Microsoft Azure, enter the key from one of the key fields, either key1 or key2.

5. Click **Save**.

   The key displays in the **Access Keys** table and can be selected when utilizing a feature that requires credentials to access a resource through the **Use existing key** option.

Deleting an access key

Delete an access key when it becomes obsolete. Ensure that you reassign a new access key to your cloud resource or repository server.

Procedure

To delete an access key, complete the following steps:

1. In the navigation menu, click **System Configuration** > **Keys and Certificates**.
2. Click the delete icon ✕ that is associated with an access key.
3. Click **Yes** to delete the access key.

Adding a certificate

Add a certificate to provide cloud resource or repository server credentials.

Procedure

To add a certificate, complete the following steps:

1. Export a certificate from your cloud resource or repository server.
2. In the navigation menu, click **System Configuration** > **Keys and Certificates**.
3. In the **Certificates** section, click **Add Certificate**.
4. Complete the fields in the **Certificate Properties** pane:

   Type

   Select the cloud resource or repository server type.

   Certificate

   Select a method to add the certificate:

   Upload
   Select to browse for the certificate locally.

   Copy and paste
   Select to enter the name of the certificate and copy and paste the contents of the certificate.

5. Click **Save**.

   The key displays in the **Certificates** table and can be selected when utilizing a feature that requires credentials to access a resource through the **Use existing certificate** option.

Deleting a certificate

Delete a certificate when it becomes obsolete. Ensure that you reassign a new certificate to your cloud resource or repository server.

Procedure

To delete a certificate, complete the following steps:

1. In the navigation menu, click **System Configuration** > **Keys and Certificates**.
2. Click the delete icon ✕ that is associated with a certificate.
3. Click **Yes** to delete the certificate.

Adding an SSH key

You can add an SSH key to provide credentials for Linux-based and AIX-based resources on virtual machines managed by vCenter and Hyper-V, as well as Oracle, Db2, and MongoDB application servers. SSH keys help to provide a secure connection between IBM Spectrum Protect Plus and target resources for file indexing and restore operations.

Before you begin

- The SSH service must be running on port 22 on the server and any firewalls must be configured to allow IBM Spectrum Protect Plus to connect to the server using SSH. The SFTP subsystem for SSH must also be enabled.
- The user account on the target resource that is used to generate the SSH key pair must have **sudo** privileges. This account, which will be assigned to IBM Spectrum Protect Plus, is known as the IBM Spectrum Protect Plus user agent (`sppagent`).
- If the environment includes virtual machines managed by vCenter, ensure that the latest VMware Tools are installed.

Procedure

To add a key, complete the following steps:

1. On the target resource, generate an SSH key by using the `ssh-keygen` command with the user account that will be assigned to IBM Spectrum Protect Plus. This account must have **sudo** privileges. For example, on an Oracle server, enter the following command in the terminal and follow the instructions:

   ```
   ssh-keygen
   ```

   If you use the default settings, two files are created in the specified directory: `id_rsa.pub` is the public key and `id_rsa` is the private key. The private key must be in PEM format. It may be necessary to explicitly use the `-m PEM` argument with `ssh-keygen` when generating the key pair.
2. When prompted enter the file name in which the key will be saved, enter a directory and file name. If you do not specify a directory and file name, the default is used:

   ```
   /home/priveleged_user/.ssh/id_rsa
   ```

   where *priveleged_user* is the account assigned to IBM Spectrum Protect Plus, `sppagent`. If a key with the default name already exists, this will be indicated with the message displayed below. Be careful not to overwrite preexisting keys if they are in use. Press **N** to enter a different file in which to save the key.

   ```
   /home/<priveleged user>/.ssh/id_rsa already exists.
   Overwrite (y/n)?
   ```

This procedure is based on the assumption that the key is saved in the default location using the default file name (`id_rsa`). If the key file is created using a different file name, use that file name in the steps that follow.

3. Supply a passphrase and press Enter. Otherwise, simply press Enter for no passphrase.
4. If a passphrase was supplied, enter it again. Press Enter.
5. Copy the contents of the `id_rsa.pub` key into the `authorized_keys` file. If the file already exists, append the public key to the `authorized_keys` file.

```
cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys
```

6. Assign the required privileges to the `authorized_keys` file by issuing the `chmod 600` command.

```
chmod 600 ~/.ssh/authorized_keys
```

7. Edit the `/etc/ssh/sshd_config` file to set the `PubkeyAuthentication` setting to `yes` by using a text editor. To ensure that the setting is not commented out, remove the number sign (#) if it appears at the beginning of the line.

```
sudo vi /etc/ssh/sshd_config
```

```
...
PubkeyAuthentication yes
...
```

8. Restart the SSH service on the target resource.

```
systemctl restart sshd
```

9. In the IBM Spectrum Protect Plus navigation panel, click **System Configuration** > **Keys and Certificates**.
10. From the **SSH Keys** section, click **Add SSH Key**.
11. Complete the fields in the **SSH Key Properties** pane:

    Name

    Enter a meaningful name to identify the SSH key.

    User

    Enter the user account that is associated with the target resource and SSH key. This is the user account used to generate the public and private keys in the previous steps.

    Encrypted

    Check this box if a passphrase was supplied when generating the public and private key.

    Passphrase

    This box is only displayed if the **Encrypted** check box is selected. If a passphrase was supplied when generating the public and private key, provide the passphrase in this box.

    Private key

    Copy and paste the private key into this box. This will be the key contained in the `id_rsa` file on the target resource. The file is similar to the following example:

```
cat ~/.ssh/id_rsa
```

```
-----BEGIN OPENSSH PRIVATE KEY-----
ZRYtuinjaHx2mKgW4LnfqzlyAIIq5Amasi/J8/AAAFiFiP4GZYj+BmAAAAB3NzaC1yc2
...
...
Q5ZqZ1Ec8N7dsAAAANdG9vckBVYnVudHVWQgECAwQFBg==
-----END OPENSSH PRIVATE KEY-----
```

12. Click **Save**.

The key is displayed in the **SSH Keys** table and can be selected when you use a feature that requires credentials to access a resource with the **Key** option.

Deleting an SSH key

Delete an SSH key when it becomes obsolete. Ensure that you reassign a new SSH key to your resources.

Procedure

To delete an SSH key, complete the following steps:

1. In the navigation menu, click **System Configuration** > **Keys and Certificates**.
2. Click the delete icon ✕ that is associated with an SSH key.
3. Click **Yes** to delete the access key.

Managing certificates for connection to the IBM Spectrum Protect Plus user interface

To establish secure connections to IBM Spectrum Protect Plus, you can upload the following TLS certificates depending on your environment and requirements: Hypertext Transfer Protocol Secure (HTTPS) and Lightweight Directory Access Protocol (LDAP).

An HTTPS certificate authority (CA) is required to establish a trusted connection to the IBM Spectrum Protect Plus user interface. You can start the IBM Spectrum Protect Plus user interface with the default self-signed certificate, but you will receive a browser notification that the certificate is not trusted.

An LDAP certificate is required if you are using LDAP authentication for IBM Spectrum Protect Plus users.

The following technotes provide introductory information for using certificates with IBM Spectrum Protect Plus:

HTTPS

Technote 739663 provides information about using a HTTPS certificate that is issued by Microsoft Certificate Authority. However, you can use a certificate that is issued by another certificate authority (CA).

For HTTPS certificates, PEM encoded certificates with `.cer` or `.crt` extensions are supported.

LDAP

Technote 791677 provides information about using an LDAP certificate.

For LDAP certificates, DER encoded certificates with `.cer` or `.crt` extensions are supported. If you are uploading an LDAP TLS certificate, ensure that IBM Spectrum Protect Plus has connectivity to the LDAP server and that the LDAP server is running.

CA certificate format example

The following example shows the format of a PEM encoded CA certificate. This example file contains a private key, a CA server certificate, one intermediate certificate, and a root certificate. The values in the private key and certificates are for example purposes.

```
# Private key
-----BEGIN PRIVATE KEY-----
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
ABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
ABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwx
-----END PRIVATE KEY-----

# CA server certificate
-----BEGIN CERTIFICATE-----
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
ABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
```

```
ABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
ABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghi
-----END CERTIFICATE-----

# Intermediate certificate
-----BEGIN CERTIFICATE-----
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
ABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
ABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
ABCDEFGHIJKLMOPQRSTUVWXWXYZabcdefghijklm
-----END CERTIFICATE-----

# Root certificate
-----BEGIN CERTIFICATE-----
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
ABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
ABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyzABCDEFGHIJKLM
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMOPQRSTUVWXYZabcdefghijklm
ABCDEFGHIJKLMOPQRSTUVWXW
-----END CERTIFICATE-----
```

Uploading a TLS certificate

You can upload TLS certificates to establish secure connections to IBM Spectrum Protect Plus.

Before you begin

To upload a TLS certificate, you must log on to IBM Spectrum Protect Plus as the superuser. The IBM Spectrum Protect Plus superuser is the user who is assigned the SUPERUSER role.

Procedure

To upload a TLS certificate, complete the following steps:

1. In the IBM Spectrum Protect Plus user interface, click the user menu in the menu bar, and then click **Manage TLS certificates**.
2. Select the TLS certificate type: **HTTPS** or **LDAP**.
3. Click **Browse**, and select the certificate that you want to upload.
4. Click **Upload**.
5. Restart IBM Spectrum Protect Plus.

Testing network connectivity

The IBM Spectrum Protect Plus Service Tool tests host addresses and ports to determine if a connection can be established. You can use the Service Tool to verify whether a connection can be established between IBM Spectrum Protect Plus and a node.

You can run the Service Tool from the IBM Spectrum Protect Plus command line or remotely by using a .jar file. If a connection can be established, the tool returns a green check mark. If a connection cannot be established, the error condition is displayed, along with possible causes and actions.

The tool provides guidance for the following error conditions:

- Timeout
- Connection refused
- Unknown host
- No route

Running the Service Tool from a command line

You can start the Service Tool from the IBM Spectrum Protect Plus virtual appliance command line interface and run the tool in a web browser. Then, you can use the Service Tool to verify network connectivity between IBM Spectrum Protect Plus and a node.

Procedure

1. Log in to the IBM Spectrum Protect Plus virtual appliance by using the `serveradmin` user ID and access the command line. Run the following command:

   ```
   # sudo bash
   ```

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

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

3. Run the tool by running the following command:

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

4. To connect to the tool, enter the following URL in a browser:

   ```
   http://hostname:9000
   ```

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

5. To specify the node to test, complete the following fields:

   Host
   The hostname or IP address of the node that you want to test.

   Port
   The connection port to test.

6. Click **Save**.
7. To run the tool, hover the cursor over the tool, and then click **Run**.

   If a connection cannot be established, the error condition is displayed, along with possible causes and actions.

8. Stop the tool by running the following command on the command line:

   ```
   ctl-c
   ```

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

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

   Note: If the `firewall-cmd` command is not available on your system, edit the firewall manually to add necessary ports and restart the firewall with `iptables`. For more information on editing firewall rules, see the **Firewall configuration with `iptables`** section here: https://www.ibm.com/support/knowledgecenter/en/STXKQY_5.0.3/com.ibm.spectrum.scale.v5r03.doc/bl1adv_firewallportopenexamples.htm.

Running the Service Tool remotely

You can download the Service Tool as a .jar file from the IBM Spectrum Protect Plus user interface. Then, you can use the Service Tool to remotely test connectivity between IBM Spectrum Protect Plus and a node.

Procedure

1. In the IBM Spectrum Protect Plus user interface, click the user menu, and then click **Download Test Tool**.

   A .jar file is downloaded to your workstation.
2. Launch the tool from a command-line interface. Java is only required on the system where the tool will be launched. Endpoints or target systems that are tested by the tool do not require Java.

   The following command launches the tool in a Linux environment:

   ```
   # java -jar -Dserver.port=9000 /<tool path >/ngxdd.jar
   ```
3. To connect to the tool, enter the following URL in a browser:

   ```
   http://hostname:9000
   ```

   where *hostname* specifies the IP address of the virtual machine where the application is deployed.
4. To specify the node to test, populate the following fields:

   Host
   The host name or IP address of the node that you want to test.

   Port
   The connection port to test.
5. Click **Save**.
6. To run the tool, hover the cursor over the tool, and then click the green **Run** button.

   If a connection cannot be established, the error condition is displayed, along with possible causes and actions.
7. Stop the tool by issuing the following command on the command line:

   ```
   ctl-c
   ```

Configuring global preferences

As the administrator, you can configure preferences that apply to all IBM Spectrum Protect Plus operations in the **Global Preferences** pane.

Before you begin

You must have administrator credentials to configure global preferences.

You can change the preference in the **Integration with other storage products** and **User interface** categories at any time.

Attention: Although you can modify the preference in the **Integrations with other storage products** and **User interface** categories, modify all other preferences only if absolutely necessary and only at the direction of IBM Support. Modifying global preferences can affect your storage environment. Preferences that require consultation with IBM Support are in the following categories: **Application**, **General**, **Job**, **Logging**, **Protection**, and **Security**.

About this task

Any changes that you make to parameter default values apply to all IBM Spectrum Protect Plus operations when you save the changes.

The figures in this task are for example purposes only and some values might differ from the default.

Procedure

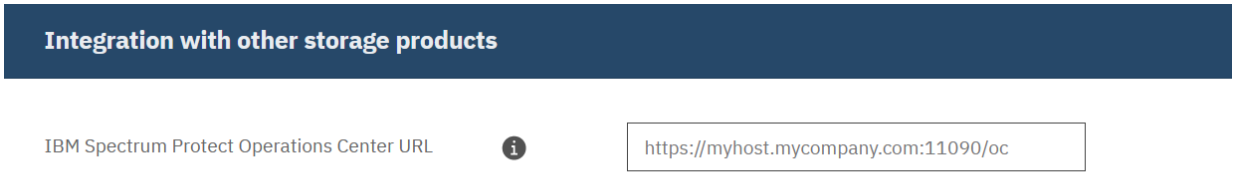
To edit the values for any setting and apply them globally, complete the following steps:

1. In the navigation panel, click **System Configuration** > **Global Preferences**.
2. To enable access to IBM Spectrum Protect Operations Center from IBM Spectrum Protect Plus, edit the **IBM Spectrum Protect Operations Center URL** preference in the **Integration with other storage products** category.

   The **IBM Spectrum Protect Operations Center URL** preference is the IP address of IBM Spectrum Protect Operations Center. The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect environment.

   When this preference is set, the IBM Spectrum Protect icon is active on the IBM Spectrum Protect Plus menu bar. When you initially set the URL for this preference or if you change it, you must log off and log back in for the preference to take effect in the user interface.

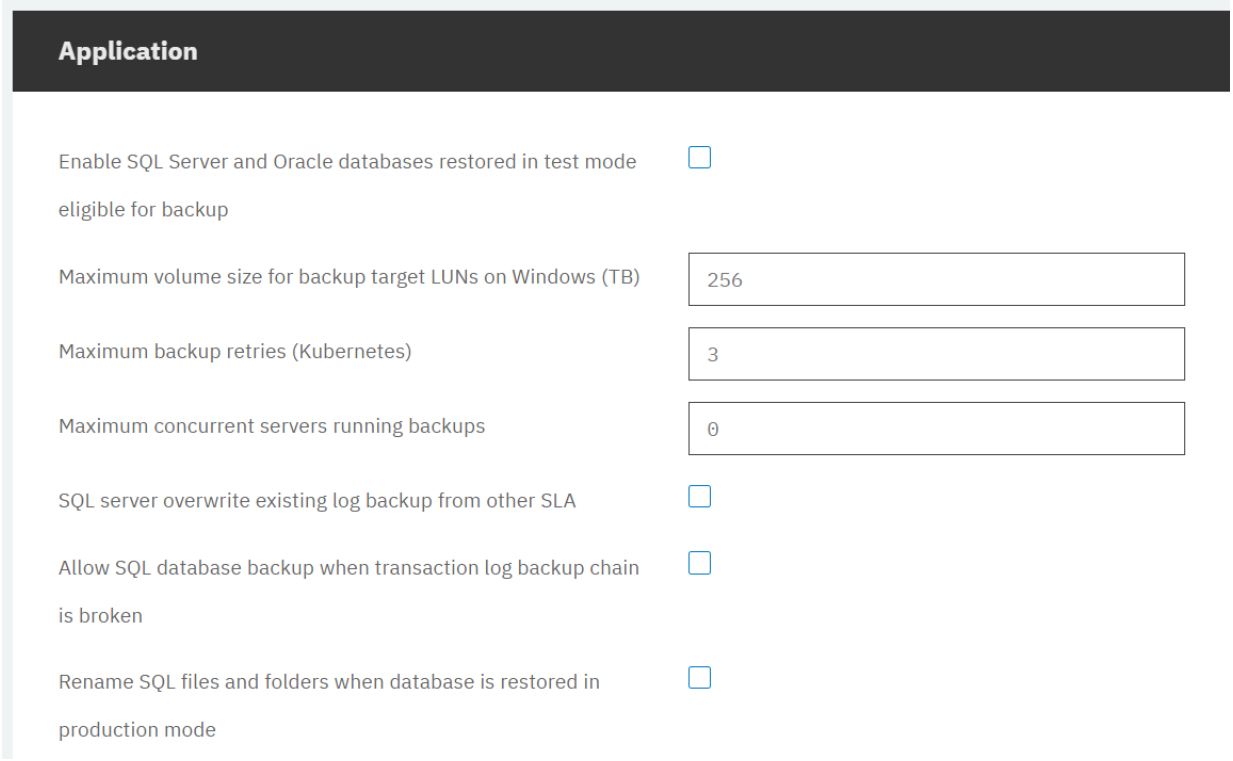
   The URL is created during the Operations Center installation process. To obtain the Operations Center URL, contact the IBM Spectrum Protect system administrator.



3. To set an automatic log out time for the IBM Spectrum Protect Plus user interface after a period of inactivity, edit the **Auto log out (minutes)** preference in the **User interface** category.



4. To apply global application preferences, edit the settings in the **Application** category.



You can edit the following application preferences:

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

Application inventory SSH timeout (hours)
The amount of time in hours after which the SSH session for an application inventory job will expire. The default setting is 12 hours with a maximum of 2048 hours.

Maximum volume size for backup target LUNs on Windows (TB)
The maximum size of the storage for a backup target.

Maximum backup retries (Kubernetes)
The maximum number of times that IBM Spectrum Protect Plus reattempts backup sessions for a copy backup job that contains multiple persistent volume claims (PVCs).

When multiple PVCs are involved in the same copy backup job, IBM Spectrum Protect Plus runs the backup operations as parallel jobs. To help prevent the backup sessions from timing out due to connection issues, specify the maximum number of times that IBM Spectrum Protect Plus reattempts the connections.

If the maximum number of retries is reached and connection failures still exist, only the PVC backups that were part of the failed sessions will be reported as failed.

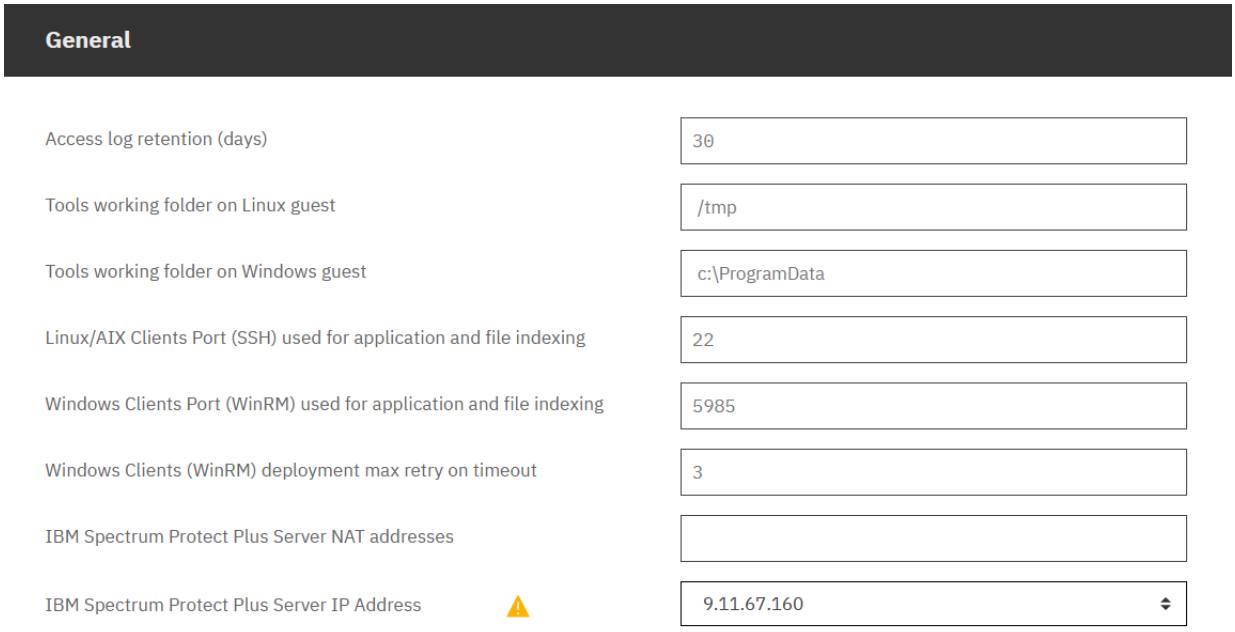
Maximum concurrent servers running backups
The maximum number of concurrent application servers per backup session.

SQL server overwrite existing log backup from other SLA
Enable to allow for the overwriting of the existing SQL Server log backup that belongs to another SLA policy.

Allow SQL database backup when transaction log backup chain is broken
Run a database backup job when IBM Spectrum Protect Plus detects a break in the log backup chain for a database.

Rename SQL files and folders when database is restored in production mode
Rename associated SQL database data and log files during a restore job in production mode. This option applies only when a new database name is provided during an SQL database restore job.

5. To apply general preferences, edit the settings in the **General** category.



You can edit the following general preferences:

Access log retention (days)
Enter the number of days that the access log should be retained. The default setting is 30 days with a maximum setting of 1825 days (five years).

Tools working folder on Linux guest
The working folder for tools on Linux VM guests.

Tools working folder on Windows guest
The working folder for tools on Windows VM guests.

Linux/AIX Clients Port (SSH) used for application and file indexing
The SSH port that is used for application and file indexing on Linux and AIX clients.

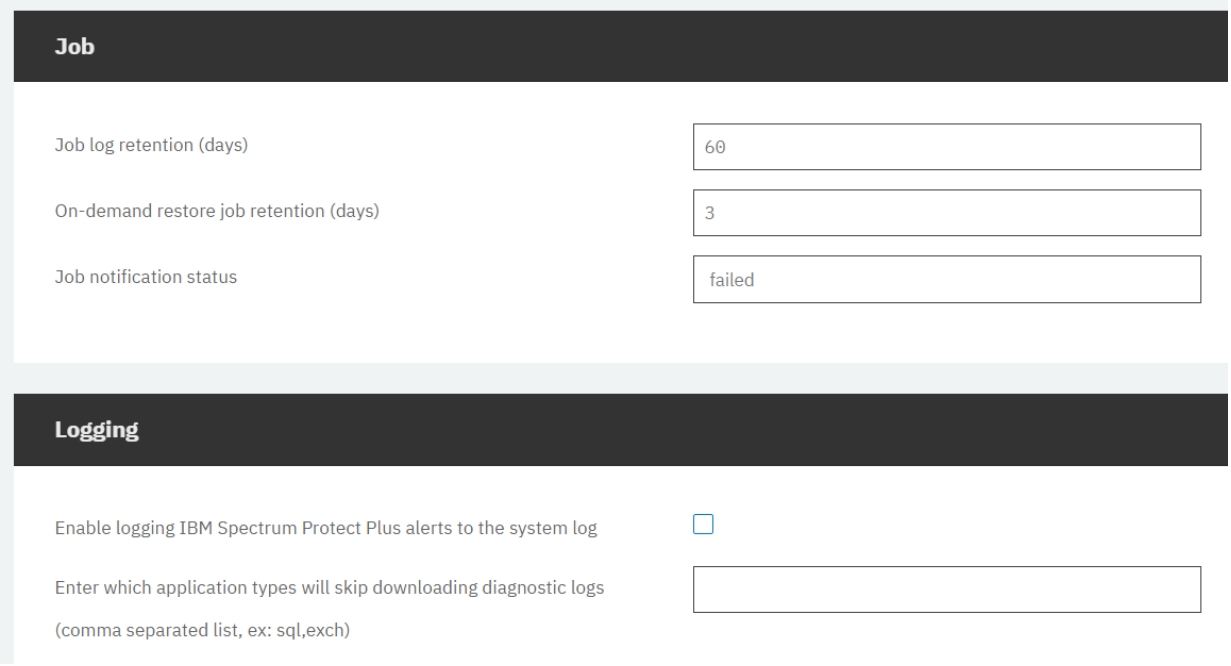
Windows Clients Port (WinRM) used for application and file indexing
The Windows Remote Management (WinRM) port that is used for application and file indexing on Windows clients.

Windows Clients (WinRM) deployment max retry on timeout
This is for Windows clients that utilize Windows Remote Management (Win Rm). If a timeout occurs, the value determines the number of retries that should occur. The value can range from 0 to 5 with a default of 3.

IBM Spectrum Protect Plus Server NAT Addresses
The list of NAT IP addresses or FQDNs for the IBM Spectrum Protect Plus server.

IBM Spectrum Protect Plus Server IP Address
The list of available IP addresses for the IBM Spectrum Protect Plus server. The addresses are used for remote agent communication.

6. To apply job or logging preferences, edit the values in the **Job** or **Logging** categories.



You can edit the following job and logging preferences:

Job log retention (days)
The number of days to retain job logs before the logs are deleted.

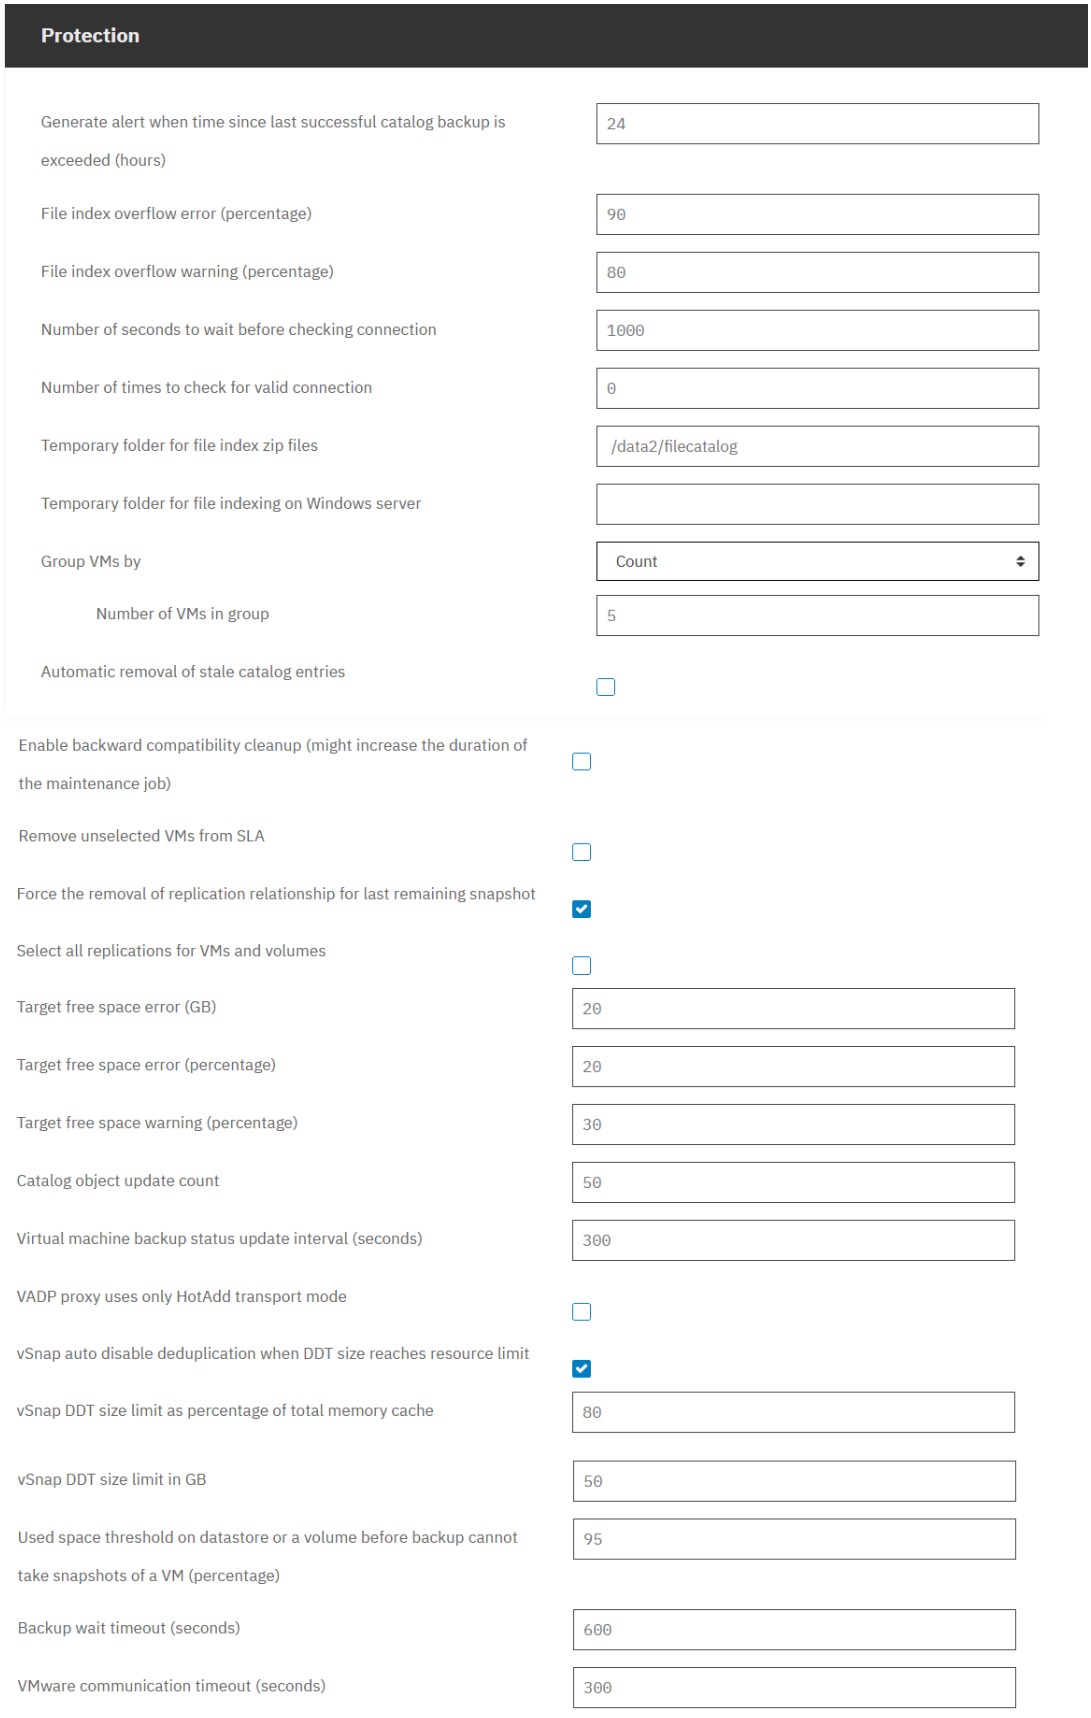
On-demand restore job retention (days)
The number of days to retain on-demand restore jobs before the jobs, sessions, and logs are removed. If a value of 0 days is used, on-demand restore jobs are not removed. The default setting is 3 days.

Job notification status
The status level for sending alerts. Alerts are sent when a job is completed with the specified status. For example, if the job notification status is **`failed`**, when the `failed` status is reported for a job, an alert is sent.

Enable logging IBM Spectrum Protect Plus alerts to the system log
Include alerts that are generated by IBM Spectrum Protect Plus in the system log. After you enable this feature, you can search the system log to find alerts.

Enter which application types will skip downloading diagnostic logs (comma separated list, ex: sql, exch)
As part of each backup, IBM Spectrum Protect Plus collects a diagnostic log from the application server if the application agent generates a diagnostic log. You can add an application to this list to disable the downloading of the diagnostic logs for specific applications, even if they are generated by the application agent. Application types include: `sql`, `oracle`, db2, `exch`, `mongo`, `kuberbetesvol`, `office365`, k8s, `file`, and `openshift`. Multiple application types can be entered and must be separated by commas. The default setting is to download all application agent diagnostic logs.

7. To apply protection preferences, edit the settings in the **Protection** category.



You can edit the following protection preferences:

Generate alert when time since last successful catalog backup is exceeded (hours)
The threshold in hours of when an alert should be generated if a successful catalog backup has not occurred. The alert is displayed on the IBM Spectrum Protect Plus product dashboard and in the **Alerts** window.

File index overflow error (percentage)
The threshold as a percentage for the file index count that when exceeded, results in an error in the job log. For example, if a value of 90 is specified, an error is displayed if the file index count reaches 90 percent of the allowed maximum of 2,147,483,647 entries.

File index overflow warning (percentage)
The threshold as a percentage for the file index count that when exceeded, results in a warning in the job log. For example, if a value of 80 is specified, a warning is displayed if the file index count reaches 80 percent of the allowed maximum of 2,147,483,647 entries.

Number of seconds to wait before checking connection
The amount of time that IBM Spectrum Protect Plus waits before checking the connection to a cloud object.

Number of times to check for valid connection
The number of times that IBM Spectrum Protect Plus checks for an available connection.

Temporary folder for file index zip files
The temporary folder for storing the compressed (`.zip`) files that contain the metadata for indexing. When the indexing is completed, the files are deleted.

Temporary folder for file indexing on Windows server
The temporary folder for storing the compressed (`.zip`) files that contain the metadata for indexing the Windows server. When the indexing is completed, the folder is deleted.

Group VMs by
Virtual machines can be grouped together. The group can be defined by a count of the VMs that are included in the group.

Number of VMs in group
For VM grouping, four VM groups are available and each VM group can have a maximum of five VMs. Each group corresponds to one destination volume (data stream). A maximum of 20 VMs (four data streams) can be grouped at a time based on size calculations.

Important: The default value for **Number of VMs in group** changed from 1 to 5 beginning in IBM Spectrum Protect Plus version 10.1.7.

If you are upgrading from version 10.1.6 and earlier and use the default value 1 for this option, the value is automatically changed to 5. If you have a custom value in this field, that value is maintained.

Automatic removal of stale catalog entries
Enables the checking and removal of stale catalog entries when compared with storage servers stored in MongoDB catalog. This action occurs during the execution of maintenance jobs. This will increase the overall time taken to complete the maintenance job.

Enable backward compatibility cleanup (might increase the duration of the maintenance job)
Enables the option to allow for the cleaning of both expired and unexpired snapshots that have reached their retention limit. Enabling this option might increase the duration of the maintenance job.

Remove unselected VMs from SLA
Enables the option to allow for VMware virtual machines (VMs) to be removed from the volume on vSnap server for corresponding SLA policies if the VM is no longer selected in the SLA policy. This will increase the overall time taken by maintenance and should only be enabled when required.

Force the removal of the replication relationship for last remaining snapshot
Remove an existing replication relationship for the last remaining snapshot that is set to expire and is locked.

Select all replications for VMs and volumes
Enables the capability for all replication jobs to select VMs and volumes from the beginning of the associated policy instead of selecting only from the last successful replication.

Target free space error (GB)
The size threshold in Gigabytes of remaining free space in the vSnap storage pool. Errors are displayed in the job log. For example, if a value of 100 is specified, an error is displayed if the vSnap storage pool has 100 GB or less of remaining free space. The error is generated when the threshold is reached based on the setting of either the size in this option or percentage as specified in **Target free space error (percentage)**. The resources will fail to back up to the affected vSnap storage pool if the threshold is reached.

Target free space error (percentage)
The percentage threshold of remaining free space in the vSnap storage pool. Errors are displayed in the job log. For example, if a value of 5 is specified, an error is displayed if the vSnap storage pool has 5% or less of remaining free space. The error is generated when the threshold is reached based on the setting of either the percentage in this option or size as specified in **Target free space error (GB)**. The resources will fail to back up to the affected vSnap storage pool if the threshold is reached.

Target free space warning (percentage)
The percentage threshold of remaining free space in the vSnap storage pool. Warnings are displayed in the job log. For example, if a value of 10 is specified, a warning is displayed if the vSnap storage pool has 10% or less of remaining free space.

Catalog object update count
The count that you can set to limit how many objects are queried and updated in the catalog. For example, if the catalog includes 100 objects and the update count is 20, IBM Spectrum Protect Plus updates the catalog in five iterations.

Virtual machine backup status update interval (seconds)
The frequency at which messages about the progress of data transfer are updated in the job log.

VADP proxy uses only HotAdd transport mode
Use the HotAdd virtual disk transport method to connect the VMware IBM Spectrum Protect Plus virtual appliance with VADP proxies. If this option is enabled, VADP proxies will use HotAdd only without falling back to an alternate transport mode.

vSnap auto disable deduplication when DDT size reaches resource limit
The deduplication table (DDT) is enabled by default. When either of the threshold limits defined by disk space (gigabytes) or percentage is exceeded, vSnap data deduplication is disabled and an alert is displayed.

vSnap DDT size limit as percentage of total memory cache
The threshold as a percentage of the vSnap deduplication table (DDT) as compared to the total memory cache. The DDT is disabled when the vSnap auto disable option is selected and the defined threshold is exceeded.

vSnap DDT size limit in GB
The threshold, in gigabytes (GB), of the vSnap DDT. The DDT is disabled when the vSnap auto disable option is selected and the defined threshold is exceeded.

Used space threshold on datastore or a volume before backup cannot take snapshots of a VM (percentage)
The percentage of used space on a datastore or a volume that is the threshold before snapshots of a VM cannot be taken for backup.

Backup wait timeout (seconds)
The amount of time that IBM Spectrum Protect Plus waits for a backup job to finish before starting another backup job. If the backup job does not finish within the wait period, the job is timed out, and the next job begins.

VMware communication timeout (seconds)
The amount of time that IBM Spectrum Protect Plus waits for commands that are issued to connected vCenters to finish. If the operations do not finish within the specified amount of time, they are logged as errors. This setting applies only to VMware hypervisors.

8. To apply a security preference, edit the setting in the **Security** category.



You can edit the following security preference:

Set minimum password length (characters)
The minimum length of passwords for IBM Spectrum Protect Plus. By default, the password has a minimum length of 8 characters, but you can specify a longer password. This value applies to all user accounts.

Configuring IBM Spectrum Protect Plus installed as a virtual appliance

There are some configuration tasks that apply only when IBM Spectrum Protect Plus is installed as a virtual appliance.

Logging on to the administrative console

Log on to the administrative console to review the configuration of the IBM Spectrum Protect Plus virtual appliance. Available information includes general system settings, network, and proxy settings.

Procedure

To log on to the administrative console, complete the following steps:

1. From a supported browser, enter the following URL:

   ```
   https://HOSTNAME:8090/
   ```

   where *HOSTNAME* is the IP address of the virtual machine where the application is deployed.
2. In the login window, from the **Authentication Type** list, select one of the following authentication types:

Authentication Type	Logon information
IBM Spectrum Protect Plus	To log on as the IBM Spectrum Protect Plus superuser, enter the username and password. The IBM Spectrum Protect Plus super user is the user who is assigned the SUPERUSER role.
System	To log in as a system user, enter the `serveradmin` password.

Related concepts

"System requirements " on page 27
Before you install IBM Spectrum Protect Plus, review the hardware and software requirements for the product and other components that you plan to install in the storage environment.

"Managing roles" on page 648

Roles define the actions that can be completed for the resources that are defined in a resource group. While a resource group defines the resources that are available to an account, a role sets the permissions to interact with the resources.

Logging on to the virtual appliance

Log on to the IBM Spectrum Protect Plus virtual appliance by using the vSphere Client to access the command line. You can access the command line in a VMware environment or in a Hyper-V environment.

Accessing the virtual appliance in VMware

In a VMware environment, log on to the IBM Spectrum Protect Plus virtual appliance through vSphere Client to access the command line.

Procedure

Complete the following steps to access the virtual appliance command line:

1. In vSphere Client, select the virtual machine where IBM Spectrum Protect Plus is deployed.
2. On the **Summary** tab, select **Open Console** and click in the console.
3. Select **Login**, and enter your user name and password. The default user name is `serveradmin` and the default password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in "Start IBM Spectrum Protect Plus" on page 257.

What to do next

Enter commands to administer the virtual appliance. To log off, type `exit`.

Accessing the virtual appliance in Hyper-V

In a Hyper-V environment, log on to the IBM Spectrum Protect Plus virtual appliance through vSphere Client to access the command line.

Procedure

Complete the following steps to access the virtual appliance command line:

1. In Hyper-V Manager, select the virtual machine where IBM Spectrum Protect Plus is deployed.
2. Right-click the virtual machine and select **Connect**.
3. Select **Login**, and enter your user name and password. The default user name is `serveradmin` and the default password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in "Start IBM Spectrum Protect Plus" on page 257.

What to do next

Enter commands to administer the virtual appliance. To log off, type `exit`.

Setting the time zone

Use the administrative console to set the time zone of the IBM Spectrum Protect Plus virtual appliance.

Procedure

To set the time zone, complete the following steps:

1. From a supported browser, enter the following URL:

```
https://HOSTNAME:8090/
```

where *HOSTNAME* is the IP address of the virtual machine where the application is deployed.

2. In the login window, from the **Authentication Type** list, select one of the following authentication types:

Authentication Type	Login information
IBM Spectrum Protect Plus	To log in as the IBM Spectrum Protect Plus superuser, enter the username and password. The IBM Spectrum Protect Plus super user is the user who is assigned the SUPERUSER role.
System	To log in as a system user, enter the `serveradmin` password.

3. Click **System Management**.
4. In the **Timezone Management** section, select your time zone.

   A message stating that the operation was successful displays. All IBM Spectrum Protect Plus logs and schedules will reflect the selected time zone. The selected time zone will also display on the IBM Spectrum Protect Plus virtual appliance when logged in with the user ID **serveradmin**.
5. Restart the IBM Spectrum Protect Plus virtual appliance from the administrative console.
6. When the IBM Spectrum Protect Plus virtual appliance has restarted, view the current time zone. Select **Product Information** from the main page of the administrative console and verify the updated time zone.

Adding virtual disks

You can add new virtual disks (hard disks) to your IBM Spectrum Protect Plus virtual appliance by using vCenter.

When you deploy the IBM Spectrum Protect Plus virtual appliance, you can deploy all virtual disks to one datastore that you specify at the time of deployment. You can add a disk within the virtual appliance and configure it as a Logical Volume Manager (LVM). You can then mount the new disk as a new volume or attach the new disk to the existing volumes within the virtual appliance.

Important: Do not add space or extend an existing volume for the IBM Spectrum Protect Plus virtual appliance.

You can review the disk partitions by using the **`fdisk -l`** command. You can review the physical volumes and the volume groups on the IBM Spectrum Protect Plus virtual appliance by using the **`pvdisplay`** and **`vgdisplay`** commands.

Adding a disk to the virtual appliance

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

Before you begin

To run commands, you must connect to the command line for the IBM Spectrum Protect Plus virtual appliance by using Secure Shell (SSH) and log in with the user ID `serveradmin`. The default initial password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in “Start IBM Spectrum Protect Plus” on page 257.

Procedure

To add a disk to an IBM Spectrum Protect Plus virtual appliance, complete the following steps from the vCenter client:

1. From the vCenter client, complete the following steps:

a) On the **Hardware** tab, click **Add**.

b) Select **Create a new virtual disk**.

c) Select the required disk size. In the **Location** section, select one of the following options:

- To use the current datastore, select **Store with the virtual machine**.
- To specify one or more datastores for the virtual disk, select **Specify a datastore or datastore cluster**. Click **Browse** to select the new datastores.

d) In the **Advanced Options** tab, leave the default values.

e) Review and save your changes.

f) Click the **Edit Settings** option for the virtual machine to view the new hard disk.

2. Add the new SCSI device without rebooting the virtual appliance. From the console of the IBM Spectrum Protect Plus appliance, issue the following commands:

```
sudo bash
```

Press Enter.

```
# for host in `ls /sys/class/scsi_host/`; do
echo "- - -" > /sys/class/scsi_host/${host}/scan;
done
```

Adding storage capacity from a new disk to the appliance volume

After you add a disk to the virtual appliance, you can attach the new disk to the existing volumes within the virtual appliance.

Before you begin

To run commands, you must connect to the console of the IBM Spectrum Protect Plus virtual appliance by using SSH and log in with the user ID `serveradmin`. The default initial password is `sppDP758-SysXyz`. You are prompted to change this password during the first logon. Certain rules are enforced when creating a new password. For more information, see the password requirement rules in "Start IBM Spectrum Protect Plus" on page 257.

About this task

You need to complete this task only if you want to add the storage capacity from a new disk to an existing appliance volume. If you added the disk as a new volume, you do not need to complete this task.

Procedure

To add storage capacity from a new disk to the appliance volume, complete the following steps from the console of the virtual appliance:

1. Complete the following steps to set up a partition for the new disk and set the partition to be of type Linux LVM:

a) Open the new disk by using the **fdisk** command. For the command below, the disk `/dev/sdd` is used as an example. Use the **fdisk** command with the appropriate disk that is to be added.

```
[serveradmin@localhost ~]# fdisk /dev/sdd
```

The **fdisk** utility starts in interactive mode. Output similar to the following output is displayed:

```
Device contains neither a valid DOS partition table, nor Sun, SGI or
OSF disklabel
Building a new DOS disklabel with disk identifier 0xb1b293df.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.
Warning: invalid flag 0x0000 of partition table 4 will be corrected by
w(rite)
WARNING: DOS-compatible mode is deprecated. It's strongly recommended
to
switch off the mode (command 'c') and change display units to
sectors (command 'u').
Command (m for help):
```

a) At the **fdisk** command line, enter the **n** subcommand to add a partition.

```
Command (m for help): n
```

The following command action choices are displayed:

```
Command (m for help): n
Command action
e extended
p primary partition (1-4)
```

b) Enter the **p** command action to select the primary partition.
You are prompted for a partition number:

```
Command (m for help): n
Command action
e extended
p primary partition (1-4)
Partition number (1-4):
```

c) At the partition number prompt, enter the partition number 1.

```
Partition number (1-4): 1
```

The following prompt is displayed:

```
First cylinder (1-2610, default 1):
```

d) Do not type anything at the First cylinder prompt. Press the **Enter** key.
The following output and prompt is displayed:

```
First cylinder (1-2610, default 1):
Using default value 1
Last cylinder, +cylinders or +size{K,M,G} (1-2610, default 2610):
```

e) Do not type anything in the Last cylinder prompt. Press the **Enter** key.
The following output is displayed:

```
Last cylinder, +cylinders or +size{K,M,G} (1-2610, default 2610):
Using default value 2610
Command (m for help):
```

f) At the **fdisk** command line, enter the **t** subcommand to change a partition's system ID.

```
Command (m for help): t
```

You are prompted for a hex code that identifies the partition type:

```
Selected partition 1
Hex code (type L to list codes):
```

g) At the Hex code prompt, enter the hex code 8e to specify the Linux LVM partition type.
   The following output is displayed:

```
Hex code (type L to list codes): 8e
Changed system type of partition 1 to 8e (Linux LVM)
Command (m for help):
```

h) At the **fdisk** command line, enter the **w** subcommand to write the partition table and to exit the **fdisk** utility.

```
Command (m for help): w
```

   The following output is displayed:

```
Command (m for help): w (write table to disk and exit)
The partition table has been altered!
Calling ioctl() to re-read partition table.
Syncing disks.
```

2. To review the changes to the disk, issue the **fdisk -l** command.
3. To review the current list of Physical Volumes (PV), issue the **pvdisplay** command.
4. To create a new Physical Volume (PV), issue the **pvcreate /dev/sdd1** command.
5. To view the new PV from /dev/sdd1, issue the **pvdisplay** command.
6. To review the Volume Group (VG), issue the **vgdisplay** command.
7. To add the Physical Volume (PV) to the Volume Group (VG) and increase the space of the VG, issue the following command:

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

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

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

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

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

Resetting the serveradmin password

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

About this task

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

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

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

Procedure

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

   Important: By ensuring that no jobs are running when you shut down the virtual appliance, you help to prevent possible issues.
2. Log in to the vSphere Client.
3. Restart the IBM Spectrum Protect Plus 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 10. Managing SLA policies for backup operations

Service level agreement (SLA) policies, also known as backup policies, define parameters for backup jobs. These parameters include the frequency and retention period of backups and the option to replicate or copy backup data. You can use predefined SLA policies, or customize them to meet your needs.

The following default SLA policies are available. Each policy specifies a frequency and retention period for the backup. You can use these policies as they are or modify them. You can also create custom SLA policies.

Gold
This policy runs every 4 hours with a retention period of 1 week. For all supported resources except for Amazon EC2 instances and containers.

Silver
This policy runs daily with a retention period of 1 month. For all supported resources except for Amazon EC2 instances and container data.

Bronze
This policy runs daily with a retention period of 1 week. For all supported resources except for Amazon EC2 instances and container data.

EC2
To protect Amazon EC2 instances, this policy runs daily snapshot backups with a retention period of 31 days.

Container
To protect container data, this policy runs the following operations:

- Snapshot backups every 6 hours with a retention period of 1 day
- Copy backups daily with a retention period of 31 days.

To view and manage backup policies and to monitor the virtual machines and databases that are protected by policies, click **Manage Protection** > **Policy Overview** in the navigation panel.

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

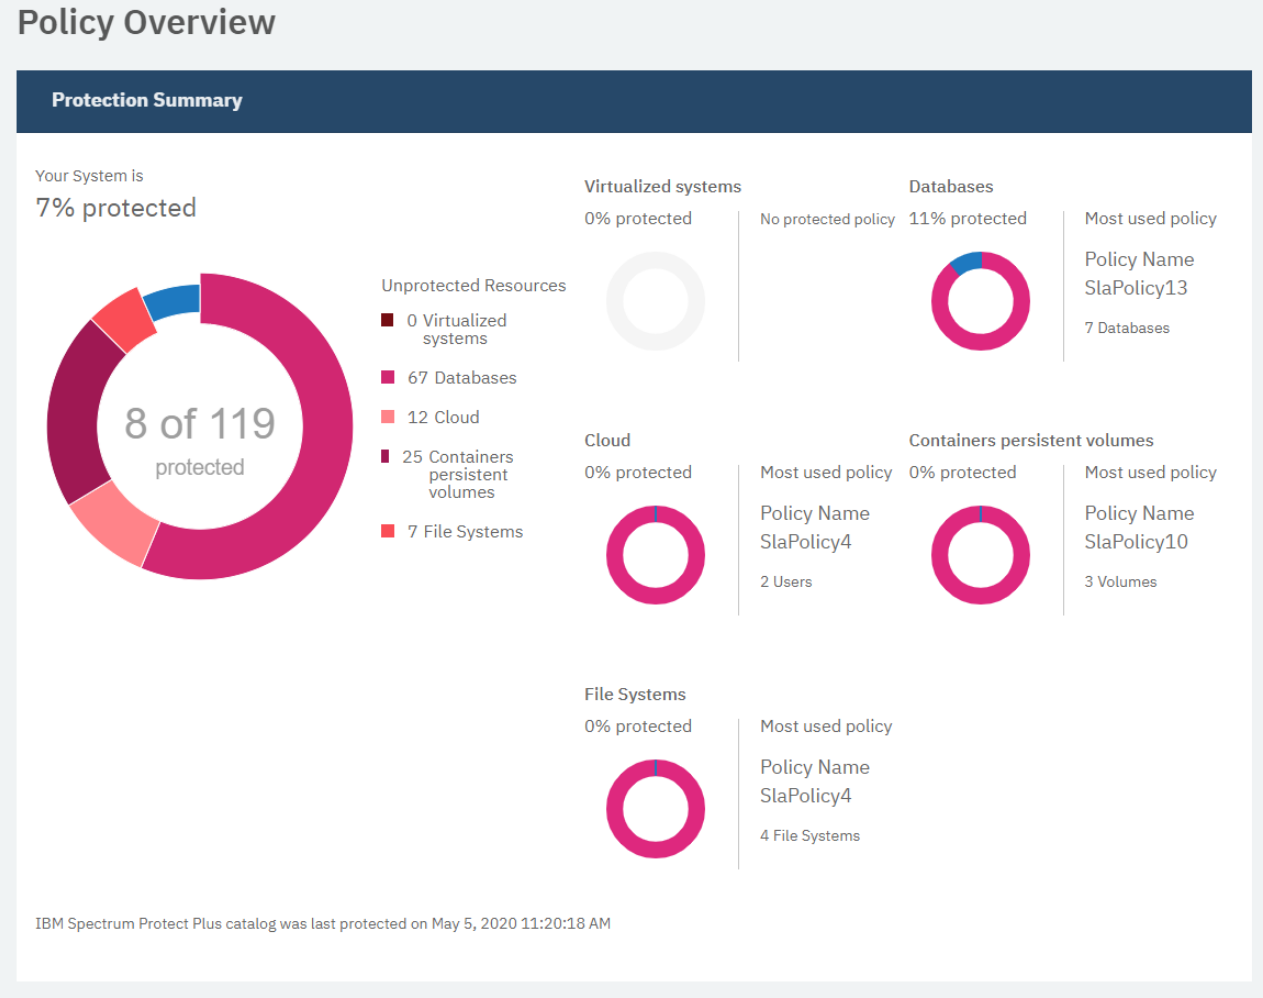
Protection Summary

You can view the protection status of the resources in your system in the **Protection Summary** pane.

The **Protection Summary** pane consists of donut charts that depict the number of protected resources versus the number of unprotected resources. For each type of resource, you can view the percentage of the resource that is protected, and the service level agreement (SLA) policy that is most frequently used for that resource.

To view the **Protection Summary** pane, from the navigation panel, click **Manage Protection** > **Policy Overview**.

© Copyright IBM Corp. 2017, 2021



System

The **Your System** chart shows the total percentage of resources in your system that are protected by IBM Spectrum Protect Plus.

% protected

Shows the percentage of resources that are protected by IBM Spectrum Protect Plus. In the donut chart, the protected resources are represented by the blue line. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected resources.

Unprotected Resources

Shows the legend of unprotected resources. In the list, data is shown only for the types of resources that are managed by your instance of IBM Spectrum Protect Plus. If a type of resource is not managed by IBM Spectrum Protect Plus, the count is 0.

Virtualized systems

The **Virtualized systems** chart shows the percentage of virtualized system that are protected by IBM Spectrum Protect Plus.

% protected

Shows the percentage of virtualized systems that are protected. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected virtualized systems.

If no virtualized systems are managed by IBM Spectrum Protect Plus, the percentage is 0.

Most used policy

Shows the name of the most frequently used SLA policy, and the number of virtualized systems that are using that policy. If no virtualized systems are managed by IBM Spectrum Protect Plus, this field is not displayed.

No protected policy
This message is shown only when no virtualized systems are managed by IBM Spectrum Protect Plus.

Databases
The **Databases** chart shows the percentage of databases that are protected by IBM Spectrum Protect Plus.

% protected
Shows the percentage of databases that are protected. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected databases.

If no application databases are managed by IBM Spectrum Protect Plus, the percentage is 0.

Most used policy
Shows the name of the most frequently used SLA policy, and the number of databases that are using that policy. If no databases are managed by IBM Spectrum Protect Plus, this field is not displayed.

No protected policy
This message is shown only when no databases are managed by IBM Spectrum Protect Plus.

Cloud
The **Cloud** chart shows the percentage of cloud-based accounts, such as Microsoft Office 365 tenants, that are protected by IBM Spectrum Protect Plus.

% protected
Shows the percentage of cloud-based accounts that are protected. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected accounts.

If no cloud-based accounts are managed by IBM Spectrum Protect Plus, the percentage is 0.

Most used policy
Shows the name of the most frequently used SLA policy, and the number of accounts that are using that policy. If no cloud-based accounts are managed by IBM Spectrum Protect Plus, this field is not displayed.

No protected policy
This message is shown only when no cloud-based accounts are managed by IBM Spectrum Protect Plus.

Containers persistent volumes
Shows the percentage of persistent volumes that are protected by IBM Spectrum Protect Plus.

% protected
Shows the percentage of persistent volumes that are protected. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected persistent volumes.

If no persistent volumes are managed by IBM Spectrum Protect Plus, the percentage is 0.

Most used policy
Shows the name of the most frequently used SLA policy, and the number of persistent volumes that are using that policy. If no persistent volumes are managed by IBM Spectrum Protect Plus, this field is not displayed.

No protected policy
This message is shown only when no persistent volumes are managed by IBM Spectrum Protect Plus.

File Systems
Shows the percentage of file systems that are protected by IBM Spectrum Protect Plus.

% protected
Shows the percentage of file systems that are protected. By hovering your cursor over the different parts of the donut, you can view the numbers of protected and unprotected file systems.

If no file systems are managed by IBM Spectrum Protect Plus, the percentage is 0.

Most used policy
Shows the name of the most frequently used SLA policy, and the number of file systems that are using that policy. If no file systems are managed by IBM Spectrum Protect Plus, this field is not displayed.

No protected policy
This message is shown only when no file systems are managed by IBM Spectrum Protect Plus.

Creating an SLA policy for hypervisors, databases, and file systems

You can create custom service level agreement (SLA) policies to define backup frequency, retention, replication, and copy policies that are specific for your environment.

About this task

If a virtual machine is associated with multiple SLA policies, ensure that the policies that you create are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

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

When copying data from a vSnap server to cloud storage, the most recent successfully completed snapshot will be copied.

Procedure

To create an SLA policy for hypervisors, databases, and file systems, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Policy Overview**.
2. Click **Add SLA Policy**.
   The **New SLA Policy** pane is displayed.
3. In the **Name** field, enter a name that provides a meaningful description of the SLA policy.
4. Click **VMware, Hyper-V, Exchange, Microsoft 365, SQL, Oracle, Db2, MongoDB, Catalog, and Windows File Systems**.
5. In the **Backup Policy** section, set the following options for backup operations. These operations occur on the vSnap servers that are defined in the **System Configuration** > **Backup Storage** > **Disk** window.

   Retention
   Specify the retention period for the backup snapshots.

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

   Repeats
   Enter a frequency for backup operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

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

   The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

   Target Site
   Select the target backup site for backing up data.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, IBM Spectrum Protect Plus server manages data placement in the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a vSnap server, are not shown.

Only use encrypted disk storage
Select this check box to back up data to encrypted vSnap servers if your environment includes a mixture of encrypted and unencrypted servers.

Restriction: If this option is selected and there are no encrypted vSnap servers available, the associated job will fail.

6. Under **Replication Policy**, set the following options to enable asynchronous replication from one vSnap server to another. For example, you can replicate data from the primary to the secondary backup site.

Replication partnerships requirement: These options apply to established replication partnerships. To add a replication partnership, see the instructions in "Configuring backup storage partners" on page 164.

Backup Storage Replication
Select this option to enable replication.

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

Frequency
Enter a frequency for replication operations.

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

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site
Select the target backup site for replicating data.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, IBM Spectrum Protect Plus server manages data placement in the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus, but are not associated with a vSnap server, are not shown.

Only use encrypted disk storage
Select this option to replicate data to encrypted vSnap servers if your environment includes a mixture of encrypted and unencrypted servers.

Restriction: If this option is selected and there are no encrypted vSnap servers available, the associated job will fail.

Same retention as source selection
Select this option to use the same retention policy as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

7. In the **Additional copies** section, set the following options to copy data to standard object storage or archive object storage.

Standard object storage (incremental copy)
Select this option to copy data to cloud storage or to a repository server.

Data is backed up to the vSnap server for short term protection, and then copied to the selected cloud storage or repository server for longer-term protection. During the first copy of a backup volume, the snapshot is backed up in full. After the first copy of the base snapshot is completed,

subsequent copies are incremental and capture cumulative changes since the last copy. Cloud or repository server restore operations can be performed from any available vSnap server.

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

Frequency
Enter a frequency for copy operations.

Start Time
Enter the date and time that you want the copy operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Same retention as source selection
Select this option to use the same retention policy as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

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

Source
Click the source for the copy operation:

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

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

This option is available only when **Backup Storage Replication** is selected.

Destination
Click **Cloud services** or **Repository servers**.

Target
Click the cloud storage system or repository server to which you want to copy data.

This list contains the secondary storage systems that you have added to IBM Spectrum Protect Plus. If you have not added secondary storage or want to add it, see "Managing secondary backup storage" on page 271 for information about the cloud storage systems and repository servers that are supported and how to add them to IBM Spectrum Protect Plus.

Archive object storage (full copy)
Select this option to archive data to cloud storage or to a repository server for long-term protection.

This operation provides a full image copy to the selected archival storage.

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

Frequency
Enter a frequency for archive operations.

Start Time
Enter the date and time that you want the archive operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Retention
Specify the retention period for the archive snapshots as a unit of time in days, months, or years.

Source
Click the source for the archive destination:

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

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

This option is available only when **Backup Storage Replication** is selected.

Destination
Click **Cloud services** or **Repository servers**.

Target
Click the cloud storage system or repository server to which you want to archive data.

Only cloud targets that have a defined archive bucket are shown in this list. To add an archive bucket for a cloud storage system, follow the instructions in "Managing cloud storage" on page 271.

8. Click **Save**. The SLA policy can now be applied to backup job definitions.

What to do next

After you create an SLA policy, complete the following actions:

Action	How to
Assign user permissions to the SLA policy.	See "Creating a role" on page 650
Create a backup job definition that uses the SLA policy.	See the backup topics in Chapter 11, "Protecting virtualized systems," on page 335, Chapter 15, "Protecting databases," on page 497, and Chapter 12, "Protecting file systems," on page 383.

Related concepts

"Replicate backup-storage data " on page 14
When you enable replication of backup data, data from one vSnap server is asynchronously replicated to another vSnap server. For example, you can replicate backup data from a vSnap server on a primary site to a vSnap server on a secondary site.

"Copy snapshots to secondary backup storage" on page 15
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.

Related tasks

"Creating an SLA policy for Amazon EC2 instances" on page 328
You can create custom service level agreement (SLA) policies to define snapshot retention and frequency policies that are specific to Amazon EC2 instances.

"Creating an SLA policy for containers" on page 329

You can create custom service level agreement (SLA) policies for persistent volumes that are attached to a Kubernetes or OpenShift cluster. You can define the frequency of snapshot and backup operations and specify policies for retention, replication, and copy jobs.

Creating an SLA policy for Amazon EC2 instances

You can create custom service level agreement (SLA) policies to define snapshot retention and frequency policies that are specific to Amazon EC2 instances.

About this task

When a scheduled backup job runs, a snapshot of the instance is created at the frequency that is defined by the snapshot policy.

If an instance is associated with multiple SLA policies, ensure that the policies that you create are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

Procedure

To create an SLA policy for your instances complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Policy Overview**.
2. Click **Add SLA Policy**.

   The **New SLA Policy** pane is displayed.
3. In the **Name** field, enter a name that provides a meaningful description of the SLA policy.
4. Click **Amazon EC2**.

   The SLA policy options for EC2 instances are displayed.
5. In the **Snapshot Protection** section, set the following options for snapshot operations:

   Retention
   Specify the retention period for the snapshots.

   Disable Schedule
   Select this checkbox to create the snapshot policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand. This field is optional.

   Repeats
   Enter a frequency for snapshot operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

   Start Time
   Enter the date and time when you want the snapshot operation to start.

   The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

   Snapshot Prefix
   Enter a prefix to add to the beginning of snapshot names. Prefixes can help you organize and easily identify snapshots. This field is optional.

   For example, if you entered the prefix "daily_", all snapshot names that are created with this SLA policy will begin with "daily_".
6. Click **Save.**

The SLA policy that you created is displayed in the table in the SLA Policies pane.

What to do next

After you create an SLA policy, complete the following actions:

- Assign user permissions to the SLA policy. For instructions, see "Creating a role" on page 650.
- Create a backup job definition that uses the SLA policy. For instructions, see "Backing up Amazon EC2 data" on page 376.

Related tasks

"Editing an SLA policy" on page 333
Edit the options for an SLA policy to reflect changes in your IBM Spectrum Protect Plus environment.

"Deleting an SLA policy" on page 334
Delete an SLA policy when it becomes obsolete.

Creating an SLA policy for containers

You can create custom service level agreement (SLA) policies for persistent volumes that are attached to a Kubernetes or OpenShift cluster. You can define the frequency of snapshot and backup operations and specify policies for retention, replication, and copy jobs.

Before you begin

If you plan to copy data to secondary storage or to archive data to a cloud storage system, take the following actions:

- If you plan to copy data to secondary storage, such as a cloud storage system or repository server, ensure that the secondary storage is configured. For information about the secondary storage systems that are supported and for configuration instructions, see "Managing secondary backup storage" on page 271
- If you plan to archive data to a cloud storage system, the cloud target must have a defined archive bucket. To add an archive bucket for a cloud storage system, follow the instructions in "Managing cloud storage" on page 271.

About this task

You can create custom SLA policies if you do not want to use the predefined **Container** policy. The **Container** policy runs the following operations:

- Snapshot backups every 6 hours with a retention period of 1 day
- Copy backups daily with a retention period of 31 days

A snapshot is required in a container backup operation. When a scheduled backup job runs, a snapshot of the persistent volume claim (PVC) is created on your storage system at the frequency that is defined by the snapshot policy. You can specify additional policy settings to copy the snapshot to the IBM Spectrum Protect Plus vSnap server, replicate the vSnap server, or copy the data to object storage in the cloud or to a repository server.

If a PVC is associated with multiple SLA policies, ensure that the policies that you create are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

Procedure

To create an SLA policy for your PVCs, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Policy Overview**.
2. Click **Add SLA Policy**.
   The **New SLA Policy** pane is displayed.
3. In the **Name** field, enter a name that provides a meaningful description of the SLA policy.

4. Click **Kubernetes, OpenShift**.

   The SLA policy options for Kubernetes or OpenShift clusters are displayed.

5. In the **Snapshot Protection** section, set the following options for snapshot operations:

   Retention
   Specify the retention period for the snapshots.

   Disable Schedule
   Select this checkbox to create the snapshot policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand. This field is optional.

   If you plan to enable the policy sections for copy backup, replication, or additional copy operations, ensure that this checkbox is not selected. Otherwise, no snapshots will be available for copying to the vSnap server.

   Repeats
   Enter a frequency for snapshot operations. Enter a frequency for snapshot operations. Choose from **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.

   Start Time
   Enter the date and time when you want the snapshot operation to start.

   The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

   Snapshot Prefix
   Enter a prefix to add to the beginning of snapshot names. You can add a prefix to snapshot names to help you organize and easily identify snapshots. This field is optional.

   You can enter up to 32 characters for the prefix.

   For example, if you entered the prefix "daily", all snapshot names that are created with this SLA policy will begin with "daily".

6. Optional: In the **Backup Policy** section, set the following options for copy backup operations to the vSnap server:

   Backup Storage
   Select this checkbox to enable copy backup operations to the vSnap server. These operations occur on the vSnap servers that are defined in the **System Configuration** > **Backup Storage** > **Disk** window.

   Retention
   Specify the retention period for the copy backups on the vSnap server.

   Disable Schedule
   Select this checkbox to create the backup policy without defining a frequency or start time. Policies that are created without a schedule can be run on demand. This field is optional.

   Frequency
   Enter a frequency for copy backup operations.

   Start Time
   Enter the date and time when you want the copy backup operation to start.

   Tip: Allot time for the snapshot backup to complete before starting the copy backup operation. For example, if the snapshot operation starts at midnight (0:00), set the copy backup operation to start 15 minutes later, at 00:15.

   The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

   Target Site
   Select the target site for backup copies.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, the IBM Spectrum Protect Plus server manages data placement in the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus but are not associated with a vSnap server are not shown.

Only use encrypted disk storage
If your environment includes encrypted and unencrypted servers, select this checkbox to back up data to encrypted vSnap servers.

Restriction: If this option is selected, but no encrypted vSnap servers are available, the associated job fails.

7. Optional: Under **Replication Policy**, set the following options to enable asynchronous replication from one vSnap server to another. For example, you can replicate data from the primary to the secondary backup site.

Replication partnerships requirement: These options apply to established replication partnerships. To add a replication partnership, see the instructions in "Configuring backup storage partners" on page 164.

Backup Storage Replication
Select this option to enable replication.

This option is enabled only when **Backup Policy** is selected.

Disable Schedule
Select this checkbox to create the replication relationship without defining a frequency or start time. This field is optional.

Frequency
Enter a frequency for replication operations.

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

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Target Site
Select the target site for replicating data.

A site can contain one or more vSnap servers. If more than one vSnap server is in a site, the IBM Spectrum Protect Plus server manages data placement in the vSnap servers.

Only sites that are associated with a vSnap server are shown in this list. Sites that are added to IBM Spectrum Protect Plus but are not associated with a vSnap server are not shown.

Only use encrypted disk storage
Select this option to replicate data to encrypted vSnap servers if your environment includes encrypted and unencrypted servers.

Restriction: If this option is selected, but no encrypted vSnap servers are available, the associated job fails.

Same retention as source selection
Select this option to use the same retention policy as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

8. Optional: In the **Additional copies** section, set the options to copy data to standard object storage or archive object storage.

When copying data from a vSnap server to cloud storage, the most recent successfully completed snapshot will be copied.

Standard object storage (incremental copy)
Select this option to copy data to cloud storage or to a repository server. This option is enabled only when **Backup Policy** is selected.

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

Disable Schedule
Select this checkbox to create the copy relationship without defining a frequency or start time. This field is optional.

Frequency
Enter a frequency for copy operations.

Start Time
Enter the date and time that you want the copy operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Same retention as source selection
Select this option to use the same retention policy as the source vSnap server. To set a different retention policy, clear this option and set a different policy.

Restriction: Copy retention options are disabled if a server that uses Write Once Read Many (WORM) retention is selected in the **Target** field.

Source
Click the source for the copy operation:

Backup Policy Destination
The source for the copy operation is the target site that is defined in the **Backup Policy** section.

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

This option is enabled only when **Backup Storage Replication** is selected.

Destination
Click **Cloud services** or **Repository servers**.

Target
Click the cloud storage system or repository server to which you want to copy data.

This list contains the secondary storage systems that you have added to IBM Spectrum Protect Plus.

Archive object storage (full copy)
Select this option to archive data to cloud storage or to a repository server for long-term protection. This option is enabled only when **Backup Policy** is selected.

This operation provides a full image copy to the selected archival storage.

Disable Schedule
Select this checkbox to create the archive relationship without defining a frequency or start time. This field is optional.

Frequency
Enter a frequency for archive operations.

Start Time
Enter the date and time that you want the archive operation to start.

The time zone is automatically populated with your browser settings. To update the time zone, click the field and select a region and city from the list, for example: **Europe/Dublin**. You can

also click the field and enter a region or city in the **Search** field, and select an item from the matching results.

Retention
Specify the retention period for the archive snapshots as a unit of time in days, months, or years.

Source
Click the source for the archive destination:

Backup Policy Destination
The source for the archive operation is the target site that is defined in the **Backup Policy** section.

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

This option is enabled only when **Backup Storage Replication** is selected.

Destination
Click **Cloud services** or **Repository servers**.

Target
Click the cloud storage system or repository server to which you want to archive data.

Only cloud targets that have a defined archive bucket are shown in this list.

9. Click **Save.**
   The SLA policy that you created is displayed in the table in the **SLA Policies** pane.

What to do next

After you create an SLA policy, take the following actions:

- Assign user permissions to the SLA policy. For instructions, see "Creating a role" on page 650.
- Create a backup job definition that uses the SLA policy. For instructions, see "Backing up persistent volumes in a Kubernetes cluster" on page 415.

Related tasks

"Editing an SLA policy" on page 333
Edit the options for an SLA policy to reflect changes in your IBM Spectrum Protect Plus environment.

"Deleting an SLA policy" on page 334
Delete an SLA policy when it becomes obsolete.

Editing an SLA policy

Edit the options for an SLA policy to reflect changes in your IBM Spectrum Protect Plus environment.

Procedure

To edit an SLA policy, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Policy Overview**.
2. Click the edit icon that is associated with a policy.
   The **Edit SLA Policy** pane is displayed.
3. Edit the policy options, and then click **Save**.

Deleting an SLA policy

Delete an SLA policy when it becomes obsolete.

Before you begin

Ensure that there are no jobs that are associated with the SLA policy.

Procedure

To delete an SLA policy, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Policy Overview**.
2. Click the delete icon that is associated with an SLA policy.
3. Click **Yes** to delete the policy.

Chapter 11. Protecting virtualized systems

You must register the virtualized systems that you want to protect in IBM Spectrum Protect Plus and then create jobs to back up and restore the resources that are associated with the systems.

Virtualized systems refers to VMware and Microsoft Hyper-V hypervisors and Amazon EC2 instances.

Backing up and restoring VMware data

To protect VMware data, first add vCenter Server instances in IBM Spectrum Protect Plus, and then create jobs for backup and restore operations for the content of the instances.

Ensure that your VMware environment meets the system requirements in “Hypervisor (Microsoft Hyper-V and VMware) and cloud instance (Amazon EC2) backup and restore requirements ” on page 57.

Support for VMware tags

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

Support for encryption

Backing up and restoring encrypted virtual machines is supported in vSphere 6.5 environments and later. Encrypted virtual machines can be backed up and restored at the virtual-machine level to their original location. If you are restoring a virtual machine to an alternative location, the encrypted virtual machine is restored without encryption, and must be encrypted manually by using the vCenter Server after the restore operation is completed.

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

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

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

Adding a vCenter Server instance

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

Procedure

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

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **VMware**.
2. Click **Manage vCenter**.
3. Click **Add vCenter**.
4. Populate the fields in the **vCenter Properties** section:

© Copyright IBM Corp. 2017, 2021

Hostname/IP
Enter the resolvable IP address or a resolvable path and machine name.

Use existing user
Enable to select a previously entered user name and password for the vCenter Server instance.

Username
Enter your user name for the vCenter Server instance.

Password
Enter your password for the vCenter Server instance.

Port
Enter the communications port of the vCenter Server instance. Select the **Use TLS** check box to enable an encrypted Transport Layer Security (TLS) connection. The typical default port is 80 for non TLS connections or 443 for TLS connections.

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

   Maximum number of VMs to process concurrently per ESX server and per SLA
   Set the maximum number of concurrent VM snapshots to process on the ESX server. The default setting is 3.

6. Click **Save**. IBM Spectrum Protect Plus confirms a network connection, adds the vCenter Server instance to the database, and then catalogs the instance.

   If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connections.

What to do next

After you add a vCenter Server instance, complete the following action:

Action	How to
Assign user permissions to the hypervisor.	See “Creating a role” on page 650.

Related concepts

“Managing identities” on page 656
Some features in IBM Spectrum Protect Plus require credentials to access your resources. For example, IBM Spectrum Protect Plus connects to Oracle servers as the local operating system user that is specified during registration to complete tasks like cataloging, data protection, and data restore.

Related tasks

“Backing up VMware data” on page 340
Use a backup job to back up VMware resources such as virtual machines (VMs), datastores, folders, vApps, and datacenters with snapshots.

“Restoring VMware data” on page 350
VMware restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source.

Virtual machine privileges

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

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

vCenter Server Object	Required Privileges
Alarm	• Acknowledge alarm • Set alarm status
Cryptographic Operations (6.5 and 6.7)	• Add disk • Direct access • Encrypt • Encrypt new • Manage encryption policies
Datastore	• Allocate space • Browse datastore • Low level file operations • Remove datastore • Remove file • Update virtual machine files
Distributed switch	• Port configuration operation • Port setting operation
Folder	• Create folder
Global	• Cancel task • Manage custom attributes • Set custom attribute
Host > Configuration	• Storage partition configuration
vSphere Tagging (6.5, 6.7, and 7.0)	• Assign or Unassign vSphere Tag • Assign or Unassign vSphere Tag on Object (7.0) • Create vSphere Tag • Create vSphere Tag Category • Modify UsedBy Field for Category • Modify UsedBy Field for Tag
Network	• Assign network
Resource	• Apply recommendation • Assign a vApp to resource pool • Assign virtual machine to resource pool • Migrate powered off virtual machine • Migrate powered on virtual machine • Query vMotion

vCenter Server Object	Required Privileges
Virtual Machine > Configuration	• Acquire disk lease (6.7 and 7.0) • Add existing disk • Add new disk • Add or remove device • Advanced (6.5) • Advanced configuration (6.7 and 7.0) • Change CPU count • Change memory (6.7 and 7.0) • Change settings (6.7 and 7.0) • Configure raw device (6.7 and 7.0) • Disk change tracking (6.5) • Disk lease (6.5) • Memory (6.5) • Modify device settings • Raw device (6.5) • Reload from path • Remove disk • Rename • Settings (6.5) • Toggle disk change tracking (6.7 and 7.0)
Virtual Machine > Guest Operations	• Guest Operation Modifications • Guest Operation Program Execution • Guest Operation Queries
Virtual Machine > Interaction	• Backup operation on virtual machine • Power Off • Power On
Virtual Machine > Inventory	• Register • Remove • Unregister
Virtual Machine > Provisioning	• Allow disk access • Allow read-only disk access • Allow virtual machine download • Allow virtual machine files upload • Mark as template • Mark as virtual machine
Virtual Machine > Snapshot management	• Create snapshot • Remove snapshot • Revert snapshot

vCenter Server Object	Required Privileges
vApp	• Add virtual machine • Assign resource pool • Assign vApp • Create • Delete • Power Off • Power On • Rename • Unregister • vApp resource configuration

Detecting VMware resources

VMware resources are automatically detected after the vCenter Server instance is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the instance was added. If a virtual machine inventory job fails, subsequent attempts to run a backup job will also fail.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **VMware**.
2. In the list of vCenters Server instances, select an instance or click the link for the instance to navigate to the resource that you want. For example, if you want to run an inventory job for an individual virtual machine in the instance, click the instance link and then select a virtual machine.
3. Click **Run Inventory**.

Testing the connection to a vCenter Server virtual machine

You can test the connection to a vCenter Server virtual machine. The test function verifies communication with the virtual machine and tests domain name server (DNS) settings between the IBM Spectrum Protect Plus virtual appliance and the virtual machine.

Procedure

To test the connection, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **VMware**.
2. In the list of vCenters Server instances, click the link for a vCenter Server to navigate to the individual virtual machines.
3. Select a virtual machine, and then click **Select Options**.
4. Select **Use existing user**.
5. Select a user in the **Select user** list.
6. Click **Test**.

Backing up VMware data

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

Before you begin

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

- Register the providers that you want to back up. For more instructions, see "Adding a vCenter Server instance" on page 335.
- Configure SLA policies. For more instructions, see "Create backup policies" on page 259.
- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see Chapter 19, "Managing user access," on page 643.
- If a VM is associated with multiple SLA policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.
- If your vCenter is a VM, to help maximize data protection, have the vCenter on a dedicated datastore and backed up in a separate backup job.
- Ensure the latest version of VMware Tools is installed on VMware VMs.

About this task

- When backing up VMware VMs, IBM Spectrum Protect Plus downloads .vmx, .vmxf, and .nvram files if necessary, and then it transfers those files to the vSnap server as needed. For this to work successfully, the IBM Spectrum Protect Plus appliance must be able to resolve and access all protected ESXi hosts. When the appliance communicates with an ESXi host, the correct IP address must be returned.
- If a VM is protected by an SLA policy, the backups of the VM will be retained based on the retention parameters of the SLA policy, even if the VM is removed from vCenter.
- If an existing VM is migrated by a vMotion operation, IBM Spectrum Protect Plus will perform a rebase operation if necessary.

Restriction: File cataloging, backup, point-in-time restores, and other operations that invoke the Windows agent will fail if a non-default local administrator is entered as the **Guest OS Username** when defining a backup job. A non-default local administrator is any user that has been created in the guest OS and has been granted the administrator role.

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

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

Procedure

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

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **VMware**.
2. Select resources to back up.

Use the search function to search for available resources and toggle the displayed resources by using the **View** filter. Available options are **VMs and Templates**, **VMs**, **Datastore**, **Tags and Categories**, and **Hosts and Clusters**. Tags are applied in vSphere, and allow a user to assign metadata to VMs.

3. Click **Select SLA Policy** to add one or more SLA policies that meet your backup criteria to the job definition.
4. To create the job definition by using default options, click **Save**.

   The job will run as defined by the SLA policies that you selected. To run the job immediately, click **Jobs and Operations** > **Schedule**. Select the job and click **Actions** > **Start**.

   Tip: When the job for the selected SLA policy runs, all resources that are associated with that SLA policy are included in the backup operation. To back up only selected resources, you can run an on-demand job. An on-demand job runs the backup operation immediately.

   - To run an on-demand backup job for a single resource, select the resource and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
   - To run an on-demand backup job for one or more resources, click **Create job**, select **Ad hoc backup**, and follow the instructions in "Running an ad hoc backup job" on page 627.

   When the job definition is saved, available virtual machine disks (VMDKs) in a VM are discovered and are shown when **VMs and Templates** is selected in the **View** filter. By default, these VMDKs are assigned to the same SLA policy as the VM. If you want a more granular backup operation, you can exclude individual VMDKs from the SLA policy. For instructions, see "Excluding VMDKs from the SLA policy for a job" on page 344.
5. To edit options before you create the job definition, click **Select Options**.

   Tips for configuring options:

   Review the following tips to help you configure options for the backup job:

   - To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
   - If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
   - Options that are shown in yellow indicate that the option value has changed from the previously saved value.
   - To close the **Options** pane without saving changes, click **Select Options**.

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

   Skip Read-only datastores
   Skip datastores that are mounted as read-only.

   Skip temporary datastores mounted for Instant Access
   Exclude temporary Instant Access datastores from the backup job definition.

   VADP Proxy
   Select a VADP proxy to balance the load.

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

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

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

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

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

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

Truncate SQL logs
To truncate application logs for SQL Server databases that are on the VM during the backup job, enable the **Truncate SQL logs** option.

Restriction: It is possible that the same databases that are on a VM might be backed up as part of a VM backup job and a SQL Server backup job. Do not select this option if you want to back up the database transaction logs during the SQL Server backup operation. The log truncation deletes all inactive logs from the log file. The deleted log sequence causes discontinuity in the log backup.

For more information about backing up SQL Server database logs, see "Log backups" on page 605.

The credentials must be established for the associated VM by using the Guest OS user name and Guest OS Password option within the backup job definition. When the VM is attached to a domain, the user identity follows the default *domain**name* format. If the user is a local administrator, the format *local_administrator* is used.

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

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

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

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

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

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

Catalog file metadata
Turn on file indexing for the associated snapshot. When file indexing is completed, individual files can be restored by using the **File Restore** pane in IBM Spectrum Protect Plus. Credentials must be established for the associated VM by using an SSH key, or the **Guest OS Username** and **Guest OS Password** options within the backup job definition. Ensure that the VM can be accessed from the IBM Spectrum Protect Plus appliance either by using DNS or a host name.

Restriction: SSH Keys are not a valid authorization mechanism for Windows platforms.

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

Use existing user
Select a previously entered user name and password for the provider.

Guest OS Username/Password
For some tasks (such as cataloging file metadata, file restore, and IP reconfiguration), credentials must be established for the associated VM. Enter the user name and password, and ensure that the VM can be accessed from the IBM Spectrum Protect Plus appliance either by using DNS or a host name.

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

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

   Pre-scripts and Post-scripts

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

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

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

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

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

   Run inventory before backup
   Run an inventory job and capture the latest data of the selected resources before starting the backup job.

   Exclude Resources

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

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

   Separate multiple filters with a semicolon.

   Exclude Resources by Tag
   Exclude specific resources based on associated VM tags from the backup job. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*). Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *. Multiple filters may be separated with a semicolon.

   Force Full Backup of Resources
   Force base backup operations for specific VMs or databases in the backup job definition. Separate multiple resources with a semicolon.

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

What to do next

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

Action	How to
If you are using a Linux environment, consider creating VADP proxies to enable load sharing.	See “Creating VADP proxies” on page 346.
Create a VMware restore job definition.	See “Restoring VMware data” on page 350.

In some cases, VMware backup jobs fail with “failed to mount” errors. To resolve this issue, increase the maximum number of NFS mounts to at least 64 by using the NFS.MaxVolumes (vSphere 5.5 and later) and NFS41.MaxVolumes (vSphere 6.0 and later) values. Follow the instructions in Increasing the default value that defines the maximum number of NFS mounts on an ESXi/ESX host.

Related concepts

“Configuring scripts for backup and restore operations” on page 628
Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

“Starting jobs on demand” on page 621
You can run any job on demand, even if the job is set to run on a schedule.

Excluding VMDKs from the SLA policy for a job

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

Before you begin

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

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

Procedure

To exclude VMDKs from the SLA policy that has been applied to a virtual machine:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **VMware**.
2. Select **VMs and Templates** in the **View** filter.
3. Click the link for the vCenter that contains the virtual machines in the **Name** field. It may be necessary to select the datacenter and folder.
4. Identify the virtual machine that contains the VMDKs to be excluded. If an SLA policy is already applied, it will be visible in the **SLA Policy** field.

5. Click the virtual machine name for the virtual machine to display the associated VMDKs.
6. Select the VMDK that is to be excluded and then click **Select SLA Policy**.
7. Clear the selected SLA policy that is applied to the VMDK in the SLA Policy pane.

   Note: If multiple VMDKs are selected that are assigned to the same policy, the SLA policy will not appear as selected in the SLA Policy pane after clicking **Select SLA Policy**. Leave all SLA policies clear to exclude the VMDKs from the policy.
8. Click **Save**.

Managing VADP backup proxies

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

The backup of a VMware virtual machine includes the following files:

- VMDKs corresponding to all disks. The base backup captures all allocated data, or all data if disks are on NFS datastores. Incremental backups will capture only changed blocks since the last successful backup.
- Virtual machine templates.
- VMware files with the following extensions:
  - .vmx
  - .vmfx (if available)
  - .nvram (stores the state of the virtual machine BIOS)

At least one VADP proxy must be enabled in the backup site that is specified in the SLA for VMware backups. For more information, see “Creating VADP proxies” on page 346.

The processing load is shifted off the host system and onto the proxies for VMware backup jobs. When more than one VADP proxy exists, throttling ensures that multiple proxies are optimally utilized to maximize data throughput. For each VMware virtual machine being backed up, IBM Spectrum Protect Plus determines which VADP proxy is the least busy and has the most available memory and free tasks. Free tasks are determined by the number of available CPU cores or by using the **Softcap task limit** option.

If a proxy server goes down or is otherwise unavailable before the start of the job, the other proxies take over and the job is complete. If a proxy server becomes unavailable when a job is running, the job may fail.

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

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

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

- SAN transport mode is preferred in a direct storage environment because this mode is typically fast and reliable.
- HotAdd transport mode is preferred if the VADP proxy is virtualized. This mode supports all vSphere storage types.

  Note: To use only the HotAdd transport mode without falling back to alternate transport modes, select **VADP proxy uses only HotAdd transport mode** in **Global Preferences**. For more information, see “Configuring global preferences” on page 304.
- NBD or NBDSSL transport mode (LAN) is the fallback mode because it works in physical, virtual, and mixed environments. However, with this mode, the data transfer speed might be compromised if network connections are slow. NBDSSL mode is similar to NBD mode except that data transferred between the VADP proxy and the ESXi server is encrypted when using NBDSSL.

Creating VADP proxies

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

Before you begin

Review the IBM Spectrum Protect Plus system requirements in “VADP proxy requirements” on page 41.

Ensure that you have the required user permissions to work with VADP proxies. For instructions about managing VADP proxy permissions, see “Permission types ” on page 650.

Restriction: For running the steps to create VADP proxies, ensure that you have a user ID with the SYSADMIN role assigned. For more information about roles, see “Managing roles” on page 648.

Tip: The IBM Spectrum Protect Plus version of the VADP proxy installer includes Virtual Disk Development Kit (VDDK) version 6.5. This version of the VADP proxy installer provides the external VADP proxy support with vSphere 6.5.

Procedure

To create VMware VADP proxies, complete the following steps:

1. In the navigation panel, click **System Configuration** > **VADP Proxy**.
2. Click **Register Proxy**.
3. Complete the following fields in the **Install VADP Proxy** pane:

   Hostname/IP

   Enter the resolvable IP address or a resolvable path and machine name.

   Select a site

   Select a site to associate with the proxy.

   Use existing user

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

   Username

   Enter the user name for the VADP proxy server.

   Password

   Enter the password name for the VADP proxy server.
4. Click **Install**.
5. Click **Yes** on the confirm screen.
6. Repeat the previous steps for each proxy you want to create.

Results

The proxy is added to the **VADP Proxy** table. You can suspend, uninstall, unregister, or edit a proxy server by clicking the ellipses icon ••• to open the actions menu. Suspending a proxy prevents upcoming backup jobs from using the proxy, and jobs that use a suspended or unregistered proxy will run locally, which may impact performance. You can complete maintenance tasks on the proxy while it is suspended. To resume usage of the proxy, click the ellipses icon ••• to open the actions menu and click **Resume**. After successful creation, the service vadp is started on the proxy machine. A log file, `vadp.log`, is generated in `/opt/IBM/SPP/logs` directory.

The connection between the IBM Spectrum Protect Plus virtual appliance and a registered VADP proxy is a bidirectional connection that requires the IBM Spectrum Protect Plus virtual appliance to have connectivity to the VADP proxy, and the VADP proxy to have connectivity to the IBM Spectrum Protect Plus virtual appliance. To ensure a proper connection from the IBM Spectrum Protect Plus virtual

appliance to the VADP proxy, verify that the IBM Spectrum Protect Plus virtual appliance can ping the VADP proxy by completing the following steps:

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

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

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

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

What to do next

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

Action	How to
Run the VMware backup job.	See "Backing up VMware data" on page 340. The proxies are indicated in the job log by a log message similar to the following text: `Run remote vmdkbackup of MicroService: http://<proxy>` *`nodename`*`, IP:`*`proxy_IP_address`*

Related tasks

"Setting options for VADP proxies" on page 348
When you create VADP proxies in IBM Spectrum Protect Plus, you can configure various options for each VADP proxy.

Registering a VADP proxy on a vSnap server

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

Before you begin

One or more stand-alone vSnap servers must be properly deployed and configured in your environment and added to IBM Spectrum Protect Plus backup storage providers. For instructions, see "Registering a vSnap server as a backup storage provider" on page 158.

For the combined system requirements of a vSnap server and the VADP proxy, see VADP proxy on vSnap server requirements.

Ensure that you have the required user permissions to work with VADP proxies. For instructions about managing VADP proxy permissions, see "Permission types " on page 650.

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

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

Procedure

1. In the navigation panel, 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.

Setting options for VADP proxies

When you create VADP proxies in IBM Spectrum Protect Plus, you can configure various options for each VADP proxy.

Before you begin

Ensure that you have the required user permissions to work with VADP proxies. For instructions about managing VADP proxy permissions, see “Permission types ” on page 650.

Procedure

To set options for VMware VADP proxies, complete the following steps:

1. In the navigation panel, click **System Configuration** > **VADP Proxy**.
2. Click the VADP proxy that you want to configure, which then displays the information in the adjacent details pane.
3. In the VADP proxy details pane, click the ellipses icon ••• and then choose **Proxy Options**.
4. Complete the following fields in the **Set VADP Proxy Options** pane:

   Site

   Assign a site to the proxy.

   User

   Select a previously entered user name for the provider.

   Transport Modes (ordered list)

   Set the transport modes to be used by the proxy. The order in which each mode is selected will determine the order in which the transport modes are used. To remove a transport mode, click the delete icon beside the transport mode. For more information about VMware transport modes, see Virtual Disk Transport Methods.

   Enable NBDSSL Compression

   If you selected the NBDSSL transport mode, enable compression to increase the performance of data transfers. Available compression types include **libz**, **fastlz**, and **skipz**.

To turn off compression, select **disabled**.

Log retention in days

Set the number of days to retain logs before they are deleted.

Read and write buffer size

Set the buffer size of the data transfer, measured in bytes.

Block size of NFS volume

Set the block size to be used by the mounted NFS volume, measured in bytes.

Softcap task limit

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

1 CPU = 1 VMDK

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

What to do next

After setting the VADP proxy options, you can complete the following actions:

Action	How to
Run the VMware backup job.	See “Backing up VMware data” on page 340.
Uninstall the proxies when you cease running the VMware backup jobs.	See “Uninstalling VADP proxies” on page 349.

Related tasks

“Creating VADP proxies” on page 346
You can create VADP proxies to run VMware backup jobs with IBM Spectrum Protect Plus in Linux environments.

Uninstalling VADP proxies

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

Procedure

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

Note: This procedure only applies to VADP proxies that have been installed in the environment. It does not apply to the VADP proxy that is deployed with the IBM Spectrum Protect Plus appliance.

1. In the navigation panel, click on **System Configuration** > **VADP Proxy**.
2. Click the VADP proxy that you want to uninstall, which then displays the information in the adjacent details pane.
3. Click the ellipses icon ••• in the details pane and select **Uninstall**.

Updating the IBM Spectrum Protect Plus IP address for VADP proxies

You can use the IBM Spectrum Protect Plus user interface to deploy vStorage API for Data Protection (VADP) proxy servers. If the IP address of the IBM Spectrum Protect Plus server changes after VADP proxy deployment, the proxies in the environment lose contact with the IBM Spectrum Protect Plus server. To resolve this issue, a script is provided to update the IBM Spectrum Protect Plus address for orphaned VADP proxies.

Before you begin

About this task

Ensure that you have the new IP address of the IBM Spectrum Protect Plus server.

Procedure

1. If the VADP proxy server and vSnap server are co-deployed, log in by using the secure shell protocol (SSH) with the `serveradmin` user account. If you are using a stand-alone VADP proxy server, log in to the VADP proxy server with the `root` user account.
2. Run the `update_vadp.sh` script with the new IBM Spectrum Protect Plus server IP address as the only argument by taking one of the following actions:
   - If the VADP proxy is co-deployed with the vSnap server, run the following script as the `serveradmin` user:

     ```
     $ sudo /opt/IBM/SPP/bin/update_vadp.sh new_ip
     ```

   - If the VADP proxy server is stand-alone, run the script as the `root` user:

     ```
     # /opt/IBM/SPP/bin/update_vadp.sh new_ip
     ```

3. When the script completes processing, the following message is displayed:

   ```
   The new IP address ( <new_ip> ) has been successfully updated.
   ```

   The IBM Spectrum Protect Plus server IP address is updated on the VADP proxy server

Restoring VMware data

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

Before you begin

Complete the following tasks:

- Ensure that a VMware backup job was run at least once. For instructions, see “Backing up VMware data” on page 340.
- Ensure that appropriate roles are assigned to IBM Spectrum Protect Plus users so that they can complete backup and restore operations. Grant users access to hypervisors and backup and restore operations through the **Accounts** pane. For more information, see Chapter 19, “Managing user access,” on page 643 and “Managing user accounts” on page 653.
- Ensure that the destination that you plan to use for the restore job is registered in IBM Spectrum Protect Plus. This requirement applies to restore jobs that restore data to original hosts or clusters.
- When restoring a virtual machine by using clone mode and by using the original IP configuration, ensure that credentials are established through the **Guest OS Username** and **Guest OS Password** options within the backup job definition.

About this task

If a VMDK is selected for restore operation, IBM Spectrum Protect Plus automatically presents options for an Instant Disk restore job, which provides instant writable access to data and application restore points. An IBM Spectrum Protect Plus snapshot is mapped to a target server where it can be accessed or copied as required.

All other sources are restored through Instant VM restore jobs, which can be run in the following modes:

Production mode
Production mode enables disaster recovery at the local site from primary storage or a remote disaster recovery site, replacing original machine images with recovery images. All configurations are carried over as part of the recovery, including names and identifiers, and all copy data jobs associated with the virtual machine continue to run. As part of a production mode restore, you can opt to replace the storage in the virtual machine with a virtual disk from a previous virtual machine backup. Select

Test mode
Test mode creates temporary virtual machines for development or testing, snapshot verification, and disaster recovery verification on a scheduled, repeatable basis without affecting production environments. Test machines are kept running as long as needed to complete testing and verification and are then cleaned up. Through fenced networking, you can establish a safe environment to test your jobs without interfering with virtual machines used for production. Virtual machines that are created in test mode are also given unique names and identifiers to avoid conflicts within your production environment.

Clone mode
Clone mode creates copies of virtual machines for use cases that require permanent or long-running copies for data mining or duplication of a test environment in a fenced network. Virtual machines created in clone mode are also given unique names and identifiers to avoid conflicts within your production environment. With clone mode, you must be sensitive to resource consumption because clone mode creates permanent or long-term virtual machines.

Note: Restoring a VM in test mode or clone mode without selecting a fenced network might interfere with production VMs and network services. For instructions for creating a fenced network, see “Creating a fenced network through a VMware restore job” on page 357.

The size of a virtual machine that is restored from a vSnap copy to an IBM Spectrum Protect restore point will be equal to the thick provisioned size of the virtual machine, regardless of source provisioning due to the use of NFS datastores during the copy operation. The full size of the data must be transferred even if it is unallocated in the source virtual machine.

When you restore VMware data from an IBM Spectrum Protect archive, files initially will be migrated from tape to a staging pool. Depending on the size of the restore operation, this process could take several hours.

Restriction: Windows file indexing and file restore on volumes residing on dynamic disks is not supported.

Procedure

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

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **VMware** > **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **VMware**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

2. On the **Select source** page, take the following actions:
   a) Review the available sources, including virtual machines (VMs) and virtual disks (VDisks). Use the **View** filter to toggle the displayed sources to show hosts and clusters, VMs, or tags and categories. You can expand a source by clicking its name.

      You can also enter all or part of a name in the **Search for** box to locate VMs that match the search criteria. You can use the wildcard character (*) to represent all or part of a name. For example, `vm2*` represents all resources that begin with "vm2".

b) Click the plus icon  next to the item that you want to add to the restore list next to the list of sources. You can add more than one item of the same type (VM or virtual disk).

To remove an item from the restore list, click the minus icon  next to the item.

c) Click **Next**.

3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand
Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

Recurring
Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resource restore or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Demo** The demonstration site from which to restore snapshots. This menu item is available only if you updated the product from IBM Spectrum Protect Plus Version 10.1.6 or earlier. **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Set destination** page, specify the instance that you would like to restore for each chosen source and click **Next**:

Original Host or Cluster
Select this option to restore data to the original host or cluster.

Alternate Host or Cluster
Select this option to restore data to a local destination that is different from the original host or cluster, and then select the alternate location from the available resources. Test and production networks can be configured on the alternate location to create a fenced network, which keeps virtual machines used for testing from interfering with virtual machines used for production. From the **vCenters** section, select an alternative location. You can filter the alternative locations by either hosts or clusters.

In the **VM Folder Destination** field, enter the virtual machine folder path on the destination datastore. Note that the directory will be created if it does not exist. Use "/" as the root virtual machine folder of the targeted datastore.

When restoring a virtual disk to a new destination VM, select the virtual machine to which the virtual disk will be restored and the **Destination Disk Mode**. You can set the **Destination Controller** to select a supported SCSI controller. Changing the SCSI controller type replaces the existing controller with a new controller, applies the common settings of the existing controller to the new controller, and reassigns all SCSI devices to the new controller. Optionally, can also set the **Destination Controller Address #** and **Destination Controller LUN #** to select specific controllers or LUNs.

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

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

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

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

To restore data to an ESXi host, the host must have a standard switch or a distributed switch with ephemeral binding. Review the information in “Restoring data when vCenter Server or other management VMs are not accessible” on page 359 to ensure that you have the correct environment configured to use this option.

6. On the **Set datastore** page, take the following actions:
   - If you are restoring data to an alternate ESXi host or cluster, select the destination datastore and click **Next**.
   - If you are restoring data to the original ESXi host or cluster, this page is not displayed.
7. On the **Set network** page, specify the network settings to use for each chosen source and click **Next**.
   - If you are restoring data to the original ESXi host or cluster, specify the following network settings:

     Allow system to define IP configuration
     Select this option to allow your operating system to define the destination IP address. During a test mode restore operation, the destination virtual machine receives a new MAC address along with an associated NIC. Depending on your operating system, a new IP address can be assigned based on the original NIC of the virtual machine, or assigned through DHCP. During a production mode restore, the MAC address does not change; therefore, the IP address should be retained.

     Use original IP configuration
     Select this option to restore data to the original host or cluster using your predefined IP address configuration. During the restore operation, the destination virtual machine receives a new MAC address, but the IP address is retained.
   - If you are restoring data to an alternate ESXi host or cluster, complete the following steps:

a. In the **Production** and **Test** fields, set virtual networks for production and test restore job runs. Destination network settings for production and test environments should point to different locations to create a fenced network, which keeps virtual machines used for testing from interfering with virtual machines used for production. The networks that are associated with test and production modes will be used when the restore job is run in the associated mode.

b. Set an IP address or subnet mask for virtual machines to be repurposed for development, testing, or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet. Virtual machines that contain multiple NICs are supported.

   Take one of the following actions:

   – To allow your operating system to define the destination subnets and IP addresses, click **Use system defined subnets and IP addresses for VM guest OS on destination**.

   – To use your predefined subnets and IP addresses, click **Use original subnets and IP addresses for VM guest OS on destination**.

   – To create a new mapping configuration, select **Add mappings for subnets and IP addresses for VM guest OS on destination**, click **Add Mapping**, and enter a subnet or IP address in the **Add Source Subnet or IP Address** field.

     Choose one of the following network protocols:

     - Select **DHCP** to automatically select an IP and related configuration information if DHCP is available on the selected source.
     - Select **Static** to enter a specific subnet or IP address, subnet mask, gateway, and DNS. The **Subnet / IP Address**, **Subnet Mask**, and **Gateway** are required fields. If a subnet is entered as a source, a subnet must also be entered as a destination.

       Note: When a mapping is added, the source IP address must be entered into the field by the **+** button. The destination IP address information should be entered into the **Subnet / IP Address, Subnet Mask,** and **Gateway** fields. Re-addressing can only be performed on machines with VMware Tools installed prior to executing the backup job that is to be restored.

     IP reconfiguration is skipped for virtual machines if a static IP is used but no suitable subnet mapping is found, or if the source virtual machine is powered off and there is more than one associated NIC. In a Windows environment, if a virtual machine uses DHCP only, then IP reconfiguration is skipped for that virtual machine. In a Linux environment, all addresses are assumed to be static, and only IP mapping will be available.

8. On the **Restore methods** page, select the restore method to be used for source selection. Set the VMware restore job to run in **Production**, **Test**, or **Clone**. When you run in production mode, you have the option to select **Replace virtual disks only and retain virtual machine configuration**. Enabling this option replaces the storage in the virtual machine with the virtual disks from a previous virtual machine backup. This restore method maintains the virtual machine configuration replacing only the storage. When production restore is selected, the virtual machine to which the disk replace is applied must be powered off and only restoring to the original location is supported. Additionally, options such as overwriting the virtual machine and restoring based on tags are not available because the virtual machine is not being recreated. After the job is created, it can be run in production or clone mode through the **Job Sessions** pane.
9. You can also change the name of the restored VM by entering the new VM name in the **Rename VM (optional)** field. The **Rename VM** option is not available when the production mode is used with the replace virtual disks option enabled. Click **Next** to continue.
10. If you are running the restore job in advanced mode, you can set additional options as follows:

    Power on after recovery
    Toggle the power state of a virtual machine after a recovery is run. Virtual machines are powered on in the order in which they are recovered, as set in the Source step. If **Use original IP configuration** is selected, the **Power on after recovery** option is not honored.

    Restriction: Restored virtual machine templates cannot be powered on after recovery.

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

Continue with restore even if it fails
Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the restore job stops if the recovery of a resource fails.

Run cleanup immediately on job failure
This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow to overwrite and force cleanup of pending old sessions
Enable this option to allow a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Restore VM tags
Enable this option to restore tags that are applied to virtual machines through vSphere.

Enable Streaming (VADP) restore
Parallel streaming for virtual machine restore operations is set by default. You can deselect this option for virtual machine restore operations. VADP proxy selection for streaming restore is based on the site of the source snapshot being restored.

Tip: When you are restoring virtual machines managed by a VMware Cloud (VMC) on AWS Software-Defined Data Center (SDDC), this option should always be enabled to allow streaming of the data.

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

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

11. Optional: On the **Apply scripts** page, choose the following script options and click **Next**.
    - Select **Pre-script** to select an uploaded script, and an application or script server where the prescript runs. To select an application server where the script will run, clear the **Use Script Server** check box. Go to the **System Configuration** > **Script** page to configure scripts and script servers.
    - Select **Post-script** to select an uploaded script and an application or script server where the postscript runs. To select an application server where the script runs, clear the **Use Script Server** check box. Navigate to the **System Configuration** > **Script** page to configure scripts and script servers.
    - Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails. When this option is enabled and the prescript completes with a nonzero return code, the backup or restore job continues to run and the prescript task status returns `COMPLETED`. If a postscript completes with a nonzero return code, the postscript task status returns `COMPLETED`. When this option is not selected, the backup or restore job does not run, and the prescript or postscript task status returns with a `FAILED` status.
12. Take one of the following actions on the **Schedule** page:
    - To run an on-demand job, click **Next**.
    - To set up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
13. On the **Review** page, review your restore job settings and click **Submit** to create the job.

    On-demand jobs will begin immediately; recurring jobs will begin at the scheduled start time.

What to do next

After the job is completed, select one of the following options from the **Actions** menu on the Jobs Sessions or Active Clones sections in the **Restore** pane:

Cleanup

Destroys the virtual machine and cleans up all associated resources. Because this is a temporary virtual machine to be used for testing, all data is lost when the virtual machine is destroyed.

Move to Production (vMotion)

Migrates the virtual machine through vMotion to the datastore and the virtual Network defined as the production network.

Clone (vMotion)

Migrates the virtual machine through vMotion to the datastore and virtual Network defined as the test network.

Related tasks

"Adding a vCenter Server instance" on page 335
When a vCenter Server instance is added to IBM Spectrum Protect Plus, an inventory of the instance is captured, enabling you to complete backup and restore jobs, as well as run reports.

Creating a fenced network through a VMware restore job

Through fenced networking, you can establish a safe environment to test your jobs without interfering with virtual machines that are used for production. Fenced networking can be used with jobs that are running in test mode and production mode.

Before you begin

Create and run a VMware Restore job. For instructions, see "Restoring VMware data" on page 350.

Procedure

To create a fenced network, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **VMware**.
2. In the **Restore** pane, review the available restore points of your VMware sources, including virtual machines, VM templates, datastores, folders, and vApps. Use the search function and filters to fine-tune your selection across specific recovery site types. Expand an entry in the **Restore** pane to view individual restore points by date.
3. Select restore points and click the add to restore list icon to add the restore point to the Restore List. Click the remove icon to remove items from the Restore List.
4. Click **Options** to set the job definition options.
5. Select **Alternate ESX Host or Cluster**, then select an alternate host or cluster from the vCenter list.
6. Expand the **Network Settings** section. From the **Production** and **Test** fields, set virtual networks for production and test Restore job runs. Destination network settings for production and test environments should be different locations to create a fenced network, which keeps virtual machines used for testing from interfering with virtual machines used for production. The networks associated with Test and Production will be utilized when the restore job is run in the associated mode. The IP addresses of the target machine can be configured by using the following options:

   Use system defined subnets and IP addresses for VM guest OS on destination

   Select to allow your operating system to define the destination IP address. During a Test Mode restore, the destination virtual machine receives a new MAC address along with an associated NIC. Depending on your operating system, a new IP address can be assigned based on the original NIC of the virtual machine, or assigned through DHCP. During a Production Mode restore operation the MAC address does not change; therefore, the IP address should be retained.

Use original subnets and IP addresses for VM guest OS on destination
Select to restore to the original host or cluster using your predefined IP address configuration. During a restore, the destination virtual machine receives a new MAC address, but the IP address is retained.

Set the network settings for a restore to an alternate or long distance ESX host or cluster:

From the **Production** and **Test** fields, set virtual networks for production and test restore job runs. Destination network settings for production and test environments should be different locations to create a fenced network, which keeps virtual machines used for testing from interfering with virtual machines used for production. The networks associated with Test and Production will be utilized when the restore job is run in the associated mode.

Set an IP address or subnet mask for virtual machines to be re-purposed for development/testing or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet. Virtual machines containing multiple NICs are supported.

By default, the **Use system defined subnets and IP addresses for VM guest OS on destination** option is enabled. To use your predefined subnets and IP addresses, select **Use original subnets and IP addresses for VM guest OS on destination**.

To create a new mapping configuration, select **Add mappings for subnets and IP addresses for VM guest OS on destination**, then click **Add Mapping**. Enter a subnet or IP address in the **Source** field. In the destination field, select **DHCP** to automatically select an IP and related configuration information if DHCP is available on the selected client. Select **Static** to enter a specific subnet or IP address, subnet mask, gateway, and DNS. Note that **Subnet / IP Address**, **Subnet Mask**, and **Gateway** are required fields. If a subnet is entered as a source, a subnet must also be entered as a destination.

IP reconfiguration is skipped for virtual machines if a static IP is used but no suitable subnet mapping is found, or if the source machine is powered off and there is more than one associated NIC. In a Windows environment, if a virtual machine is DHCP only, then IP reconfiguration is skipped for that virtual machine. In a Linux environment all addresses are assumed to be static, and only IP mapping will be available.

Destination Datastore
Set the destination datastore for a restore to an alternate ESX host or cluster.

VM Folder Destination
Enter the VM folder path on the destination datastore. Note that the directory will be created if it does not exist. Use "/" as the root VM folder of the targeted datastore.

7. Click **Save** to save the policy options.
8. After the job is complete, select one of the following options from the **Actions** menu on the Jobs Sessions or Active Clones sections on the **Restore** pane:

Cleanup

Destroys the virtual machine and cleans up all associated resources. Since this is a temporary/testing virtual machine, all data is lost when the virtual machine is destroyed.

Move to Production (vMotion)

Migrates the virtual machine through vMotion to the Datastore and the Virtual Network defined as the "Production" Network.

Clone (vMotion)

Migrates the virtual machine through vMotion to the Datastore and Virtual Network defined as the "Test" network.

Related tasks

“Adding a vCenter Server instance” on page 335

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

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

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

Before you begin

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

About this task

vCenter Server uses the following components:

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

To use the **ESX host if vCenter is down** option, the ESXi host must have a standard switch or a distributed switch. The distributed switch must have ephemeral binding. If one or both of these switches are available, you can run a restore operation in IBM Spectrum Protect Plus with the option enabled as described in "Restoring VMware data" on page 350 and no further manual configuration is required.

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

Procedure

1. Connect to the destination ESXi host user interface and create a standard virtual switch.

   The new switch has no port groups or uplinks.
2. Use the Secure Shell (SSH) protocol to connect to the ESXi host.
3. List the distributed switches that are configured on the ESXi host by issuing the following command:

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

   ```
   #esxcfg-vswitch -Q physical_vnic -V port_group switch_name
   ```
6. Add the physical NIC and port group to the new standard switch by issuing the following command:

   ```
   #esxcli network vswitch standard uplink add –-uplink-name=physical_vnic --vswitch-
   name=new_standard_vswitch
   ```
7. In the ESXi host user interface, add a temporary port group and select the standard switch that you created in step "1" on page 359.

   The standard switch has one port group and one uplink.
8. Run a restore operation in IBM Spectrum Protect Plus with the **ESX host if vCenter is down** option enabled.

   For instructions about running a restore operation, see "Restoring VMware data" on page 350.
9. In the ESXi host user interface for the ESXi host, power on the VMs that are restored.

10. Log in to the vCenter Server user interface and start the migration of the management VMs from the temporary port group that you created in step “7” on page 359 to an available distributed port group.
11. After all of the VMs are migrated to the original port group, reincorporate the physical NIC and the port group into the original distributed switch by taking the following actions. For example purposes, the following commands reference a virtualized Network Interface Card (VNIC) named vmnic0 that is part of port group 64.
    a. Remove the network cards (known as vmnics) from a standard switch by issuing the following command:

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

    For example:

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

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

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

    For example:

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

12. Delete the temporary port group and the standard switch from the ESXi host user interface.
13. After the VMs are migrated and accessible, use the ESXi host user interface to unregister, but not delete, the old VMs if the original host is reachable.

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

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

Backing up and restoring Hyper-V data

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

Ensure that your Hyper-V environment meets the system requirements in “Hypervisor (Microsoft Hyper-V and VMware) and cloud instance (Amazon EC2) backup and restore requirements ” on page 57.

Adding a Hyper-V server

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

Before you begin

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

- Hyper-V servers can be registered using a DNS name or IP address. DNS names must be resolvable by IBM Spectrum Protect Plus. If the Hyper-V server is part of a cluster, all nodes in the cluster must be resolvable through DNS. If DNS is not available, the server must be added to the `/etc/hosts` file on the IBM Spectrum Protect Plus appliance. If more than one Hyper-V server is set up in a cluster

environment, all of the servers must be added to `/etc/hosts`. When registering the cluster in IBM Spectrum Protect Plus, register the Failover Cluster Manager.

- All Hyper-V servers, including cluster nodes, must have the Microsoft iSCSI initiator Service running in their Services list. Set the service to Automatic so that it is available when the machine boots.
- Add the user to the local administrator group on the Hyper-V server.

Procedure

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

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **Hyper-V**.
2. Click **Manage Hyper-V Server**.
3. Click **Add Hyper-V Server**.
4. Populate the fields in the **Server Properties** pane:

   Hostname/IP
   Enter the resolvable IP address or a resolvable path and machine name.

   Use existing user
   Enable to select a previously entered user name and password for the server.

   Username
   Enter your user name for the server.

   Password
   Enter your password for the server.

   Port
   Enter the communications port of the server you are adding. The typical default port is 5985.

   Select the **Use TLS** check box to enable an encrypted Transport Layer Security (TLS) connection.

   If you do not select **Use TLS**, you must complete additional steps on the Hyper-V server. See "Enabling WinRM for connection to Hyper-V servers" on page 362.
5. In the **Options** section, configure the following option:

   Maximum number of VMs to process concurrently per Hyper-V server
   Set the maximum number of concurrent virtual machine snapshots to process on the Hyper-V server.
6. Click **Save**. IBM Spectrum Protect Plus confirms a network connection, adds the server to the database, and then catalogs the server.

   If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

What to do next

After you add the Hyper-V server, complete the following action:

Action	How to
Assign user permissions to the hypervisor.	See "Creating a role" on page 650.

Related tasks

"Backing up Hyper-V data" on page 363
Use a backup job to back up Hyper-V data with snapshots.

"Restoring Hyper-V data" on page 367

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

Enabling WinRM for connection to Hyper-V servers

If you cannot use TLS to enable encrypted network traffic between IBM Spectrum Protect Plus Hyper-V servers, you must configure WinRM on the host to allow unencrypted network traffic. Ensure that you understand the security risks that are associated with allowing unencrypted network traffic.

Procedure

To configure WinRM for connection to Hyper-V hosts:

1. On the Hyper-V host system, log in with an administrator account.
2. Open a Windows command prompt. If User Account Control (UAC) is enabled, you must open the command prompt with elevated privileges by running with the **Run as administrator** option enabled.
3. Enter the following command to configure WinRM to allow unencrypted network traffic:

```
winrm s winrm/config/service @{AllowUnencrypted="true"}
```

4. Verify that the `AllowUnencrypted` option is set to true through the following command:

```
winrm g winrm/config/service
```

Detecting Hyper-V resources

Hyper-V resources are automatically detected after the Hyper-V server is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the server was added. If a virtual machine inventory job fails, subsequent attempts to run a backup job will also fail.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **Hyper-V**.
2. In the list of Hyper-V servers, select a server or click the link for the server to navigate to the resource that you want. For example, if you want to run an inventory job for an individual virtual machine in a server, click the server link and then select a virtual machine.
3. Click **Run Inventory**.

Testing the connection to a Hyper-V Server virtual machine

You can test the connection to Hyper-V Server virtual machine. The test function verifies communication with the virtual machine and tests DNS settings between the IBM Spectrum Protect Plus virtual appliance and the virtual machine.

Procedure

To test the connection, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **Hyper-V**.
2. In the list of Hyper-V Servers, click the link for a Hyper-V Server virtual machine to navigate to the individual virtual machines.
3. Select a virtual machine, and then click **Select Options**.
4. Select **Use existing user**.
5. Select a user in the **Select user** list.
6. Click **Test**.

Backing up Hyper-V data

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

Before you begin

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

- Register the providers that you want to back up. For more information see “Adding a Hyper-V server” on page 360
- Configure SLA policies. For instructions, see “Create backup policies” on page 259.
- Hyper-V Backup and Restore jobs require the installation of the latest Hyper-V integration services.

  For Microsoft Windows environments, see Supported Windows guest operating systems for Hyper-V on Windows Server.

  For Linux environments, see Supported Linux and FreeBSD virtual machines for Hyper-V on Windows.
- IBM Spectrum Protect Plus uses Resilient Change Tracking (RCT) for tracking the changed blocks of Hyper-V virtual machine disks. For more information, see Resilient Change Tracking.
- All Hyper-V servers, including cluster nodes, must have the Microsoft iSCSI initiator Service running in their Services list. Set the service to Automatic so that it is available when the machine boots.
- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see Chapter 19, “Managing user access,” on page 643.
- If a virtual machine (VM) is associated with multiple SLA Policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA Policies to run with a significant amount of time between them, or combine them into a single SLA policy.
- If the IP address of the IBM Spectrum Protect Plus appliance is changed after an initial Hyper-V base backup is created, the target IQN of the Hyper-V resource may be left in a bad state. To correct this issue, from the Microsoft iSCSI Initiator tool, click the **Discovery** tab. Select the old IP address, then click **Remove**. Click the **Target** tab and disconnect the reconnecting session.
- If a VM is protected by an SLA policy, the backups of the VM will be retained based on the retention parameters of the SLA policy, even if the VM is removed.

About this task

Restriction: File cataloging, backup, point-in-time restores, and other operations that invoke the Windows agent will fail if a non-default local administrator is entered as the **Guest OS Username** when defining a backup job. A non-default local administrator is any user that has been created in the guest OS and has been granted the administrator role.

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

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

Procedure

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

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **Hyper-V**.
2. Select resources to back up.

Use the search function to search for available resources and toggle the displayed resources through the **View** filter. Available options are **VMs** and **Datastore**.

3. Click **Select SLA Policy** to add one or more SLA policies that meet your backup criteria to the job definition.
4. To create the job definition by using default options, click **Save**.

   The job runs as defined by the SLA policies that you selected. To run the job manually, click **Jobs and Operations** > **Schedule**. Select the job and click **Actions** > **Start**.

   Tip: When the job for the selected SLA policy runs, all resources that are associated with that SLA policy are included in the backup operation. To back up only selected resources, you can run an on-demand job. An on-demand job runs the backup operation immediately.

   - To run an on-demand backup job for a single resource, select the resource and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
   - To run an on-demand backup job for one or more resources, click **Create job**, select **Ad hoc backup**, and follow the instructions in "Running an ad hoc backup job" on page 627.
5. To edit options before you start the job, click the edit icon in the table **Select Options**.

   Tips for configuring options:

   Review the following tips to help you configure options for the backup job:

   - To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
   - If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
   - Options that are shown in yellow indicate that the option value has changed from the previously saved value.
   - To close the **Options** pane without saving changes, click **Select Options**.

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

   Skip Read-only datastores

   Enable to skip datastores mounted as read-only.

   Skip temporary datastores mounted for Instant Access

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

   Priority

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

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

   Make VM snapshot application/file system consistent

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

   VM Snapshot retry attempts

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

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

   Truncate SQL logs

   To truncate application logs for SQL during the Backup job, enable the **Truncate SQL logs** option.

   Restriction: It is possible that the same databases that are on a VM might be backed up as part of a VM backup job and a SQL Server backup job. Do not select this option if you want to back up

the database transaction logs during the SQL Server backup operation. The log truncation deletes all inactive logs from the log file. The deleted log sequence causes discontinuity in the log backup.

For more information about backing up SQL Server database logs, see "Log backups" on page 605.

The credentials must be established for the associated VM through the Guest OS Username and Guest OS Password option within the backup job definition. The user identity follows the default *domain**name* format if the VM is attached to a domain. The format *local_administrator* is used if the user is a local administrator.

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

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

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

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

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

Catalog file metadata

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

Exclude Files

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

Use existing user

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

Guest OS Username/Password

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

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

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

   The **Test** function verifies communication with the VM and tests DNS settings between the IBM Spectrum Protect Plus appliance and the VM. To test a connection, select a single VM, then click **Select Options**. Select **Use existing user** and select a previously entered user name and password for the resource, and then click **Test**.
7. Click **Save**.

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

   Pre-scripts and Post-scripts

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

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

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

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

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

   Run inventory before backup

   Run an inventory job and capture the latest data of the selected resources before starting the backup job.

   Exclude Resources

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

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

   Separate multiple filters with a semicolon.

   Force Full Backup of Resources

   Force base backup operations for specific VMs or databases in the backup job definition. Separate multiple resources with a semicolon.
9. To save any additional options that you configured, click **Save**.

What to do next

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

Action	How to
Create a Hyper-V restore job definition.	See “Restoring Hyper-V data” on page 367.

Related concepts

“Configuring scripts for backup and restore operations” on page 628
Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

“Starting jobs on demand” on page 621

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

Excluding virtual disks from the SLA policy for a job

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

Before you begin

Excluding one or more virtual disks from a backup operation can impact the success of recovery. Consider the following information prior to excluding a virtual disk from a Hyper-V backup operation.

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

Procedure

To exclude virtual disks from the SLA policy that has been applied to a virtual machine:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **Hyper-V**.
2. Select **VMs** in the **View** filter.
3. Click the link for the Hyper-V server that contains the virtual machine in the **Name** field.
4. Identify the virtual machine that contains the virtual disks to be excluded. If an SLA policy is already applied, it will be visible in the **SLA Policy** field.
5. Click the virtual machine name for the virtual machine to display the associated virtual disks.
6. Select the virtual disk that is to be excluded and then click **Select SLA Policy**.
7. Clear the selected SLA policy that is applied to the virtual disk in the SLA Policy pane.

   Note: If multiple virtual disks are selected that are assigned to the same policy, the SLA policy will not appear as selected in the SLA Policy pane after clicking **Select SLA Policy**. Leave all SLA policies clear to exclude the virtual disks from the policy.
8. Click **Save**.

Restoring Hyper-V data

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

Before you begin

Complete the following tasks:

- Ensure that a Hyper-V backup job was run at least once. For instructions, see “Backing up Hyper-V data” on page 363.
- Ensure that the destination that you plan to use for the restore job is registered in IBM Spectrum Protect Plus. This requirement applies to restore jobs that restore data to original hosts or clusters.
- Ensure that the latest Hyper-V integration services are installed.

For Microsoft Windows environments, see Supported Windows guest operating systems for Hyper-V on Windows Server.

For Linux environments, see Supported Linux and FreeBSD virtual machines for Hyper-V on Windows.

- Ensure that the appropriate roles for restore operations are assigned to the affected users. Grant users access to hypervisors and backup and restore operations in the **Accounts** pane. Roles and associated permissions are assigned during user account creation. For instructions, see Chapter 19, "Managing user access," on page 643 and "Managing user accounts" on page 653.
- Windows file indexing and file restore on volumes residing on dynamic disks is not supported.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.
- When restoring a virtual machine by using clone mode and by using the original IP configuration, ensure that credentials are established through the **Guest OS Username** and **Guest OS Password** options within the backup job definition.

About this task

If a Virtual Hard Disk (VHDX) is selected for a restore job, IBM Spectrum Protect Plus automatically presents options for an Instant Disk Restore job, which provides instant writable access to data and application restore points.

An IBM Spectrum Protect Plus snapshot is mapped to a target server where the snapshot can be accessed or copied as required. All other sources are restored by using Instant VM restore jobs, which can be run in the following modes:

Production mode
Production mode enables disaster recovery at the local site from primary storage or a remote disaster recovery site, replacing original machine images with recovery images. All configurations are carried over as part of the recovery, including names and identifiers, and all copy data jobs that are associated with the virtual machine continue to run.

Test mode
Test mode creates temporary virtual machines for development, testing, snapshot verification, and disaster recovery verification on a scheduled, repeatable basis without affecting production environments. Test machines are kept running while they are needed to complete testing and verification and are then cleaned up. Through fenced networking, you can establish a safe environment to test your jobs without interfering with virtual machines that are used for production. Virtual machines that are created in test mode are also given unique names and identifiers to avoid conflicts within your production environment.

Clone mode
Clone mode creates copies of virtual machines for use cases that require permanent or long-running copies for data mining or duplication of a test environment in a fenced network. Virtual machines that are created in clone mode are also given unique names and identifiers to avoid conflicts within your production environment. With clone mode, you must be sensitive to resource consumption because clone mode creates permanent or long-term virtual machines.

Restriction: Moving from test mode to production mode is not supported for Hyper-V.

Procedure

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

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **Hyper-V** > **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **Hyper-V**.

- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
- The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

2. On the **Select source** page, take the following actions:
   a) Review the available sources, including virtual machines (VMs) and virtual disks (VDisks). You can expand a source by clicking its name.

      You can also enter all or part of a name in the **Search for** box to locate VMs that match the search criteria. You can use the wildcard character (*) to represent all or part of a name. For example, `vm2*` represents all resources that begin with "vm2".

   b) Click the plus icon next to the item that you want to add to the restore list next to the list of sources. You can add more than one item of the same type (VM or virtual disk).

      To remove an item from the restore list, click the minus icon next to the item.

   c) Click **Next**.
3. On the **Source snapshot** page, select the type of restore job that you want to create:

   On-demand
   Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

   Recurring
   Creates a repeating point-in-time restore job that runs on a schedule.
4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.

Option	Description
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resource restore or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Demo** The demonstration site from which to restore snapshots. This menu item is available only if you updated the product from IBM Spectrum Protect Plus Version 10.1.6 or earlier. **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.

Option	Description
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Set destination** page, choose the instance to be restored for the selected source and click **Next**:

   Original Host or Cluster
   Select this option to restore data to the original host or cluster.

   Alternate Host or Cluster
   Select this option to restore data to a local destination that is different from the original host or cluster, then select the alternative location from the available resources.

   In the **VM Folder Destination** field, enter the virtual machine folder path on the destination datastore. Note that the directory will be created if it does not exist. Use "/" as the root virtual machine folder of the targeted datastore.
6. On the **Set datastore** page, take the following actions:
   - If you are restoring data to an alternate Hyper-V host or cluster, select the destination datastore and click **Next**.
   - If you are restoring data to the original Hyper-V host or cluster, this page is not displayed.
7. On the **Set network** page, specify the network settings to use for each chosen source and click **Next**.
   - If you are restoring data to the original Hyper-V host or cluster, specify the following network settings:

     Allow system to define IP configuration
     Select this option to allow your operating system to define the destination IP address. During a test mode restore operation, the destination virtual machine receives a new MAC address along with an associated NIC. Depending on your operating system, a new IP address can be assigned based on the original NIC of the virtual machine, or assigned through DHCP. During a production mode restore the MAC address does not change; therefore the IP address should be retained.

     Use original IP configuration
     Select this option to restore to the original host or cluster using your predefined IP address configuration. During the restore operation, the destination virtual machine receives a new MAC address, but the IP address is retained.
   - If you are restoring data to an alternate Hyper-V host or cluster, complete the following step:

     In the **Production** and **Test** fields, set virtual networks for production and test restore job runs. Destination network settings for production and test environments should point to different locations to create a fenced network, which keeps virtual machines used for testing from interfering with virtual machines used for production. The networks that are associated with test and production modes will be used when the restore job is run in the associated mode.
8. On the **Restore methods** page, select the restore method to be used for source selection. Set the Hyper-V restore job to run in **Production**, **Test**, or **Clone**. When you run in production mode, you have the option to select **Replace virtual disks only and retain virtual machine configuration**. Enabling this option replaces the storage in the virtual machine with the virtual disks from a previous virtual machine backup. This restore method maintains the virtual machine configuration replacing only the storage. When production restore is selected, the virtual machine to which the disk replace is applied must be powered off and only restoring to the original location is supported. Additionally, options

such as overwriting the virtual machine and restoring based on tags are not available because the virtual machine is not being recreated. After the job is created, it can be run in production or clone mode through the **Job Sessions** pane.

9. You can also change the name of the restored VM by entering the new VM name in the **Rename VM (optional)** field. The **Rename VM** option is not available when the production mode is used with the replace virtual disks option enabled. Click **Next** to continue.
10. Optional: On the **Job Options (optional)** page, configure advanced options and click **Next**.

    Make IA clone resource permanent
    Enable this option to move the virtual disk to permanent storage and clean up temporary resources. This action is accomplished by starting a Live Migration operation for the resources in the background. The destination of the Live Migration operation is the VM Configuration Datastore. The Instant Access disk is still available for read/write operations during this operation.

    Power on after recovery
    Toggle the power state of a virtual machine after a recovery is run. Virtual machines are powered on in the order in which they are recovered, as set in the Source step. If **Use original IP configuration** is selected, the **Power on after recovery** option is not honored.

    Restriction: Restored virtual machine templates cannot be powered on after recovery.

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

    Continue with restore even if it fails
    Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the restore job stops if the recovery of a resource fails.

    Run cleanup immediately on job failure
    This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

    Allow to overwrite and force cleanup of pending old sessions
    Enable this option to allow a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

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

    Prepend prefix to virtual machine name
    Enter a prefix to add to the names of restored virtual machines. Click Save to save the policy options.
11. Optional: On the **Apply scripts** page, choose the following script options and click **Next**.

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

12. Take one of the following actions on the **Schedule** page:
    - To run an on-demand job, click **Next**.
    - To set up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
13. On the **Review** page, review your restore job settings and click **Submit** to create the job.

    On-demand jobs will begin immediately; recurring jobs will begin at the scheduled start time.

What to do next

After the job is complete, select one of the following options from the **Actions** menu on the **Jobs Sessions** or **Active Clones** sections on the **Restore** pane:

Cleanup
Destroys the virtual machine and cleans up all associated resources. Because this is a temporary virtual machine to be used for testing, all data is lost when the virtual machine is destroyed.

Clone (migrate)
Migrates the virtual machine to the datastore and virtual network that are defined as the test network.

Related tasks

"Backing up Hyper-V data" on page 363
Use a backup job to back up Hyper-V data with snapshots.

"Adding a Hyper-V server" on page 360
When a Hyper-V server is added to IBM Spectrum Protect Plus, an inventory of the server is captured, enabling you to complete backup and restore jobs, as well as run reports.

Backing up and restoring Amazon EC2 data

To protect Amazon EC2 data, first add an account for your EC2 instances in IBM Spectrum Protect Plus, and then create jobs for backup and restore operations for those instances.

To add an EC2 account to IBM Spectrum Protect Plus, access keys are required. Access keys are long-term credentials for an Identity and Access Management (IAM) user or the Amazon Web Services (AWS) account root user.

For information about how to create an IAM user with access keys and the permissions that are required for IBM Spectrum Protect Plus, see "Creating an AWS IAM user " on page 373.

For increased security, is recommended that the AWS account root user is not used for IBM Spectrum Protect Plus. For more information about the root user, refer to the AWS Identity and Access Management User Guide.

EC2 data is stored in Amazon Web Services (AWS) Elastic Block Store (EBS) snapshots rather than the vSnap server. IBM Spectrum Protect Plus manages these snapshots for backup and restore operations.

Ensure that your EC2 environment meets the system requirements in "Hypervisor (Microsoft Hyper-V and VMware) and cloud instance (Amazon EC2) backup and restore requirements " on page 57.

Creating an AWS IAM user

To complete tasks in the IBM Spectrum Protect Plus user interface, IAM users must have access keys and required permissions.

About this task

You can use the AWS Management Console to create an IAM user by using the following steps. These steps are condensed from the steps that are documented in the AWS Identity and Access Management User Guide to show settings that are required for IBM Spectrum Protect Plus. For the complete and detailed steps for creating an IAM user, refer to this guide.

To create a user, you muse have IAM administrative permissions.

Procedure

1. Sign in to the AWS Management Console and click **Services** > **IAM** to open the IAM Management Console.
2. In the console navigation panel, click **Users** > **Add user**.
3. Type the user name for the new user.
4. Select **Programmatic access** for the AWS access type.

   This access type is required to create an access key, which is required by IBM Spectrum Protect Plus. IBM Spectrum Protect Plus does not require the access type **AWS Management Consol access**.
5. Click **Next: Permissions**.
6. Click **Attach existing policies directly**, and then click **Create policy**.

   The **Create policy** page opens in a new browser window.
7. Click the **JSON** tab and enter the following actions:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "ec2:DetachVolume",
                "ec2:AttachVolume",
                "ec2:DeregisterImage",
                "ec2:DeleteSnapshot",
                "ec2:DescribeInstances",
                "ec2:CreateVolume",
                "ec2:DescribeTags",
                "ec2:CreateTags",
                "ec2:RegisterImage",
                "ec2:DescribeRegions",
                "ec2:RunInstances",
                "ec2:DescribeSnapshots",
                "ec2:DescribeAvailabilityZones",
                "ec2:CreateSnapshots",
                "ec2:DescribeVolumes",
                "ec2:CreateSnapshot",
                "ec2:DescribeSubnets",
                "iam:PassRole"
            ],
            "Resource": "*"
        }
    ]
}
```

8. Click **Review Policy**.
9. Type a name and description (optional) for the policy that you are creating.
10. Review the **Summary** section to see the permissions that are granted by the policy.
11. Click **Create policy**.
12. Close the browser window and return to the window that contains the **Add user** page.
13. Select the policy that you created from the list of policies.
14. Optional: Set a permissions boundary.
15. Click **Next: Tags**.
16. Optional: Add metadata to the user by attaching tags as key-value pairs.

    You can use tags to filter resources when you back up or restore EC2 data.
17. Click **Next: Review**.
18. Review your choices, and then click **Create user**.

    A new window opens showing the user name, access key, and secret key.
19. To view the secret key, click show **Show** next to the secret key.

20. Click **Download.csv** to save the access key ID and secret access key to a CSV file on your computer.

    Store the file in a secure location. You cannot access the secret access key again after this dialog box closes.
21. Click **Close** close the window.

What to do next

Add an account for EC2. To create an account, follow the instructions in "Adding an Amazon EC2 account" on page 375.

Adding an Amazon EC2 account

When an Amazon EC2 account is added to IBM Spectrum Protect Plus, an inventory of the instances that are associated with the account is captured. You can then run backup and restore jobs and generate reports for the instances.

Before you begin

An access key is required to add an EC2 account. The access key enables IBM Spectrum Protect Plus to connect to and inventory EC2 instances for data protection. Access keys that are already entered in IBM Spectrum Protect Plus are provided in a selection list. If the access key that you want to use is not in the list, you must add the access key and security key. Ensure that you have the access key and secret key that you want to add.

Procedure

To add an EC2 account, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **Amazon EC2**.
2. Click **Manage Accounts**.
3. Click **Add Account**.
4. Populate the fields in the **Account Properties** section:

   Account Name
   Enter a meaningful name to identify the access key that you select for the account.

   Use existing access key
   To specify a previously entered access key for the account, select this option and then select the key from the **Select a key** list.

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

   Access Key
   Enter the access key.

   Secret Key
   Enter the secret key.
5. Click **Save**.

   IBM Spectrum Protect Plus confirms a network connection, adds the EC2 account to the database, and then catalogs the account instances.

   If a message indicates that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connection.

What to do next

When you add an EC2 account to IBM Spectrum Protect Plus, an inventory is automatically run on each instance that is associated with the account. Instances must be detected to ensure that they can be backed up. You can run a manual inventory at any time to detect updates. For instructions about running a manual inventory, see "Detecting Amazon EC2 instances" on page 376.

Related tasks

"Backing up Amazon EC2 data" on page 376
Use a backup job to back up data in an Amazon EC2 instance.

"Restoring Amazon EC2 data" on page 378
Use a restore job to restore EC2 data from a backup copy. For example, if data on an instance is lost or corrupted. You can define a job that restores data to the original availability zone or to a different availability zone in the same region, with different types of recovery options and configurations available.

Detecting Amazon EC2 instances

Amazon EC2 instances are automatically detected after an EC2 account is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the account was added.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **Amazon EC2**.
2. In the list of EC2 accounts, select an account or accounts or click the link for an account to navigate to the regions or instances that you want to inventory.

   The navigation is in the order account > region > instance.
3. Click **Run Inventory**.

Backing up Amazon EC2 data

Use a backup job to back up data in an Amazon EC2 instance.

Before you begin

Complete the following steps:

1. Ensure that the accounts to be backed up are added to IBM Spectrum Protect Plus. For more instructions, see "Adding an Amazon EC2 account" on page 375.
2. Ensure that one or more SLA policies are configured for the EC2 instances. For more instructions, see "Creating an SLA policy for Amazon EC2 instances" on page 328.
3. Ensure that IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the backup job. For more information about assigning roles, see Chapter 19, "Managing user access," on page 643.
4. If an account is associated with multiple SLA policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

Procedure

To define an EC2 backup job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **Amazon EC2**.
2. Select the instances to back up in the Amazon EC2 Backup pane by taking one of the following actions:

   - To select all instances that are associated with an EC2 account, select the check box for the account. Any instances that are added to this account are automatically assigned to the SLA policy that you choose.
   - To select instances by region or specific instances, click the account name and navigate to the region or instance. The navigation is in the order account > region > instance. If an instance does not have an assigned name, the instance ID is shown as the instance name.

   To search for available instances, use the search function and toggle the displayed instances by using the **View** filter. Available options are **Instances** and **Tags**.

3. Click **Select SLA Policy** to add one or more SLA policies that meet your backup criteria to the job definition from the **SLA Policy Status** table.
4. Optional: To configure additional options for the SLA policies that you have added to the definition, in the **Policy Options** column of the **SLA Policy Status** table, click the clipboard icon for an SLA policy  and set the following options.

   If the job is already configured, click the icon to edit the configuration.

   Pre-scripts and Post-scripts

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

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

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

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

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

   Run inventory before backup
   Run an inventory job and capture the latest data of the selected instances before starting the backup job.

   Exclude Resources

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

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

   Separate multiple filters with a semicolon.

   Exclude Resources by Tag
   Exclude specific resources based on associated VM tags from the backup job. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*). Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *. Multiple filters may be separated with a semicolon.

   Force Full Backup of Resources
   This option is not used for EC2 backup operations.
5. Click **Save** to create the job definition.

   The job will run as defined by the SLA policies that you selected. To run the job immediately, click **Jobs and Operations** > **Schedule**. Select the job and click **Actions** > **Start**.

   Tip: When the job for the selected SLA policy runs, all instances that are associated with that SLA policy are included in the backup operation. To back up only selected instances, you can run an on-demand job. An on-demand job runs the backup operation immediately.

   - To run an on-demand backup job for a single instance, select the instance and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.

- To run an on-demand backup job for one or more instances, click **Create job**, select **Ad hoc backup**, and follow the instructions in "Running an ad hoc backup job" on page 627.

What to do next

After you define an EC2 backup job, create an EC2 restore job definition.

Related concepts

"Configuring scripts for backup and restore operations" on page 628
Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

"Restoring Amazon EC2 data" on page 378
Use a restore job to restore EC2 data from a backup copy. For example, if data on an instance is lost or corrupted. You can define a job that restores data to the original availability zone or to a different availability zone in the same region, with different types of recovery options and configurations available.

"Starting jobs on demand" on page 621
You can run any job on demand, even if the job is set to run on a schedule.

Restoring Amazon EC2 data

Use a restore job to restore EC2 data from a backup copy. For example, if data on an instance is lost or corrupted. You can define a job that restores data to the original availability zone or to a different availability zone in the same region, with different types of recovery options and configurations available.

Before you begin

Complete the following tasks:

1. Ensure that an EC2 backup job was run at least once. For instructions, see "Backing up Amazon EC2 data" on page 376.
2. Ensure that IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see Chapter 19, "Managing user access," on page 643.

About this task

IBM Spectrum Protect Plus uses clone mode to create long-term copies of instances.

Procedure

To define an EC2 restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Virtualized Systems** > **Amazon EC2** > **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **Amazon EC2**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
   a) Click an account in the list to show the instances that are available for restore operations. You can also use the search function to search for available instances. Enter all or part of a name to locate

instances that match the search criteria. You can use the wildcard character (*) to represent all or part of a name.

Use the **View** filter to toggle displayed instances.

b) Click the plus icon next to the instance that you want to use as the source of the restore operation.

You can select more than one instance from the list. However, all selected instances must be in the same region.

If the instance has attached volumes, you can navigate to the volumes and select them for the restore operation. You cannot select both instances and attached volumes.

The selected instances or attached volumes are added to the restore list next to the account list. To remove an item from the list, click the minus icon next to the item.

c) Click **Next** to continue.

3. Complete the fields on the **Source snapshot** page to select the instance snapshots that you want to restore and click **Next** to continue.

   The fields that are shown depend on the number of instances that were selected on the **Select source** page.

   - If a single instance is selected, select the date range for the snapshots that you want to restore. The snapshots that are available for that date range are listed. Select the snapshot that you want to restore.
   - If multiple instances are selected, select the date range for the snapshots they you want to restore. The instances that have snapshots within that date range are listed. For each instance, select the restore point that you want to restore.

4. On the **Set destination** page, specify the Availability Zone that you want to restore instances to and click **Next**:

   Original Availability Zone
   Select this option to restore instances to the original Availability Zone.

   Alternate Availability Zone
   Select this option to restore instances to an Availability Zone that is different from the original Availability Zone, and then select the alternate location from the available resources.

   If you are restoring an attached volume, select the destination instance in the alternate Availability Zone and enter an optional device name in the **Destination Attachment** section.

5. On the **Set network** page, change the subnet for each Availability Zone if you selected **Alternate Availability Zones** on the **Set destination** page. If you selected **Original Availability Zone**, no settings are provided on this page. Click **Next** to continue.

   The Availability Zone subnet must be in the same region as the instances that are selected in step “2” on page 378.

6. On the **Restore method** page, you can change the name of the restored instance by entering the new instance name in the **Rename Instance (optional)** field. Click **Next** to continue.

7. If you are running the restore job in advanced mode, you can set additional options as follows:

   Power on after recovery
   Toggle the power state of an instance after a recovery is run. Instances are powered on in the order in which they are recovered.

   Continue with restore even if it fails
   Toggle the recovery of an instance in a series if the previous instance recovery fails. If disabled, the restore job stops if the recovery of an instance fails.

   Run cleanup immediately on job failure
   This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Restore instance tags
Enable this option to restore tags that are applied to instances through EC2.

Prepend prefix to instance name
Enter a prefix to add to the names of restored instances.

Append suffix to instance name
Enter a suffix to add to the names of restored instances.

8. Optional: On the **Apply scripts** page, choose the following script options and click **Next**.
   - Select **Pre-script** to select an uploaded script, and an application or script server where the prescript runs. To select an application server where the script will run, clear the **Use Script Server** check box. Go to the **System Configuration** > **Script** page to configure scripts and script servers.
   - Select **Post-script** to select an uploaded script and an application or script server where the postscript runs. To select an application server where the script runs, clear the **Use Script Server** check box. Navigate to the **System Configuration** > **Script** page to configure scripts and script servers.
   - Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails. When this option is enabled and the prescript completes with a nonzero return code, the backup or restore job continues to run and the prescript task status returns `COMPLETED`. If a postscript completes with a nonzero return code, the postscript task status returns `COMPLETED`. When this option is not selected, the backup or restore job does not run, and the prescript or postscript task status returns with a `FAILED` status.
9. On the **Review** page, review your restore job settings and click **Submit** to create the job.

Results

The begins after you click **Submit**, and an **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon  .

All running jobs are viewable in the **Jobs and Operations** > **Running Jobs** page.

Related tasks

“Adding an Amazon EC2 account” on page 375
When an Amazon EC2 account is added to IBM Spectrum Protect Plus, an inventory of the instances that are associated with the account is captured. You can then run backup and restore jobs and generate reports for the instances.

Restoring files

Recover files from snapshots that are created by IBM Spectrum Protect Plus backup jobs. Files can be restored to their original or an alternate location.

Before you begin

Note the following procedures and considerations before restoring a file:

- Review the file indexing and restore requirements in “File indexing and restore requirements” on page 64.
- Run a backup job with catalog file metadata enabled. Follow these guidelines:
  - Ensure that credentials are established for the associated virtual machine as well as the alternate virtual machine destination through the Guest OS Username and Guest OS Password option within the backup job definition.
  - Ensure that the virtual machine can be accessed from the IBM Spectrum Protect Plus appliance either through DNS or hostname. In a Windows environment, the default security policy uses the Windows NTLM protocol, and the user identity follows the default *domain\name* format if the Hyper-

V virtual machine is attached to a domain. The format *local_administrator* is used if the user is a local administrator.

– For a file restore to complete successfully, ensure that the user ID that is on the target machine has the necessary ownership permissions for the file that is being restored. If a file was created by a user that differs from the user ID that is restoring the file based on Windows security credentials, the file restore job fails.

About this task

Restrictions:

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

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

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

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

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

Hyper-V considerations

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

Linux considerations

If data is located on LVM volumes, the *lvm2-lvmetad* service must be disabled because it can interfere with the ability of IBM Spectrum Protect Plus to mount and resign volume group snapshots or clones. To disable the service, complete the following steps:

1. Run the following commands:

```
systemctl stop lvm2-lvmetad
```

```
systemctl disable lvm2-lvmetad
```

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

```
use_lvmetad = 0
```

If data resides on XFS file systems and the version of the xfsprogs package is between 3.2.0 and 4.1.9, the file restore can fail due to a known issue in xfsprogs that causes corruption of a clone or

snapshot file system when its UUID is modified. To resolve this issue, update xfsprogs to version 4.2.0 or later. For more information, see Debian Bug report logs.

Procedure

To restore a file, complete the following steps.

1. In the navigation panel, click **Manage Protection** > **File Restore**.
2. Enter a search string to search for a file by name, and then click the search icon.

   For more information about using the search function, see Appendix A, “Search guidelines,” on page 681.
3. Optional: You can use filters to fine-tune your search across specific virtual machines, date range in which the file was protected, and virtual machine operating system types.

   Searches can also be limited to a specific folder through the **Folder path** field. The **Folder path** field supports wildcards. Position wildcards at the beginning, middle, or end of a string. For example, enter `*Downloads` to search within the `Downloads` folder without entering the preceding path.

   Note: Only file objects for which a snapshot was taken during the date range that is specified will be visible. For those objects, when the arrow is clicked beside the file object, all previous snapshots for that file object are displayed.
4. To restore the file by using default options, click **Restore**. The file is restored to its original location.
5. To edit options before restoring the file, click **Options**. Set the file restore options.

   Overwrite existing files/folder

   Replace the existing file or folder with the restored file or folder.

   Destination

   Select to replace the existing file or folder with the restored file or folder.

   To restore the file to its original location, select **Restore files to original location**.

   To restore to a local destination different from the original location, select **Restore files to alternative location**. Then select the alternate location from available resources by using the navigation menu or the search function.

   Restriction: A file can be restored to an alternate location only if credentials were established for the alternate virtual machine through the **Guest OS Username/Password** option in the backup job definition.

   Enter the virtual machine folder path on the alternate destination in the **Destination Folder** field. If the directory does not exist, it will be created.

   Click **Save** to save the options.
6. To restore the file by using defined options, click **Restore**.

Related tasks

“Backing up VMware data” on page 340
Use a backup job to back up VMware resources such as virtual machines (VMs), datastores, folders, vApps, and datacenters with snapshots.

“Restoring VMware data” on page 350
VMware restore jobs support Instant VM Restore and Instant Disk Restore scenarios, which are created automatically based on the selected source.

Chapter 12. Protecting file systems

File systems that contain directories and files that you want to protect can be registered with IBM Spectrum Protect Plus. Select the file system servers and the drives that contain data that you want to protect. Microsoft Windows ReFS and NTFS file systems can be registered with IBM Spectrum Protect Plus so that you can set up backup jobs or regularly scheduled service level agreement (SLA) policies.

You can protect local file systems that are assigned to a drive letter. Clustered volumes and drive shares are not protected by IBM Spectrum Protect Plus.

Windows file systems

After you successfully register the machine that hosts the Microsoft Windows NTFS or ReFS file system with IBM Spectrum Protect Plus, you can start to protect your data on the listed volumes and drives. You can also create an on-demand backup of your file systems data, or set up service level agreement (SLA) policies to run regular scheduled backup jobs.

Ensure that your environment in which the file system is located, meets the minimum system requirements. For more information about the system requirements, see "File system requirements" on page 71.

The IP address of the machine you register must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. Both must have a Windows Remote Management service that is listening on port 5985.

The fully qualified domain name must be resolvable and route-able from the IBM Spectrum Protect Plus appliance server and from the vSnap server.

Prerequisites for file systems

All prerequisites for using IBM Spectrum Protect Plus with file systems must be met before you start protecting your resources.

Requirements for working with file systems with IBM Spectrum Protect Plus are available here, "File system requirements" on page 71.

Note: The user ID for registering Windows file servers can be set up with one of the following Windows configurations:

- The *Local System Administrator* user account with the User Account Control (UAC) security component set to `Disabled`. With this user you must open the Windows system **Control Panel** > **User Account Control Settings**, and move the slider to **Never notify**.
- A user who is a member of the Local Administrator Group with the Admin Approval Mode security policy setting disabled. With this user, you must open the Windows system **Local Security Policy**. From the **Security Settings** menu, choose **Local Policies** > **Security Options** > **`User Account Control: Run all administrators in Admin Approval Mode policy`**, and set this option to `Disabled`. Ensure that your Local Administrator Group includes the `Log on as Service` policy option.

Space prerequisites

Ensure that you have enough space on the machine that hosts the file system you are protecting. For more information about space requirements, see "Space requirements for protecting file systems" on page 384. When you are restoring data to an alternative location, allow for extra space. No files are overwritten during the restore process. When files of identical names are found, both copies are retained.

© Copyright IBM Corp. 2017, 2021

Handling a security certificate for Windows

To secure access for protecting file system files with IBM Spectrum Protect Plus, you must create a certificate and manage its placement.

About this task

Note: If the restore service cannot load the certificate, files are deleted and a new self-signed certificate and key are created.

Tip: If the IBM Spectrum Protect Plus file systems agent has run, you will find a self-signed certificate and key in the following location: `%LOCALAPPDATA%\FSPA\`. If the agent has not run yet, follow the steps to create and move the self-signed certificate and key.

The administrator can access this directory at the following path:
`C:\Users\Administrator\AppData\Local\`

Procedure

1. Create a key and signed certificate for the client machine.
   Neither the key or the certificate can have password protection as this affects the loading of files.
2. Create a directory folder called FSPA at a location like this `%LOCALAPPDATA%\FSPA`.
3. Copy the key and certificate and place them in the FSPA folder.
4. Copy the key and certificate in this folder.
5. Rename the key to `localfspagent.key`.
6. Rename the certificate to `localfspagent.crt`.

Space requirements for protecting file systems

Before you start backing up data that is stored on the registered file system, ensure that you have enough free disk space on the machine and in the vSnap repository for backup and restore operations.

Adding a file system

To start protecting the data on an ReFS or NTFS file system, you must add the host address where the file system is located. You can repeat the procedure to add every host that you want to protect with IBM Spectrum Protect Plus.

Before you begin

Restriction: In an IBM Spectrum Protect Plus environment, you can assign only one application server or file server per host. For example, if you register a host as a Microsoft Windows file system, you cannot register the same host as a Microsoft SQL Server or a Microsoft Exchange Server.

Note: The user ID for registering Windows file servers can be set up with one of the following Windows configurations:

- The *Local System Administrator* user account with the User Account Control (UAC) security component set to `Disabled`. With this user you must open the Windows system **Control Panel** > **User Account Control Settings**, and move the slider to **Never notify**.
- A user who is a member of the Local Administrator Group with the Admin Approval Mode security policy setting disabled. With this user, you must open the Windows system **Local Security Policy**. From the **Security Settings** menu, choose **Local Policies** > **Security Options** > **`User Account Control: Run all administrators in Admin Approval Mode policy`**, and set this option to `Disabled`. Ensure that your Local Administrator Group includes the `Log on as Service` policy option.

About this task

To add a file system to IBM Spectrum Protect Plus, you must have the DNS name or the IP address of the machine, a user ID, and the password.

Procedure

1. In the navigation, expand **Manage Protection** > **File Systems**.
2. In the **File Systems** page, click **Manage file servers**, and click **Add file server** to add the host server.

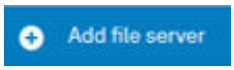


Figure 24. Adding a file system server

3. In the **File server properties** section, enter the DNS name or the IP address of the machine.
4. Specify the type of user for the Windows server you are adding.
   - Use an existing user ID and password.
   - Enter a new user ID and password.

   Note: The user ID for registering Windows file systems must be set up with one of the following Windows configurations:
   - The Local System Administrator user account with the User Account Control (UAC) security component disabled. With this user, you must access the User Account Control Settings dialog in your Windows system **Control Panel**, and move the slider to **Never**.
   - A user who is a member of the Local Administrator Group with the Admin Approval Mode security policy setting disabled. With this user you must access the Local Security Settings dialog on your Windows system and disable the **User Account Control: Run all administrators in Admin Approval Mode policy** setting. Ensure that your Local Administrator Group includes the **Log on as Service** policy option.

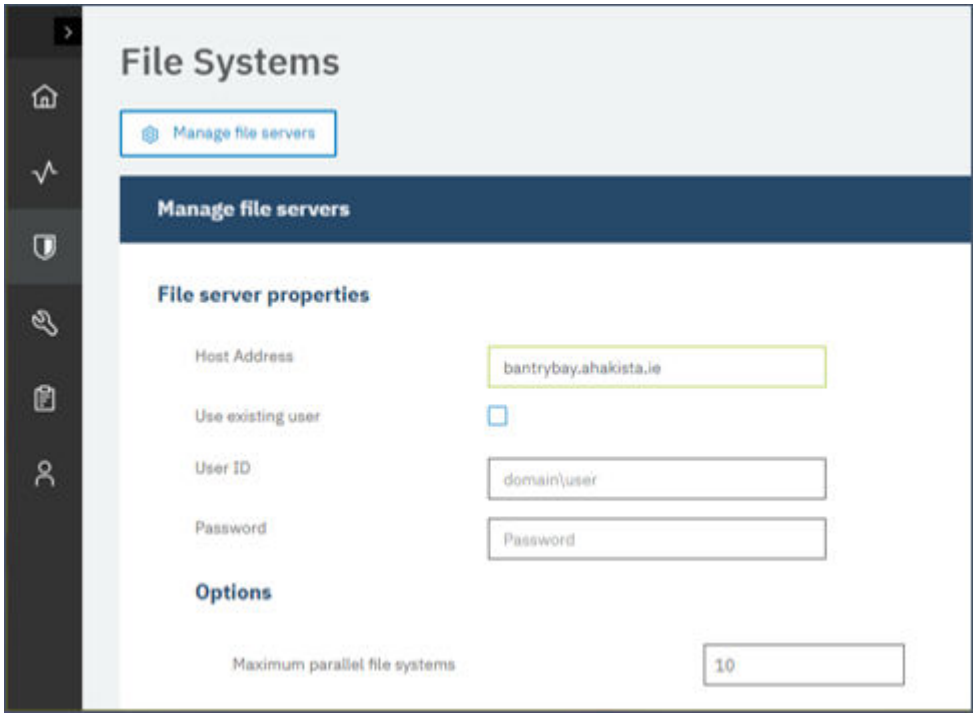


Figure 25. Managing agent users

Important: When you are entering the User ID, you do not need to enter the domain.

5. Set the maximum number of parallel file systems that are to be used for backing up data from the file system that is protected.

   This setting applies to each file system on this host. Multiple resources can be backed up in parallel when the value of the option is set to more than 1. Multiple parallel file systems can speed up restore operations.
6. Save the form.

What to do next

After you add the file system host to IBM Spectrum Protect Plus, an inventory is automatically run to detect the relevant volumes and drives.

To verify that the drives and volumes are added, review the job log. Go to **Jobs and Operations**,  . Click the **Running Jobs** tab, and look for the `Application Server Inventory` log entry that corresponds to the inventory that was started.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.

File systems must be detected to ensure that they can be protected. For instructions about running an inventory, see Detecting file systems.

Running an inventory to detect file systems

After you add a file system to IBM Spectrum Protect Plus, an inventory to detect volumes, drives, and mount points is run automatically. The inventory detects, lists, and stores the file system resources that are found on the selected host, and makes the data available for protection with IBM Spectrum Protect Plus.

Before you begin

Ensure that you added the file system to IBM Spectrum Protect Plus. For instructions, see Adding a file system.

Procedure

1. In the navigation panel, expand **Manage Protection** > **File Systems**.

   Tip: To add file systems to the **Servers** pane, follow the instructions in Adding a file system.
2. Click **Run Inventory**,  .

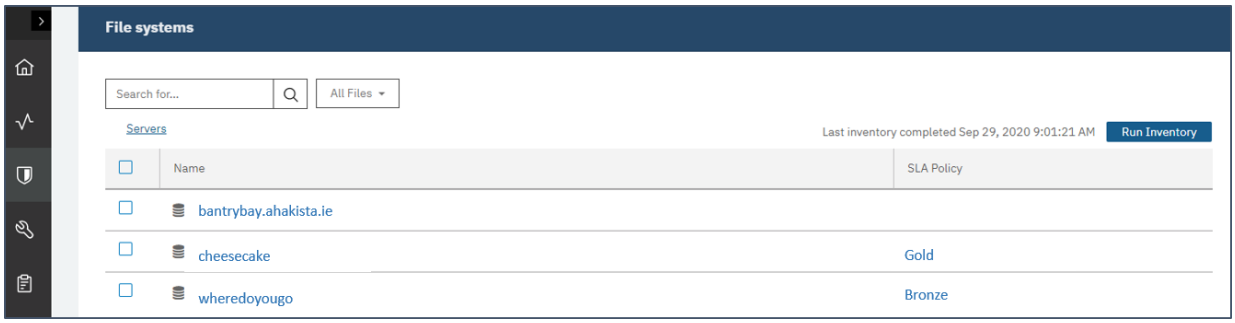


Figure 26. Detecting file systems

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

To view the job log, go to **Jobs and Operations**,  . Click the **Running Jobs** tab, and look for the newest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name. If the job is not displayed, adjust the **Job History Period** to a longer time interval.

3. Click a server name to open a view that shows the volumes, drives, and mount points that are detected for that server. If any entries are missing from the **Servers** list, check your file systems and rerun the inventory. In some cases, certain entries are marked as ineligible for backup; hover over the entry to reveal the reason why.

   Tip: To return to the list of servers, click the **Servers** hypertext.

Testing the file systems connection

After you add a file systems, you can test the connection. The test verifies communication with the server and the DNS settings between IBM Spectrum Protect Plus and the file systems server.

Procedure

1. In the navigation panel, click **Manage Protection** > **File Systems**.
2. In the **Microsoft Windows** window, click **Manage file servers**, and select the **Host Address** you want to test.

   A list of the machine hosts that are available are shown.
3. Click **Actions** and choose **Test** to start the verification tests for physical network connection, remote access, and Windows privileges connections and settings.

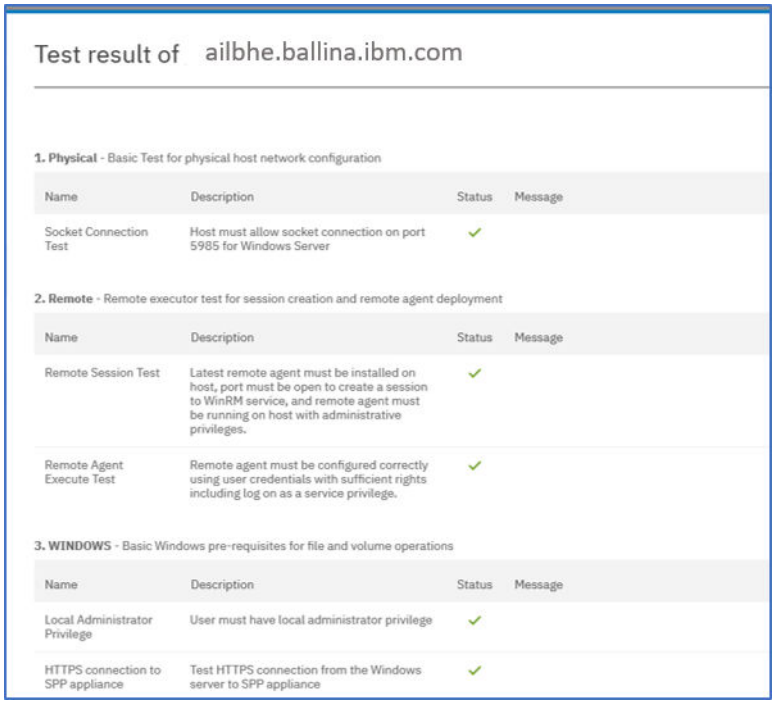


Figure 27. Testing the connection

The test report shows a list of the tests that were run. It consists of a test for the physical host network configuration, for the remote server installation on the host, and the Windows connections and privileges.

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

Backing up file system data

Define regular backup jobs and specify options to run and create backup copies to protect your file system data.

Before you begin

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

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

- Add the file system servers that you want to back up. For the procedure, see Adding a file system server.
- Configure a Service Level Agreement (SLA) Policy as described in this task.
- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see Chapter 19, “Managing user access,” on page 643.

A backup operation fails if the path is longer than 255 characters. If your paths are longer than 255 characters, you must enable longer paths by using the `Enable Win32 long paths` option in the Windows policy editor.

Note: Neither file system shares, or Microsoft cluster volumes can be protected with IBM Spectrum Protect Plus.

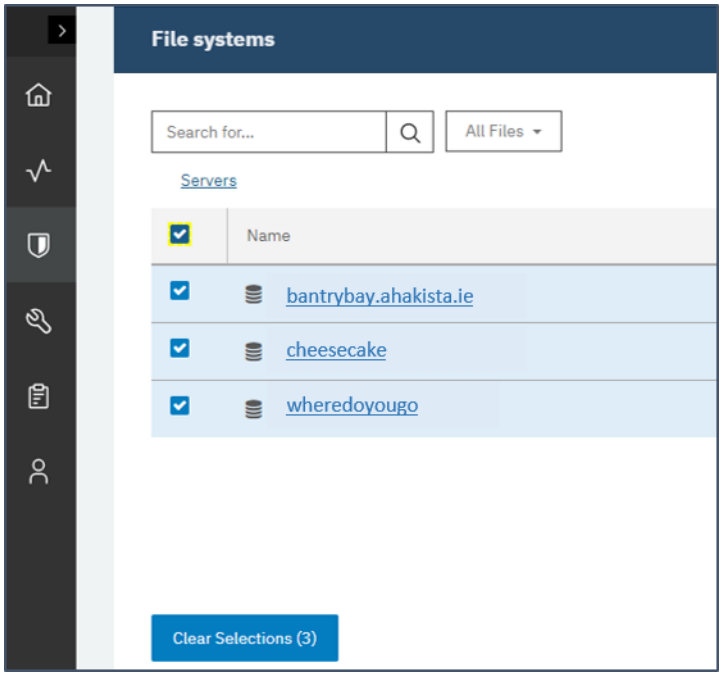
About this task

The following steps describe how to back up resources that are assigned to an SLA policy. To run an on-demand backup job for one or more resources regardless of whether those resources are already associated with an SLA policy, click **Create job**, select **Ad hoc backup**, and follow the instructions in “Running an ad hoc backup job” on page 627.

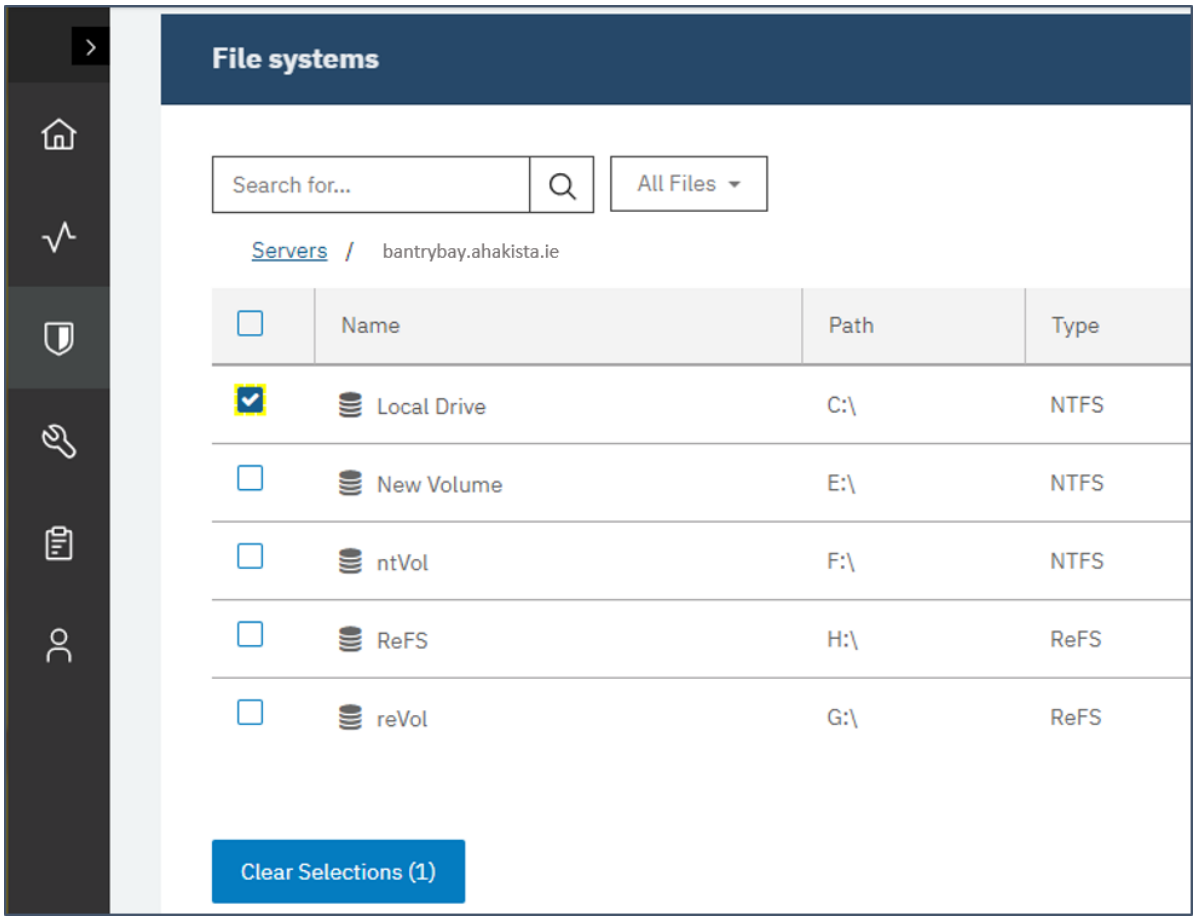
Procedure

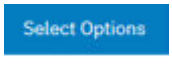
1. In the navigation panel, expand **Manage Protection** > **File Systems**.
2. Select a file system server to back up in the **Windows Backup** pane.
   - You can select an entire file system server by clicking the server name check-box. You can also select all file servers that are listed by clicking the check-box as shown. Any data

added to the servers selected is automatically assigned to the SLA policy that you choose.

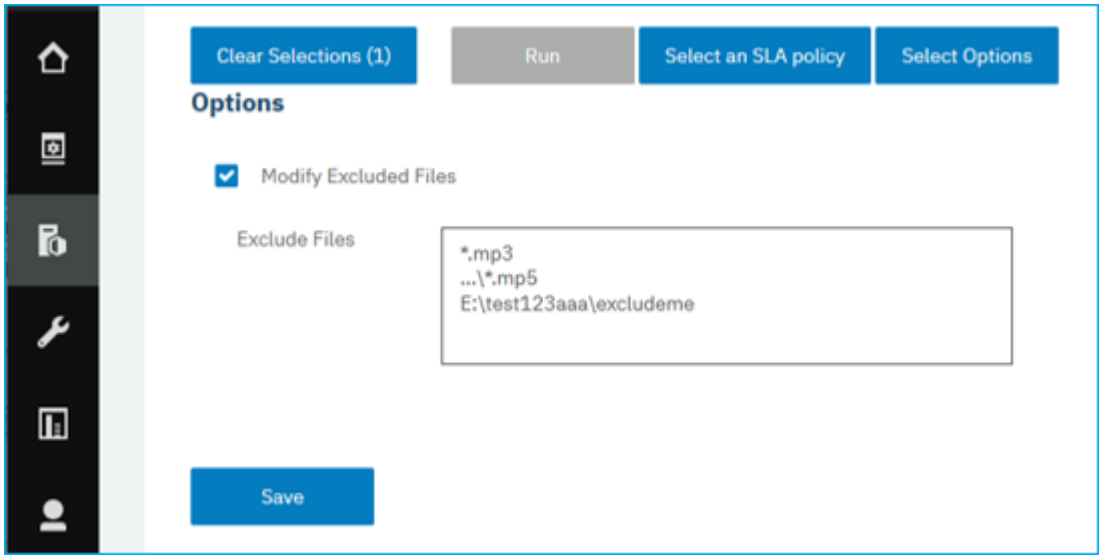


- Or, you can select a specific drive or mount point from a specific file system server by clicking the server name, and choosing a drive or mount point from the list.



3. Click **Select Options**  to specify files to be excluded from the backup job you are setting up. Alternatively, you can click **Modify Excluded Files** to leave the exclude rules as they are already defined. Click **Save** to commit your changes.

If you want to exclude all the files from a drive, you can specify the drive or a folder in a drive like this `Z:\test`. If you would like to exclude all files of a certain type from your backup job, you can specify that exclusion by using a string like this example `*.png`.

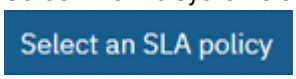


Tips for configuring options:

Review the following tips to help you configure options for the backup job:

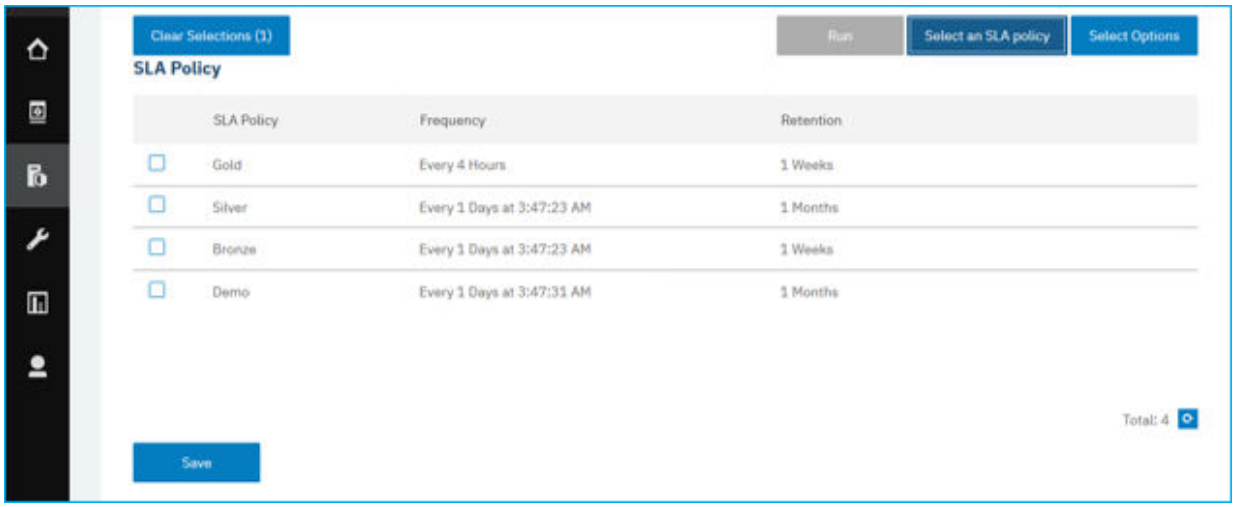
- To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
- If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
- Options that are shown in yellow indicate that the option value has changed from the previously saved value.
- To close the **Options** pane without saving changes, click **Select Options**.

4. Select the file systems server, drive, or mount point for backing up, and click **Select an SLA policy**



to choose an SLA policy for that item.

You can choose from the following options: `Gold`, `Silver`, or `Bronze`. Each policy type has different frequencies and retention rates as shown in the following picture:



If you want to define a new SLA policy, select **Manage Protection** > **Policy Overview**. In the **SLA Policies** pane, click **Add SLA Policy**, and define your policy preferences. To edit an existing policy with

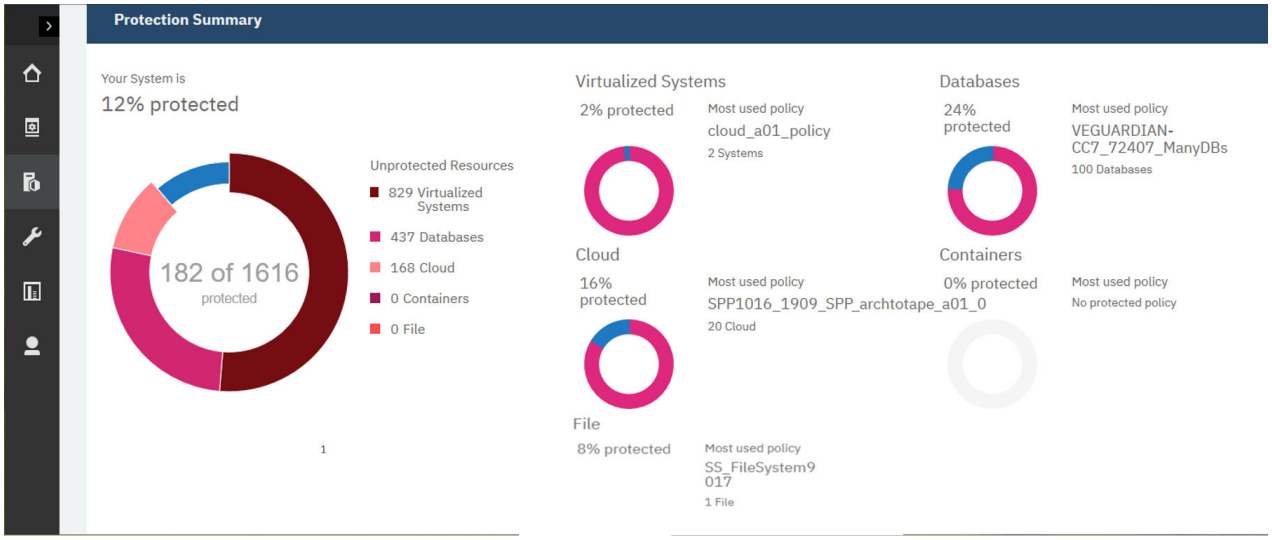
custom retention and frequency rates, click the edit icon and define your preferences. Click **Save** to commit your changes.

5. Click **Save** to save the SLA policy.

   If you want to run the backup job immediately, click **Actions** > **Start**. The status in the log changes to show that the backup is `Running`.

What to do next

To view the status of your existing file system SLA policies, select **Manage Protection** > **Policy Overview** to view a summary of your protection as shown in the following picture:



Exclude rules syntax

You can define exclusion rules to exclude specified drives, directories, or files from backup jobs. After you define the rule, the specified items are not backed up as part of your SLA policy or as part of an ad hoc backup job. When you run a restore job, the drives, directories, or files that are specified in the exclude rules are not restored to the new copy.

Exclude rules can be defined for the entire Windows file systems application. Rules that define the excluded resources are inherited by each file system that is being protected. To define new rules for a particular file system instance, you can add a rule in the **File systems** window. Select the file system servers that you want to add the rules for. The new rules that you define for that file system backup job override the exclude rules that are set for Windows file systems.

For more information about defining a backup job, see "Backing up file system data" on page 388.

To exclude a file, specify a rule as shown in the following example: `Z:\test\excludedFile.txt`.

To exclude all files in a folder, specify a rule as shown in the following example: `Z:\test\*`.

To exclude a folder and its files, specify a rule as shown in the following example: `DIR Z:\excludedFolder`.

If you use global variables for certain directories, you can use those variables to exclude items regardless of their location. For example, if you wanted to exclude a directory that is called `Blues` owned by the registered user, you can specify `DIR %USERPROFILE%\Blues`. For more information about using global variables to exclude resources, see Table 81 on page 393.

Table 80. Exclude rules syntax for Windows	
Syntax	**Syntax behavior**
:\	• Indicates a file system and Windows drive. • Must be included in all rules except for the FS rule. • A rule cannot start or end with this syntax. • A rule must start with a drive letter or wildcard followed by this sequence.
\	• Indicates the next directory level. • A rule cannot end with a backslash (\) character.
\...\	• Indicates that the rule applies to all directories below this level. • A rule cannot start or end with a \ . . . \ sequence. • This sequence must be after the drive specification sequence.
*	• This syntax is the wildcard for any character or number of characters. It is also used when no character is defined. • A rule can start or end with this syntax. • When used to indicate a drive letter, this syntax must represent one alphabetic character. • This wildcard cannot be a backslash (\) character.
?	• This syntax is used as a wildcard for any character for one occurrence only. • A rule can start and end with this syntax. • When this syntax is used to indicate a drive letter, it must be an alphabetic character between A and Z.
DIR	• This syntax indicates a directory rule, but it does not exclude any files in the affected directory. • This syntax must be a heading rule followed by a blank.
FS	• Indicates that a full file system drive is excluded from the job. • This syntax must be followed by a drive letter that can be a single character or a wildcard.
Spaces	• Spaces are allowed in file names or directory names. • A blank is not allowed before a backslash, \, or in a heading or trailing in a rule row. • Spaces are validated as single characters.

Table 80. Exclude rules syntax for Windows (continued)	
Syntax	**Syntax behavior**
Uppercase and lowercase text	Microsoft Windows is case-sensitive. Exclude rules ignore case.

Table 81. Exclude rules that use global variables	
Syntax	**Syntax behavior**
`DIR %PROGRAMDATA%`	• Indicates directory in the Windows `ProgramData` directory for the registered user. • This rule must be followed by a directory name, or wildcard to identify the resource to be excluded. For example, you can specify the following rule: `DIR %PROGRAMDATA%\WinZip` excludes the WinZip directory and all its content.
`DIR %USERPROFILE%`	• Indicates directory in the Windows `userprofile` directory for the registered user. • This rule must be followed by a directory name, or wildcard to identify the resource to be excluded. For example, you can specify the following rule: `DIR %USERPROFILE%\Elvis`. This rule excludes the `Elvis` directory in that user's directory structure.
`DIR %PROGRAMFILES%`	• Indicates a directory in the Windows `Program Files` directory for the registered user. • This rule must be followed by a directory name, or wildcard to identify the resource to be excluded. For example, you can specify the following rule: `DIR %PROGRAMFILES%\*` to exclude all directories from the program files directory structure for the registered user.
`DIR %WINDIR%`	• Indicates any specified Windows directory for the registered user. • This rule must be followed by a directory name, or wildcard to identify the resource to be excluded. • An example is, `DIR %WINDIR%\README`. This rule excludes the README directory and all its content for the registered user.

Table 82. Valid exclude statements	
Rule example	**Result**
`*:\*`	This rule excludes all files from the file system root directory from all drives, but does not exclude the directories.
`DIR *:\*`	This rule excludes all directories from all drives, but does not exclude the files in the root directory.

Table 82. Valid exclude statements (continued)	
Rule example	**Result**
`DIR E:\...\*temp*`	This rule excludes all directories that start with temp in the directory name in all directories of the E: drive.
`DIR F:\Users\Bobby\*`	This rule excludes all content from the Bobby directory without excluding that directory itself. Files in the Bobby directory are excluded.
`DIR F:\Users`	This rule excludes all users who are listed in the Users directories and also excludes the Users directory.
`DIR F:\Users\Bobby M?gee`	This rule excludes all directories that match the name with a wildcard for one letter. This rule excludes users with names like Magee, Megee, Migee, and so on.
`DIR F:\Users\Bobby Magee`	This rule excludes the directory for the user who is defined, in this case Bobby Magee. With this rule, the directory for that user and all its content, including files and subfolders, are excluded.
`F:\...\*`	This rule excludes all files from the F:\ drive, but it does not exclude the directories.
`F:\Bobby.mp?`	This rule excludes all files that match Bobby `.mp?` in the file system root directory, such as Bobby.MP3, Bobby.MP4, and so on.
`F:\Bobby.txt`	This rule excludes the file Bobby `.txt` in the file system root directory.
`F:\Users\...\*.mp3`	This rule excludes all MP3 files for all users that are listed in the F drive.
`F:\Users\Bobby\...\*.mp3`	This rule excludes all MP3 files from the user directory Bobby.
`F:\Users\Bobby\...\*music*\...\*.mp?`	This rule excludes all MP files in all directories that have the word music in the directory name for user Bobby. The excluded files are MP2, MP3, MP4, and so on.
`F:\Users\John\*` `DIR F:\Users\John\*`	This rule combination excludes all files and all subdirectories for the user John, but does not exclude the John directory itself.
`F:\Users\John\tax\Tax_20??.pdf`	This rule excludes all documents that match the pattern Tax_20 in the `John\tax` directory. Files that are named like these are excluded, TAX_2000.pdf, TAX_2019.pdf, and so on.
`FS F`	This rule excludes the file system F drive.
`FS *`	This rule excludes all drives in the file system.
`FS ?`	This rule excludes all drives.

Invalid exclude syntax

The following syntax is invalid for exclude rules:

- `\no`
- `*`
- `\*`
- `F:\no\`
- `DIR \no`
- `DIR F:\no\`
- `DIR *`
- `DIR F:\*\`

Tip:

To verify the results of an exclude rule, view the job log file. In the navigation panel, click **Jobs and Operations** and open the **Running Jobs** tab. In the **Application Server Backup** section, find the newest log entry.

Restoring file system data

To restore file system data from a vSnap repository, define a job that restores data from either the newest backup or an earlier backup copy. By using the File Systems File-Level Restore browser, you can select the file system resources to add to the job, and specify whether to restore data to another file system or to a different directory in the same file system on the same host.

Before you begin

Important: File are always restored to the same host machine.

- Ensure that at least one file system backup job was run successfully. For instructions about setting up a backup job, see "Backing up file system data" on page 388.
- Ensure that appropriate IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For instructions about assigning roles, see Chapter 19, "Managing user access," on page 643.

Ensure that the system has sufficient space to allow the restore operation to complete. For more information about space requirements, see Space requirements for file system protection. For more information about prerequisites and setup, see Prerequisites for file system protection.

About this task

Attention: When you define the restore job, the **Run cleanup immediately on job failure** option is selected by default. This option must not be cleared unless you are instructed by the IBM Software Support team to do so.

Important: Restoring of files is always on the same host machine.

Procedure

1. In the navigation panel, expand **Manage Protection** > **File Systems** and click **Create job**

   .
2. Select **Restore**.

   The **Restore** wizard opens.
3. Optional: If you started the restore wizard from the **Jobs and Operations** page, click **file system** as the source type and click **Next**.

Tips:

- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
- The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

4. On the **Select source** page, click a file system server to show the volumes that are available on that server. Select a volume by clicking the plus icon next to that volume name . Only one file system volume from the backup can be added. Click **Next** to continue.
5. On the **Source snapshot** page, select the snapshot that you want to restore. Click **Next** to continue.

   The available snapshots for the selected volume are listed with a timestamp, the associated SLA policy, and the source type that is available: backup, archive, or replication copy.

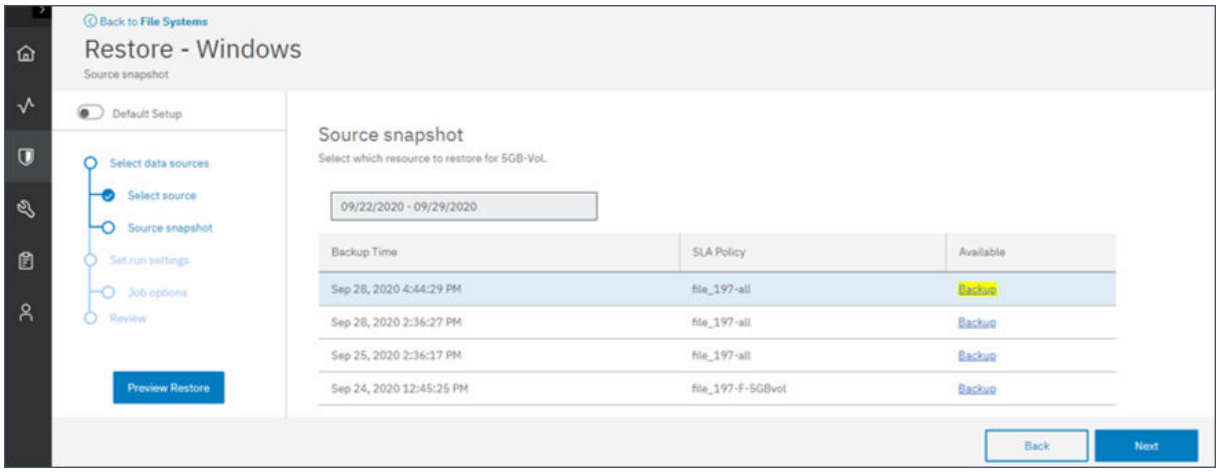


Figure 28. Selecting a source snapshot

6. On the **Review** page, review your selections for the restore job. If all selections are correct, click **Submit**, or click **Back** to edit the selections.

   The **Active Resources** tab in Jobs and Operations is opened to show the active resource that is prepared when you exit the restore wizard.

   Note: The active resource for the restore job that is submitted is not immediate and takes some time to display.
7. Open the File Systems File-Level Restore browser by clicking **Open Browser** on the **Active Resources** tab.

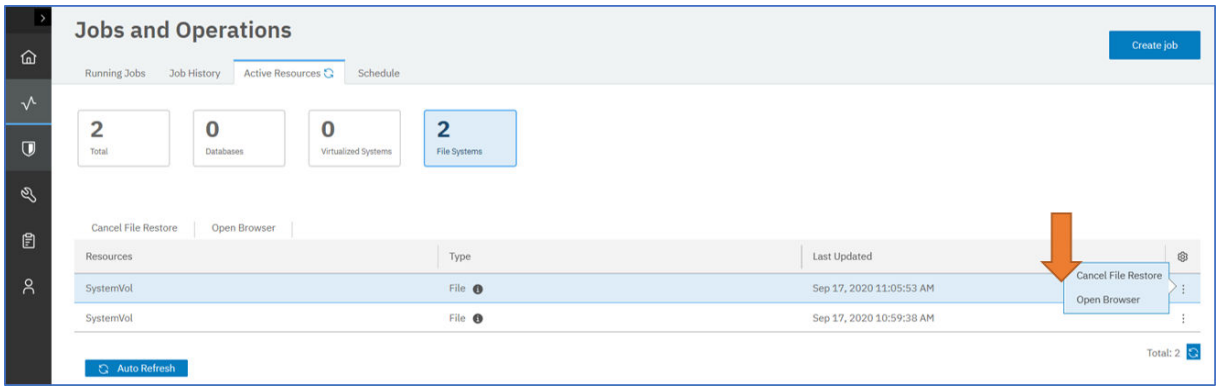


Figure 29. Opening the File Systems File-Level Restore browser from the Active Resources tab

8. In the File Systems File-Level Restore browser, select the file system resources to add to the restore job. Add items by clicking the add icon  next to the appropriate item.

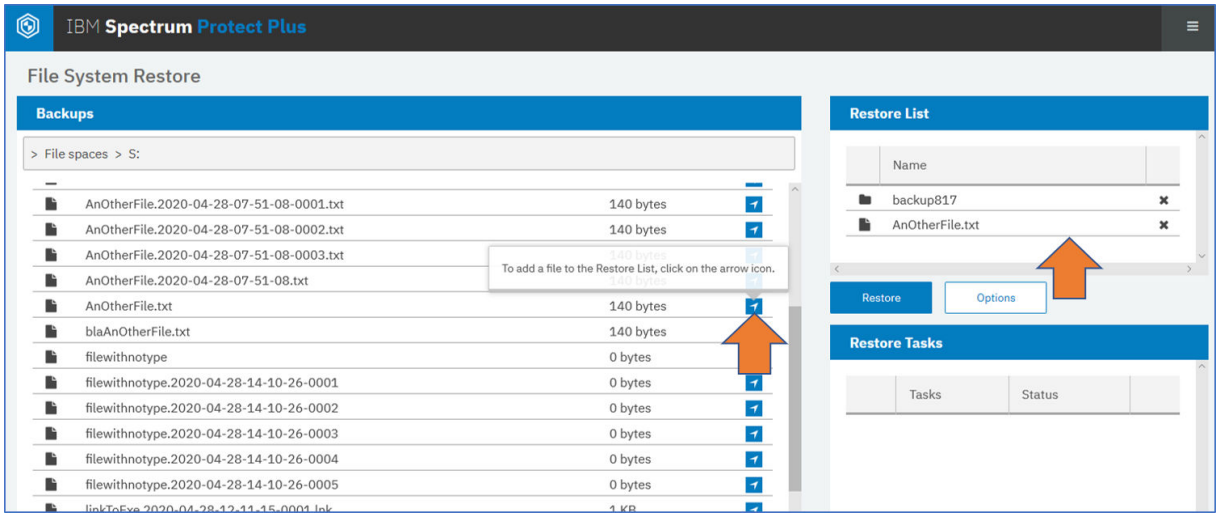
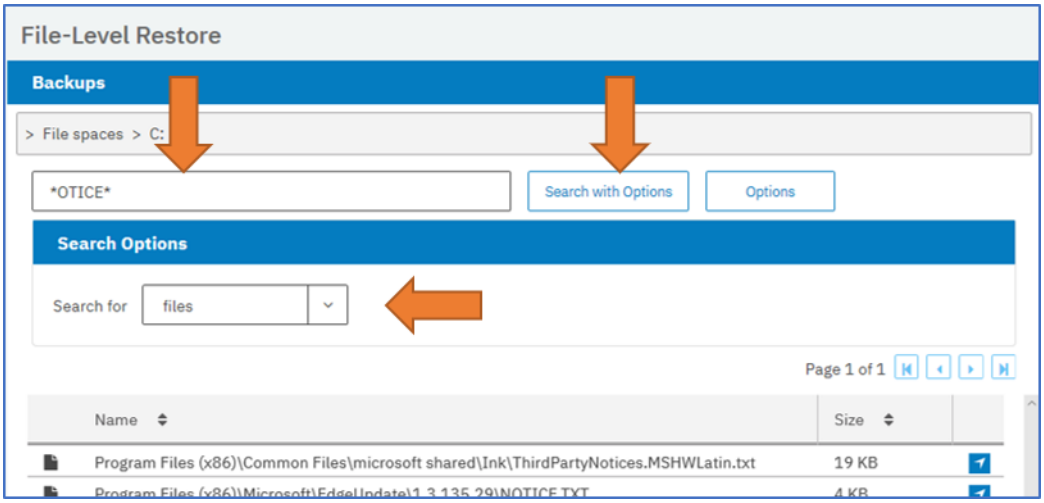


Figure 30. File Systems File-Level Restore browser: adding resources to the Restore List section

9. You can specify more options for the restore operation as follows:
   - To overwrite an existing copy of the file or directory for the restore job, click **Overwrite**.
   - To specify an alternative location for the restore job, click **Options** and enter a valid Windows local volume path as the target.
   - To overwrite the existing copy of the file or directory at the alternative location, click **Overwrite**.
   - If you cannot find the file or directory that you want to add to the restore list, use the search capability. Search with options to specify searching for files or directories. Use wildcard symbols * to broaden the search for partial name strings, as shown.



Restriction: Network shares are not valid alternative locations for restore jobs.

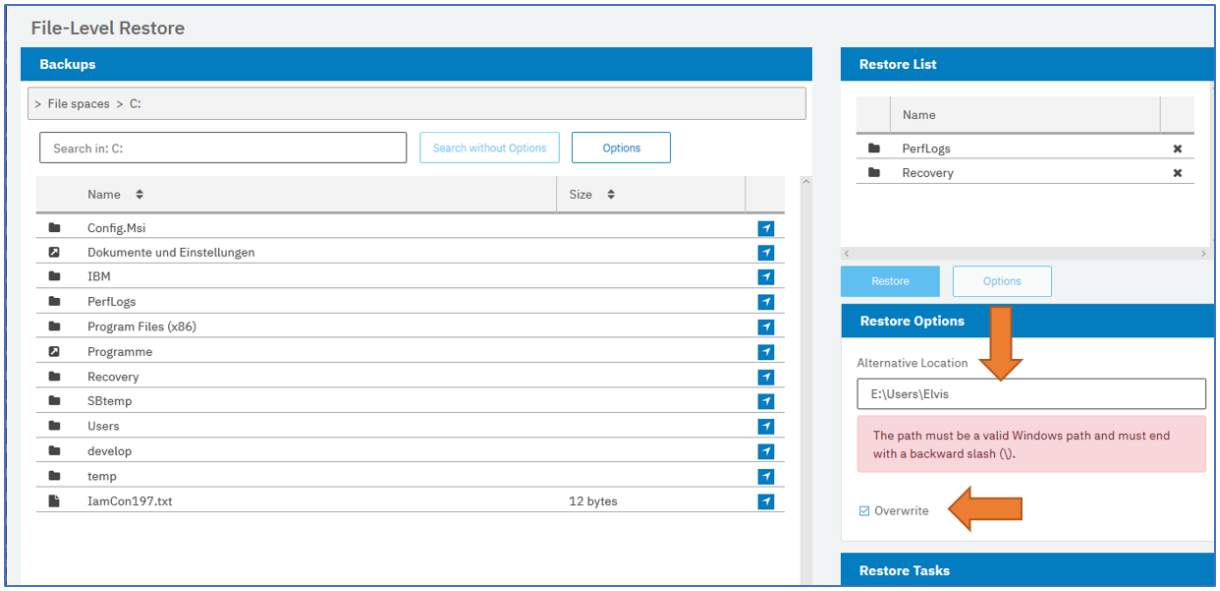


Figure 31. Specifying an alternative location for the restore job in the File Systems File-Level Restore browser

10. Click **Restore** to start the restore process.

    Existing files are overwritten only if the **Overwrite** option is selected. If files with identical names are detected during the operation, a timestamp is added to the new file and both files are stored at the restore location.

11. Optional: Monitor the progress of the restore operation in the **Restore Tasks** pane.

    Tip: The restore process is not tracked on the IBM Spectrum Protect Plus **Jobs and Operations** page. Progress of the restore job is tracked in the File Systems File-Level Restore browser.

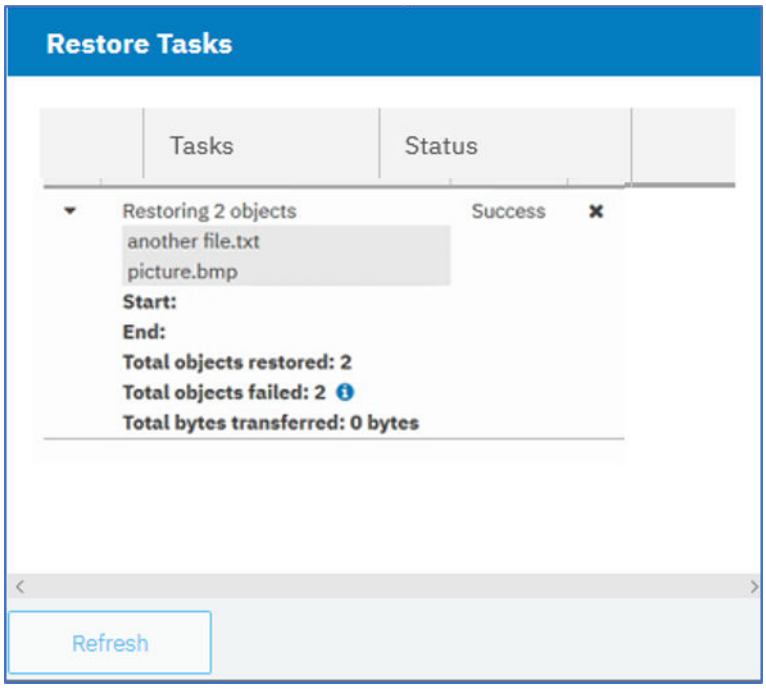


Figure 32. Monitoring the restore job in the File Systems File-Level Restore browser

What to do next

When the restore job is completed, remove the active resource by taking the following actions:

1. In the navigation panel, click **Jobs and Operations** > **Active Resources**.

2. Select the active resource that you no longer require, and click **Cancel File System Restore**.

File Systems File-Level Restore browser

When you prepare a restore job for a specific file system, the active resource that is created can be viewed in the **File Systems File-Level Restore** browser so that you can define the items to be restored. Use the browser to find and specify the directories or files that you want to restore from that file system. You can then specify an alternative location to direct the restored resources to a different location than the source.

Opening the File Systems File-Level Restore browser

After you click **Submit** in the Restore wizard, the restore job is prepared and the **Active Resources** tab in the **Jobs and Operations** page opens. To open the File Systems File-Level Restore browser, click the actions icon in the **Resources** table ⋮ and select **Open Browser** as shown.

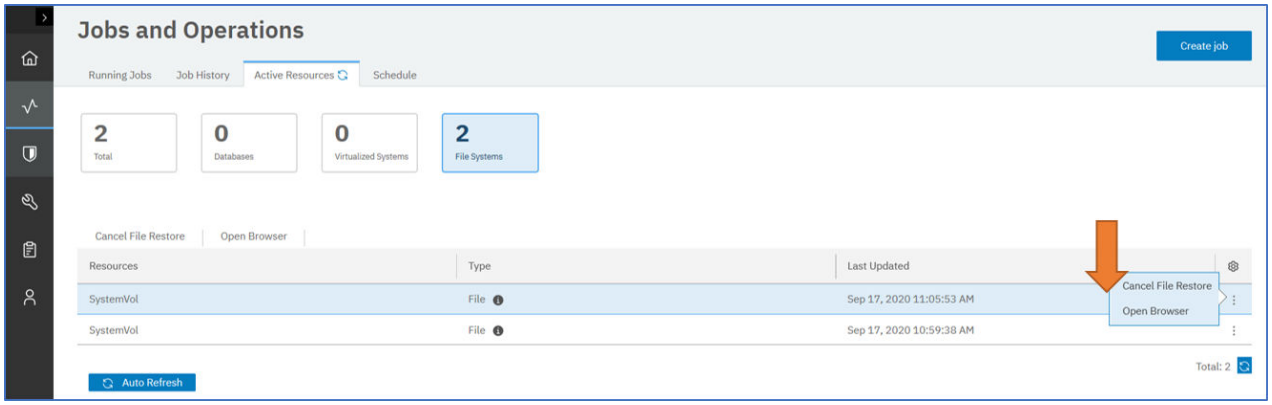


Figure 33. Opening the File Systems File-Level Restore from the Active Resources tab.

Adding resources to the restore job by using the File Systems File-Level Restore browser

To add specific file system resources to a restore job, navigate to the required file system, directories, or files. Add items to the Restore List section by clicking the icon next to the file system item  .

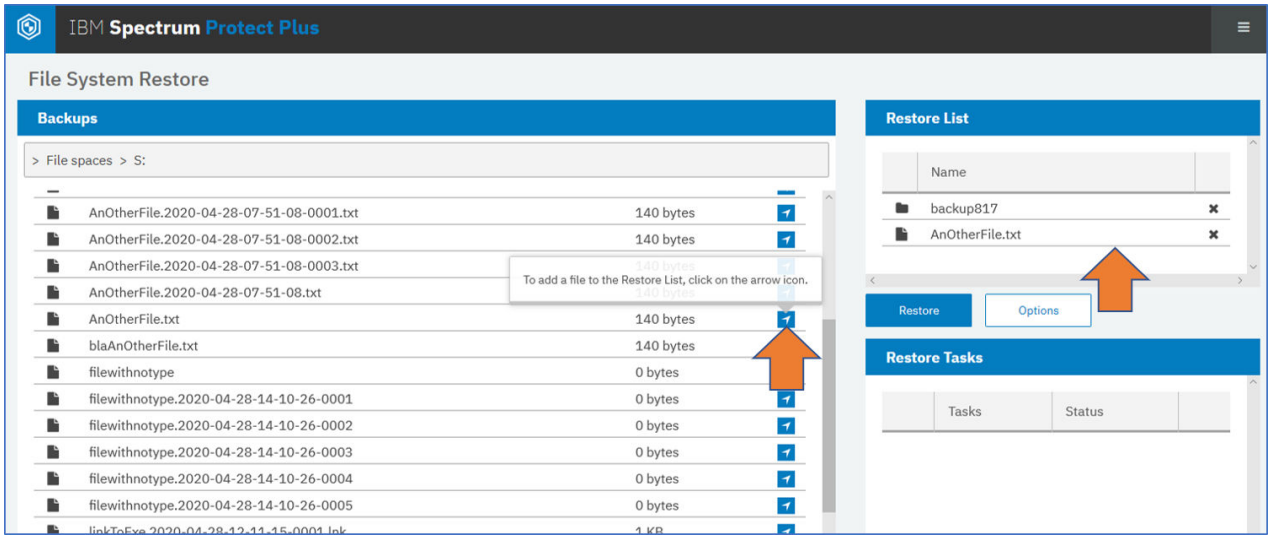


Figure 34. Adding file system objects to the restore job in the File Systems File-Level Restore browser

Restoring file system resources to an alternative location

To clone or copy resources, and to restore those resources to a different location on the same file system, you can specify a valid Windows path as the target in the **Alternative Location** in the **Options** pane.

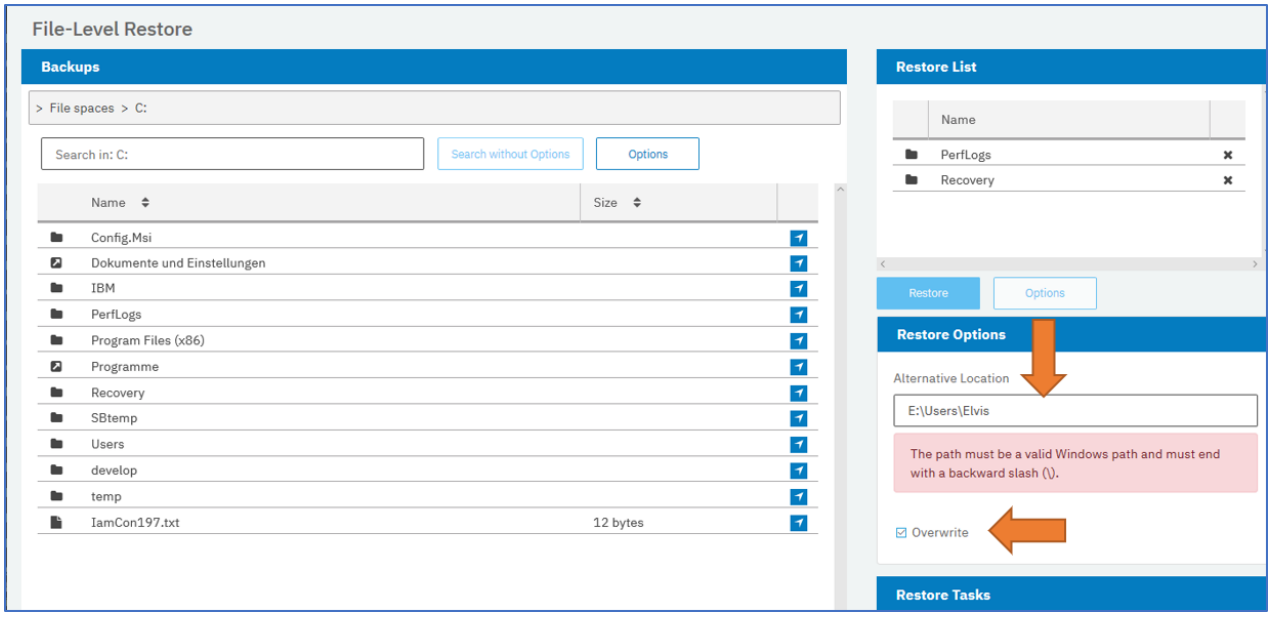
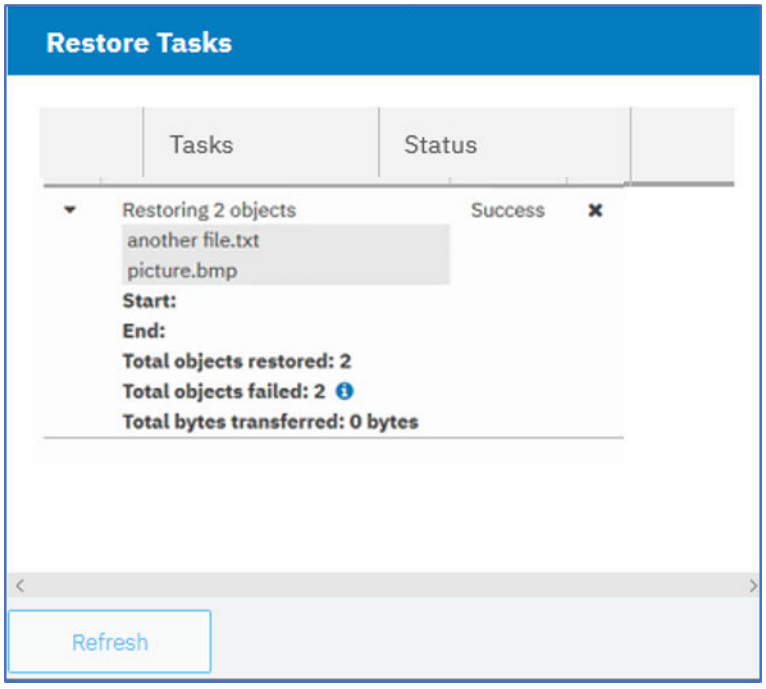


Figure 35. Specifying an alternative location on the same file system for the restore job in the File Systems File-Level Restore browser

Monitoring a restore job

When you click **Restore** in the File Systems File-Level Restore browser, you can monitor the progress of the restore job in the **Restore Tasks** pane.

Figure 36. Monitoring a restore job in the File Systems File-Level Restore browser



Chapter 13. Protecting containers

Container Backup Support is a feature of IBM Spectrum Protect Plus that extends data protection to containers on Kubernetes and Red Hat OpenShift clusters. Kubernetes is a system for orchestrating containers across clusters of hosts.

To protect persistent volumes, namespace-scoped resources, and cluster-scoped resources in the Kubernetes or OpenShift environment, first create service level agreement polices that specify the backup frequency and retention period. Then, create jobs for backup and restore operations.

Overview of Container Backup Support

IBM Spectrum Protect Plus Container Backup Support protects data of persistent volumes, namespace-scoped resources, and cluster-scoped resources that are associated with containers in a Kubernetes or Red Hat OpenShift environment. You can run snapshot backup operations to create locally stored backup copies in the cluster, or you can run backup copies to the vSnap server for longer-term retention.

Data of persistent volumes, namespace-scoped, and cluster-scoped resources can be protected by using a container service level agreement (SLA) policy that specifies how often snapshot and copy backups are created and how long they are retained. If data on the original volume is damaged or lost, the volume can be restored from either the snapshot or copy backups on the vSnap server. If data in any resource is damaged or lost, that data can also be restored.

Supported storage types

Container Backup Support protects volume data that was allocated by a storage plug-in that supports the Container Storage Interface (CSI) provided for Kubernetes. Container Backup Support is fully tested with Ceph RADOS Block Device (RBD), Ceph File System (CephFS), IBM Spectrum Scale, and IBM Spectrum Virtualized storage environments. The CSI plug-in provides snapshot capabilities that are used for backup operations.

For persistent volumes with a block-based storage type, such as Ceph RBD and IBM Spectrum Virtualized, block-based copy backup and restore operations are performed. For persistent volumes with a file system-based storage type, such as CephFS and IBM Spectrum Scale, file-based, incremental copy backup and restore operations are performed. During incremental backups, only new and changed data is copied to the IBM Spectrum Protect Plus vSnap server.

For IBM Spectrum Scale backup operations, snapshots can be created only from independent fileset-based persistent volume claims (PVCs). PVCs that are based on lightweight directories and dependent file sets are not supported. These types of PVCs are automatically filtered and are not displayed in the container inventory in the IBM Spectrum Protect Plus user interface.

Deployment overview

The following figure shows how Container Backup Support is deployed in the Kubernetes or OpenShift environment and how it interacts with IBM Spectrum Protect Plus:

© Copyright IBM Corp. 2017, 2021

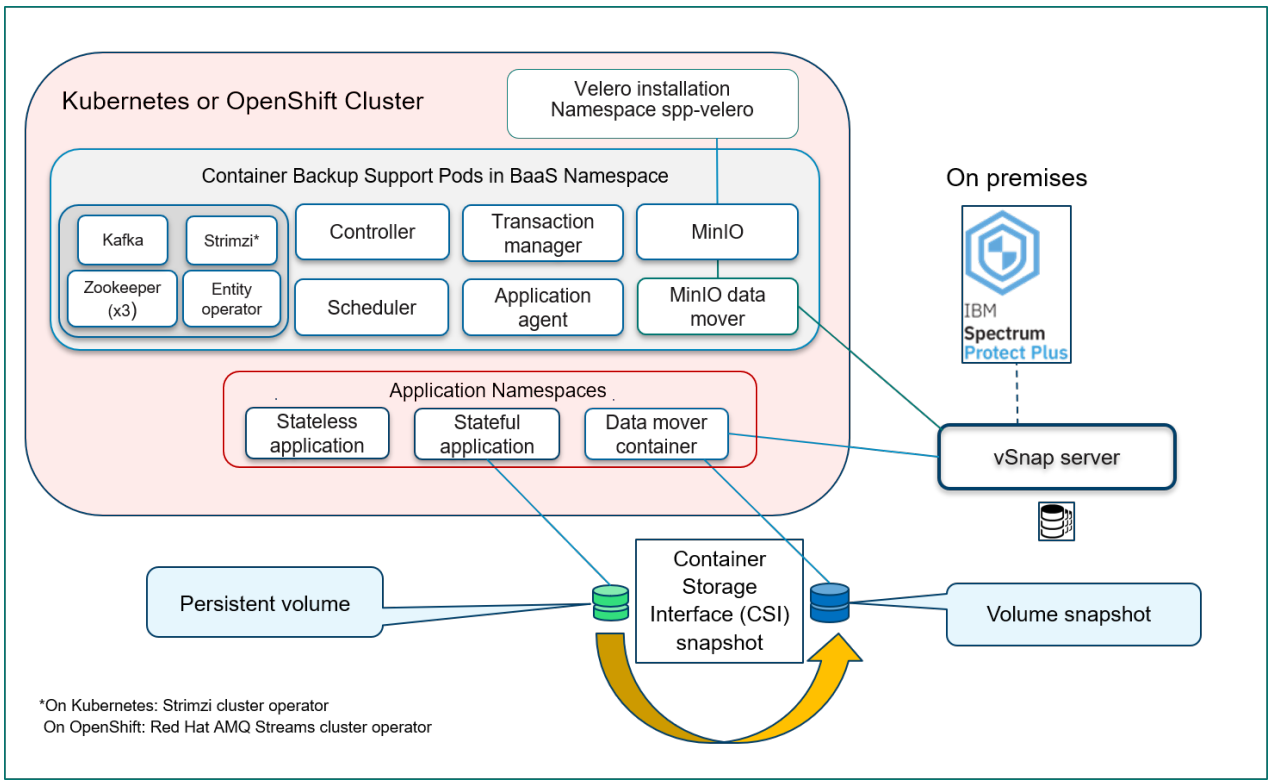


Figure 37. Container Backup Support deployment diagram

If you want to deploy Container Backup Support as a snapshot-only solution, the installation of the IBM Spectrum Protect Plus vSnap server is not required. When a schedule is run, snapshots are saved only on the storage system in your cluster; data is not copied to the vSnap server. With a snapshot-only deployment, data cannot be restored to another cluster.

Data mover container

Two types of data movers are deployed with IBM Spectrum Protect Plus. One is deployed as a container in a namespace where persistent volume claims (PVCs) exist. The other type, MinIO for resources, is deployed to the BaaS namespace. Data mover containers communicate with the IBM Spectrum Protect Plus instance outside of the Kubernetes or OpenShift environment for copy backup support as follows:

- The first type of data mover is deployed in the application namespaces.
- The second type of data mover is deployed in the BaaS namespace and copies resource data from MinIO to the vSnap server.

Container Backup Support uses PVCs to identify the persistent volumes to back up. For copy backup operations, when a schedule is run, snapshots and copy backups of a PVC are created at the time intervals that are specified by the SLA. The data mover copies the data and records the snapshot backups in the IBM Spectrum Protect Plus **Jobs and Operations** window. Snapshots that are created by on-demand backups are also recorded in IBM Spectrum Protect Plus.

Kafka cluster

The Kafka cluster handles messaging operations between the application agent and data movers. The Kafka cluster is managed by Red Hat AMQ Streams, which is an implementation of Strimzi, an operator that implements clusters of Apache Kafka. An operator is a container that configures, installs, maintains, and uninstalls, in this case, the Apache Kafka containers.

For example, in the OpenShift environment, the Kafka cluster is described by the following pods:

```
amq-streams-cluster-operator-v1.5.3-5b795f4c69-gdsrx   1/1     Running            0          24m
baas-entity-operator-c99f4c49b-p9v9c                   3/3     Running            1          24m
```

```
baas-kafka-0                                  2/2     Running            0          23m
baas-zookeeper-0                              1/1     Running            0          23m
baas-zookeeper-1                              1/1     Running            0          35m
baas-zookeeper-2                              1/1     Running            0          30m
```

The Kafka cluster consists of three zookeeper pods that form the storage system for Kafka, and a single Kafka application pod that sends and retrieves messages. The entity-operator pod is installed by the cluster-operator pod to manage local changes to the cluster. The cluster-operator pod is the only deployment that is described in the installation. On Kubernetes, the cluster-operator pod is called `baas-strimzi-cluster-operator`. On OpenShift, the cluster-operator pod is called `amq-streams-cluster-operator`.

On Kubernetes, Strimzi is installed along with the Container Backup Support product. When you update Container Backup Support, Strimzi is updated automatically.

On OpenShift, the installation files for Red Hat AMQ Streams are not provided by the Container Backup Support software. Instead, the Container Backup Support deployment requests an installation of Red Hat AMQ Streams from the OpenShift Operator Hub. Over time, the Operator Hub provides maintenance and updates of the Red Hat AMQ Streams operator and associated containers. A new version of the Container Backup Support software is not necessary to install patches and security updates to Red Hat AMQ Streams.

Multitenancy support

Container Backup Support manages backup and restore operations by using custom resources. All backup and restore objects belong to a Kubernetes or OpenShift namespace. The cluster administrator can restrict access to these objects. With controlled access, multiple users can run backup and restore requests in the same Kubernetes or OpenShift cluster. The backup and restore objects inherit a namespace from the PVC that identifies the persistent volume for backup and restore operations. For more information about multitenancy, see “Security features in Container Backup Support” on page 407.

Red Hat OpenShift Virtualization support

In Red Hat OpenShift clusters with the OpenShift Virtualization feature, virtual machines (VMs) that are running within an OpenShift container are protected during Container Backup Support backup jobs. The VMs must be allocated on storage that supports CSI.

All backup operations are PVC based. VMs are protected as part of PVC backup jobs. The backup operation has no explicit knowledge of workloads that are running within the PVC. To back up or restore a VM, use Container Backup Support to back up or restore the relevant PVC. You can also back up and restore cluster-scoped or namespace-scoped resources. Custom resource data for VMs is saved during resource backup jobs.

Important: For data to be consistent, the VMs must be powered off before a backup or restore job begins. After restoring a PVC, use the restored PVC as the source to re-create the VM. For more information, see “Preparing to restore persistent volumes that have VMs on OpenShift clusters” on page 444.

Backup and restore types

Container Backup Support provides multiple types of backup and restore functions for your PVC and other cluster resources. You can use the IBM Spectrum Protect Plus user interface or the Kubernetes or OpenShift command line to initiate backup and restore operations.

Backup types

The following types of backup operations are available:

Snapshot backup of PVCs
Creates a backup of a persistent volume by using Container Storage Interface (CSI) storage plug-in snapshot capabilities. The snapshot is stored in a location that is assigned by a Kubernetes snapshot class as defined by the backup administrator. Typically, this location is the same storage site as the persistent volume that is being backed up. The snapshot class must be compatible with the storage

class of the persistent volume. In other words, the snapshot class and storage class are defined and provided by the same CSI storage plug-in.

Snapshot backups are created by scheduled backup requests and on-demand backup requests.

During scheduled backups, snapshot backups are created at intervals that are defined by a service level agreement (SLA) policy.

During an on-demand backup request, a snapshot is taken immediately but no copy backup is created. After the initial snapshot backup, the volume is protected by the specified SLA policy.

Snapshot backup of namespace-scoped and cluster-scoped resources
Creates a backup of the cluster-scoped or namespace-scoped resources, which are backed up to the local MinIO in the backup as a service (BaaS) namespace. The MinIO object store that runs in the BaaS namespace claims a persistent volume to store the snapshot backups. This PVC is a single holding area for the data for all resource snapshot backups. The storage administrator must ensure that a default storage class exists in the cluster or a suitable storage class is defined by the MinIO storage class parameter when the BaaS is installed.

Snapshot backups are created by scheduled backup requests and on-demand backup requests. During scheduled backups, snapshot backups are created at intervals that are defined by a service level agreement (SLA) policy.

During an on-demand backup job, a snapshot is taken immediately but no copy backup is created. After the initial snapshot backup is finished, the volume is protected by the specified SLA policy thereafter.

Copy backup of PVCs
Copies the full persistent volume to an IBM Spectrum Protect Plus vSnap server. Based on predefined SLA policies, IBM Spectrum Protect Plus offers longer retention of copy backups compared to snapshot backups.

During scheduled backups, snapshot and copy backups are created at intervals that are defined by the SLA policy.

Copy backup of namespace-scoped and cluster-scoped resources
Copies the namespace-scoped or cluster-scoped resources from BaaS MinIO to the vSnap server. Based on predefined SLA policies, IBM Spectrum Protect Plus offers longer retention of copy backups compared to snapshot backups.

During scheduled backups, snapshot and copy backups are created at intervals that are defined by the SLA policy.

Snapshot-only protection
In this configuration, Container Backup Support can be deployed as a snapshot-only solution. The installation of an IBM Spectrum Protect Plus vSnap server is not required.

When you define SLA policies for containers, define the snapshot schedule but do not define a copy backup schedule. During scheduled backups, snapshots are created and stored locally but data is not copied from the cluster to the vSnap server.

With a snapshot-only deployment, data cannot be restored to another cluster.

Application-consistent PVC backup
In this configuration, Container Backup Support can be deployed in a Red Hat OpenShift environment as a solution to ensure consistency for application-level backup operations. You define SLA policies for application namespaces, which include the corresponding volumes and resources. You must specify Velero pre-backup and post-backup hooks as pod annotations.

For example, to ensure application consistency during a PVC backup operation, a pre hook can be run automatically to suspend or quiesce an application. This action can prevent any critical I/O processes from occurring on a volume while snapshots are created sequentially. After the snapshots are completed, a post hook can be run to resume the quiesced application.

Restore types

The following types of restore operations are available:

Snapshot restore of PVCs
Restores a snapshot to a new persistent volume. This type of operation is suitable for rapidly restoring recent snapshot backups.

Snapshot restore of namespace-scoped and cluster-scoped resources
Restores a snapshot copy from the BaaS MinIO to the same namespace or cluster. Snapshot copy backups cannot be restored to another cluster; they can be restored only to the same cluster or the same namespace. This type of operation is suitable for rapidly restoring recent snapshot backups.

Copy backup restore of PVCs
Restores a copy backup to the original persistent volume or to a new persistent volume. You can select to restore to an alternate cluster or namespace. If you want to restore a copy backup to the original persistent volume, the container to which the persistent volume is attached must not be running.

This type of operation is suitable for restoring persistent volumes from copy backups that are retained for a longer period on IBM Spectrum Protect Plus.

Copy backup restore of namespace-scoped and cluster-scoped resources
Restores a backup from the vSnap server to MinIO, and then to the original namespace or cluster. Alternatively, the restore operation can run to a new namespace or cluster, depending on what you specify in the restore job.

This type of restore operation is suitable for restoring resources from copy backups that are retained for a longer period on IBM Spectrum Protect Plus.

Application-consistent PVC restore
Restores an application-consistent copy backup to the original or alternative namespace in an OpenShift environment. You can run application-consistent restore jobs of both PVCs and resources. If you want to restore a copy backup to the original namespace, the container to which the persistent volume is attached must not be running. You can restore resources only if the resources do not exist in the target namespace.

SLA policies

Service level agreement (SLA) policies define how often snapshot backup and copy backup operations are run, and how long snapshots and copy backups are retained. You can set up custom SLAs that meet your operational requirements.

The storage administrator can create SLA policies by using the IBM Spectrum Protect Plus user interface. For instructions, see “Creating an SLA policy for containers” on page 329.

To view the list of SLA policies that are created for containers, use one of the following methods:

- In the IBM Spectrum Protect Plus user interface, click **Manage Protection** > **Policy Overview**. The **SLA Policies** section lists all the policies that are available. A predefined SLA policy, **Container**, is available to help you protect your persistent volumes. The **Container** policy runs the following operations:
  - Snapshot backups every 6 hours with a retention period of 1 day
  - Copy backups daily with a retention period of 31 days
- In the Kubernetes or OpenShift environment, issue the following command to view the SLA policies in the ConfigMap object `baas-sla` in the `baas` namespace:

  ```
  kubectl describe configmap baas-sla -n baas
  ```

  This command shows the available SLA policies for containers. If no SLA policy was created for containers, the output is empty.

The output is similar to the following example:

```
Name:         baas-sla
Namesapce:    baas
Labels:       app=baas
              component=scheduler
              release=10.1.8
Annotations:  <none>

Data
====
SLAs:
----
daily_midnight:
Snapshots are performed every 1 days and retained for 7 days.
No copy backups are performed.
----
every_4hours:
Snapshots are performed every 4 hours and retained for 1 days.
No copy backups are performed.
----
hourly:
Snapshots are performed every 1 hours and retained for 1 days.
No copy backups are performed.
```

The SLA is assigned to a volume or resource in the backup schedule definition. You can assign more than one SLA to a volume or resource.

When snapshot and copy backups expire, they are marked for expiration on IBM Spectrum Protect Plus and are deleted by IBM Spectrum Protect Plus maintenance jobs.

Related tasks

“Backing up persistent volumes in a Kubernetes cluster” on page 415
You can use IBM Spectrum Protect Plus to define backup jobs that run according to a service level agreement (SLA) policy. The SLA policy specifies how often backup operations are run, and how long snapshot or copy backups are retained.

“Scheduling backups of persistent volumes by using the command line” on page 468
By using the Kubernetes or OpenShift command line, you can schedule backup requests based on service level agreement (SLA) policies. SLA policies specify how often backup operations are run and how long snapshot and copy backups are retained.

User roles

Depending on their role, enterprise application developers and backup administrators interact with different user interfaces to protect persistent data in containers.

Application developer

The enterprise application developer uses the Kubernetes command-line tool (**kubectl**) or OpenShift command-line tool (**oc**) to complete the following tasks independent of the backup administrator:

- Initiates self-service backup and restore requests
- Selects a service level agreement (SLA) policy to use in backup requests to protect their volumes or resources
- Restores volumes and resources
- Views the status of backup and restore requests
- Queries information about snapshot and copy backups
- Removes SLA policy assignments from PVCs and resources
- Removes obsolete scheduled backup requests and on-demand snapshot requests

Backup administrator

The IBM Spectrum Protect Plus administrator with the **Containers admin** role completes the following tasks:

- Deploys and sets up Container Backup Support software in the Kubernetes or OpenShift environment
- Creates the storage class for persistent volumes and the snapshot class for storing snapshots
- Installs and configures IBM Spectrum Protect Plus
- Completes the following tasks in the IBM Spectrum Protect Plus user interface:
  - Manually registers a Kubernetes or OpenShift cluster or updates the cluster properties
  - Manually runs an inventory to detect cluster resources
  - Creates SLA policies
  - Defines SLA backup jobs to protect volumes and resources
  - Removes SLA policy assignments from PVCs and resources
  - Restores volumes and resources
  - Monitors inventory, backup, and restore jobs by using the IBM Spectrum Protect Plus user interface
  - Generates reports that show the history of container backup jobs by using the IBM Spectrum Protect Plus user interface
- Completes troubleshooting tasks, such as collecting log files for debugging the Kubernetes or OpenShift environment and viewing trace log files for Container Backup Support.

Security features in Container Backup Support

In addition to basic security features that are integrated into Container Backup Support, advanced security features are provided to help protect containers, secure network connections, encrypt data, and verify installation packages.

Security scanning of containers

Container Backup Support components are built on containers that are derived from the Red Hat Universal Based Image (UBI). The Container Backup Support software on each container was statically scanned for vulnerable components or libraries. In addition, the containers are dynamically scanned to help prevent runtime vulnerabilities such as code injection. After the scan, the software is tested by using an automated test suite to verify that Container Backup Support can operate as expected and correctly process erroneous input.

All containers, except for the data mover container, run in a dedicated namespace that provides further security isolation. The data mover must run in the same namespace as the persistent volume claim (PVC) for backup or restore operations because the mounting of the volume is limited to containers in a single namespace.

Least privileged containers

Each of the components in Container Backup Support runs under the principle of least privilege. The actions of the containers are constrained by the role-based authentication control rules that are associated with their service accounts in their separate namespace. In addition, the software in each container runs as a non-root user. The data mover runs as a privileged container because the data mover requires access to the device location on the host system of the volume that is backed up or restored. The application agent also runs as a privileged container, but it accesses no host resources. The application agent container is privileged in order to gain access to the **sudo** command to set up the data mover user account in the container at runtime. All other containers are not privileged.

Authentication of network connections

The network connections between Container Backup Support components are controlled by network policies that limit the connections to the ones that are required for correct operation. Connections to IBM Spectrum Protect Plus rely upon the security protocols that are provided by IBM Spectrum Protect Plus.

Multitenancy

Multitenancy is supported in Container Backup Support, which relies extensively on the authentication and authorization that is provided by the Kubernetes or OpenShift cluster for namespaces. Because the authorization is related to a namespace, any user who is authorized to create a BaaSReq object in that namespace can request a backup or restore for any PVC that is associated with that namespace. A BaaSReq object is a custom resource that is used in Container Backup Support requests.

Snapshots are protected by the Container Storage Interface (CSI) to restrict access to the namespace of the original PVC. Container Backup Support associates the namespace with the backup copies that are stored in IBM Spectrum Protect Plus, and the backup copies must be restored to volumes in the same namespace.

Encryption of data at rest

The cluster and storage administrators are responsible for enabling the mechanisms for protecting data at rest through encryption. The sensitive data includes the copy backup data and Container Backup Support secrets, which consist of user IDs and passwords that were specified during the installation process. The cluster administrator can specify that secrets are encrypted when stored in the cluster `etcd` database. For more information, see Encrypting Secret Data at Rest.

Container Backup Support does not implement additional encryption beyond what is provided by the cluster. However, the storage administrator can deploy an IBM Spectrum Protect Plus vSnap server that is enabled for encryption.

By using the IBM Spectrum Protect Plus user interface, the storage administrator can define service level agreements (SLAs) that store backup data on encrypted disks. When backup requests are created that specify encryption-enabled SLAs, data is directed to a vSnap server for encryption if the vSnap server is enabled for encryption of data at rest.

Code signing

The cluster administrator can verify that the Container Backup Support installation package has not been modified since it was generated by IBM. This process is accomplished by verifying the signature file that is included with the installation package against the appropriate signature and certificates. The verification process is described in the installation documentation.

For more information, see Chapter 6, “Installing Container Backup Support,” on page 197.

Backing up and restoring Kubernetes clusters

To protect persistent volumes and other Kubernetes resources that are attached to clusters, you can create service level agreement (SLA) polices and create jobs for backup and restore operations in the IBM Spectrum Protect Plus user interface.

Ensure that your Kubernetes environment meets the system requirements that are outlined in “Container Backup Support requirements” on page 75.

To protect cluster-scoped and namespace-scoped resources in your Kubernetes container workloads, ensure that the Velero tool is installed. For instructions, see “Installing Velero to protect resources” on page 198.

The following table shows the icon and naming convention for each type of resource that is displayed in the **Kubernetes Backup** pane, as well as a description of each type of resource:

Table 83. Kubernetes resource types

Icon and naming convention	Resource type	Description
 *clustername*	Cluster	A Kubernetes cluster.

Table 83. Kubernetes resource types (continued)		
Icon and naming convention	**Resource type**	**Description**
clustername:resources	Cluster resources	Cluster-scoped resources such as persistent volumes, cluster roles, storage classes, CSI drivers, volume snapshot classes, and custom resource definitions.
namespace	Namespace	A namespace in the cluster.
clustername:*namespace*:resources	Namespace resources	Namespace-scoped resources such as persistent volume claims, pods, containers, configuration maps, secrets, services, and deployments.
clustername:*namespace*:*pvcname*	Persistent volume claim (PVC)	PVCs in a namespace.

Related concepts

"Overview of Container Backup Support" on page 401
IBM Spectrum Protect Plus Container Backup Support protects data of persistent volumes, namespace-scoped resources, and cluster-scoped resources that are associated with containers in a Kubernetes or Red Hat OpenShift environment. You can run snapshot backup operations to create locally stored backup copies in the cluster, or you can run backup copies to the vSnap server for longer-term retention.

"Protecting containers by using the command line" on page 466
As an application developer in a Kubernetes or OpenShift environment, you can use the command-line interface to back up and restore container data, and to query the status of Container Backup Support requests.

Registering a Kubernetes cluster

If necessary, you can use the IBM Spectrum Protect Plus user interface to manually register a Kubernetes cluster or to modify the properties of a registered Kubernetes cluster.

About this task

After Container Backup Support is installed, the application host for the Container Backup Support container is automatically registered upon startup of the cluster host in Kubernetes. When a cluster is registered with IBM Spectrum Protect Plus, an inventory of the resources in the cluster is automatically captured, enabling you to complete backup and restore jobs, as well as run reports.

However, if the automatic registration was unsuccessful or if a registered cluster was accidentally unregistered, you can manually register the cluster by using the IBM Spectrum Protect Plus user interface.

You can also modify the properties of the registered cluster, such as changing the SSH port that is used to connect to the Container Backup Support container agent service.

For example, if you use a load balancer to distribute the workload in your cluster, you can edit the load balancer to use the port number for the Container Backup Support agent container service. You can then register the load balancer and port number with IBM Spectrum Protect Plus so that you do not have to configure the port number again.

Procedure

To manually register a cluster or to modify cluster properties, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **Kubernetes**.
2. In the **Kubernetes** page, click **Manage clusters**.

3. Take one of the following actions:
   - To manually register a cluster, click **Add cluster**.
   - To update existing cluster properties, in the list of host addresses, click the edit icon for the cluster host that you want to update.
4. Update the fields in the **Application Properties** section:

   Cluster Name
   The name of the cluster host or load balancer for the Container Backup Support container. You can enter a host name or IP address.

   The cluster name must match the value that is used for the **clusterName** parameter in the `baas-values.yaml` file. For more information, see "Setting up the installation variables" on page 206.

   Host Address
   The host address for the cluster host. Enter the IP address that is listed in the **clusterAPIServerips** parameter in the `baas-values.yaml` file. You can enter an IP address or a fully qualified domain name.

   Port Number
   The SSH port for the connection to the Container Backup Support agent container service.

   By default, the port is automatically assigned by Kubernetes during installation of Container Backup Support. To obtain this port number, issue the following command at the **kubectl** command line:

   ```
   kubectl get service -n baas | grep baas-spp-agent
   ```

   The output is similar to the following example:

   ```
   baas-spp-agent            NodePort    10.110.235.90    <none>        22:31299/TCP   111m
   ```

   The port number is the numerical string that follows `22:`. In the example, the port number is 31299.

   Use existing user
   Select this checkbox to use a previously entered username and password for the cluster host. Select a username from the **Select user** list.

   User ID
   Enter the username for the application host. The username must match the value that you specified for the DATAMOVER_USERNAME installation variable in the `baas-options.sh` file. For more information, see "Setting up the installation variables" on page 206.

   The credentials will be added to the list of existing users. This field is not available if you are using an existing user.

   Password
   Enter the password for the application host. The password must match the value that you specified for the DATAMOVER_PASSWORD installation variable in the `baas-options.sh` file. For more information, see "Setting up the installation variables" on page 206.

   The credentials will be added to the list of existing users. This field is not available if you are using an existing user.
5. Optional: Set additional options in the **Options** section:

   Maximum concurrent PVCs
   Set the maximum number of PVC snapshots or copy backups to create concurrently. Cluster performance is impacted when you back up many PVCs concurrently, as each PVC uses multiple threads and consumes bandwidth when copying data. Use this option to control the impact on cluster resources and minimize the impact on production operations.

   The default value is 10.

6. Click **Save**. IBM Spectrum Protect Plus confirms a network connection, adds the cluster to the IBM Spectrum Protect Plus database, and then catalogs the cluster resources, including namespaces and PVCs.

   If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connection.

What to do next

To verify that the clusters are updated, review the job log. In the navigation panel, click **Jobs and Operations**. Click the **Running Jobs** tab, and look for the most recent `Application Server Inventory` log entry. You can specify a filter to show only inventory jobs by clicking the filter icon, selecting **Inventory**, and clicking **Apply**.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name. If the status of inventory job status is `Partial`, click **Job Log** and review the log entries to find the error.

Clusters must be detected to ensure that their resources can be backed up. You can run a manual inventory at any time to detect updates in cluster resources. For instructions about running a manual inventory, see “Detecting Kubernetes resources” on page 411. For instructions about scheduling Kubernetes backup jobs, see “Backing up Kubernetes container data” on page 415.

Detecting Kubernetes resources

Kubernetes cluster resources are automatically detected after the cluster is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the cluster was added.

About this task

Run an inventory job periodically to help to ensure that all cluster resources are detected and that they are available for backup operations. To verify that namespaces in a cluster are detected, you can click the name of the cluster to expand it.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **Kubernetes**.
2. Optional: In the list of clusters, click a cluster name link to show the resources in that cluster.
3. Click **Run Inventory**.

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

   If you do not select a cluster from the list and click **Run Inventory**, an inventory job is started for all clusters.
4. To view the resources that are detected, click a hyperlink-enabled name to navigate to the different levels of a cluster. Click a breadcrumb to move to a previous level in the cluster.

   You can view the cluster resources by namespace or by label by changing the View options,



.

For more information about working with labels, see "Working with Kubernetes labels" on page 413.

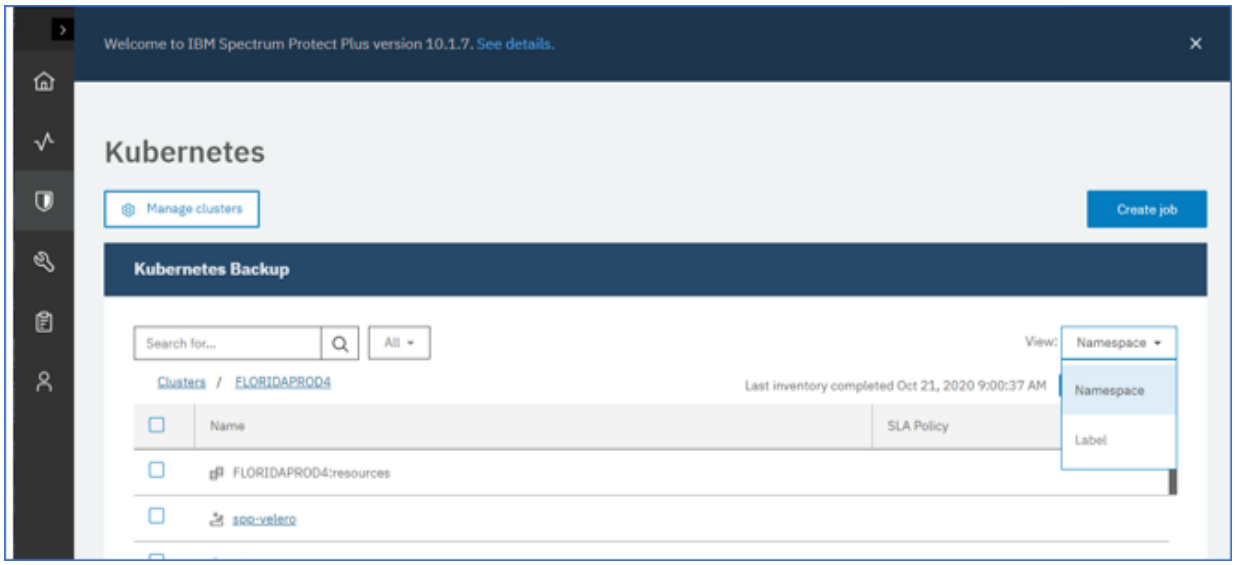


Figure 38. Kubernetes cluster resources

The following table shows the icon and naming convention for each type of resource that is displayed in the **Kubernetes Backup** pane, as well as a description of each type of resource:

Table 84. Kubernetes resource types

Icon and naming convention	Resource type	Description
clustername	Cluster	A Kubernetes cluster.
clustername:resources	Cluster resources	Cluster-scoped resources such as persistent volumes, cluster roles, storage classes, CSI drivers, volume snapshot classes, and custom resource definitions.
namespace	Namespace	A namespace in the cluster.
clustername:*namespace*:resources	Namespace resources	Namespace-scoped resources such as persistent volume claims, pods, containers, configuration maps, secrets, services, and deployments.
clustername:*namespace*:*pvcname*	Persistent volume claim (PVC)	PVCs in a namespace.

An example might include a cluster, a namespace in that cluster, and a PVC in that namespace such as *CLUSTERPROD1*:*db2-namespace*:*pvc-1*.

What to do next

To monitor the inventory job, in the navigation panel, click **Jobs and Operations**. Click the **Running Jobs** tab, and look for the most recent `Application Server Inventory` log entry. You can specify a filter to show only inventory jobs by clicking the filter icon, selecting **Inventory**, and clicking **Apply**.

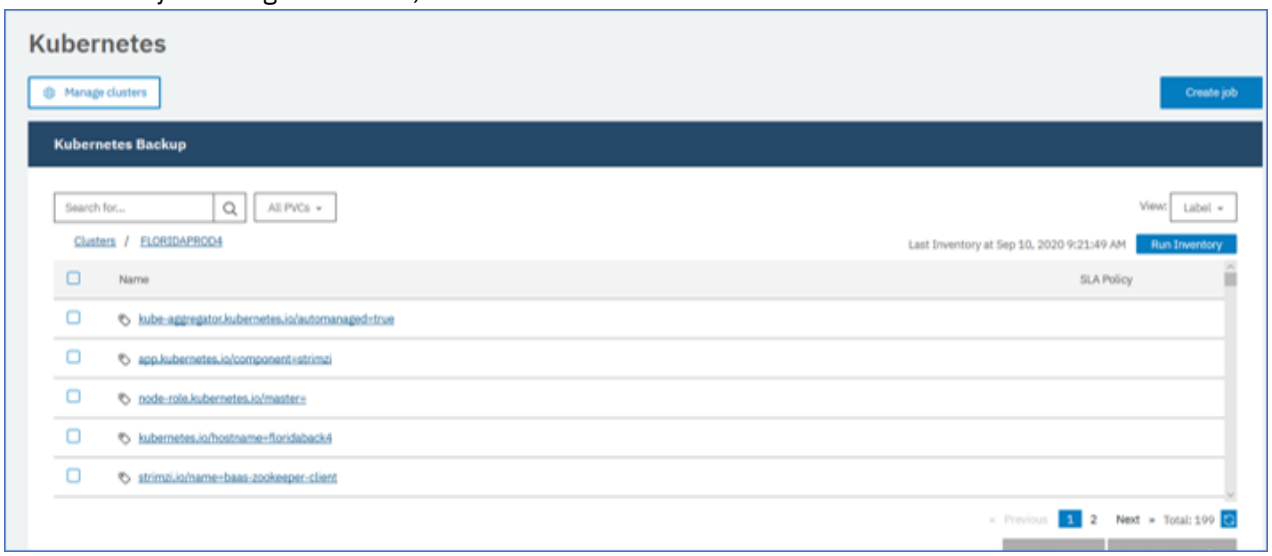
Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name. If the status of an inventory job is `Partial`, click **Job Log** and review the log entries to find the error.

Working with Kubernetes labels

Labels are used in Kubernetes environments to label key pairs, such as pods. Labels are used to specify identifying attributes of objects that are meaningful and relevant within the environment, and are generally used to organize subsets of objects. With IBM Spectrum Protect Plus, you can view and manage your Kubernetes resources by using their labels.

Inventory

To view your resources by their assigned labels, change the **View** setting to `Label`. When you expand the cluster after you changed the view, the labeled resources in that cluster are listed in the table as shown.



Backing up and SLA policy

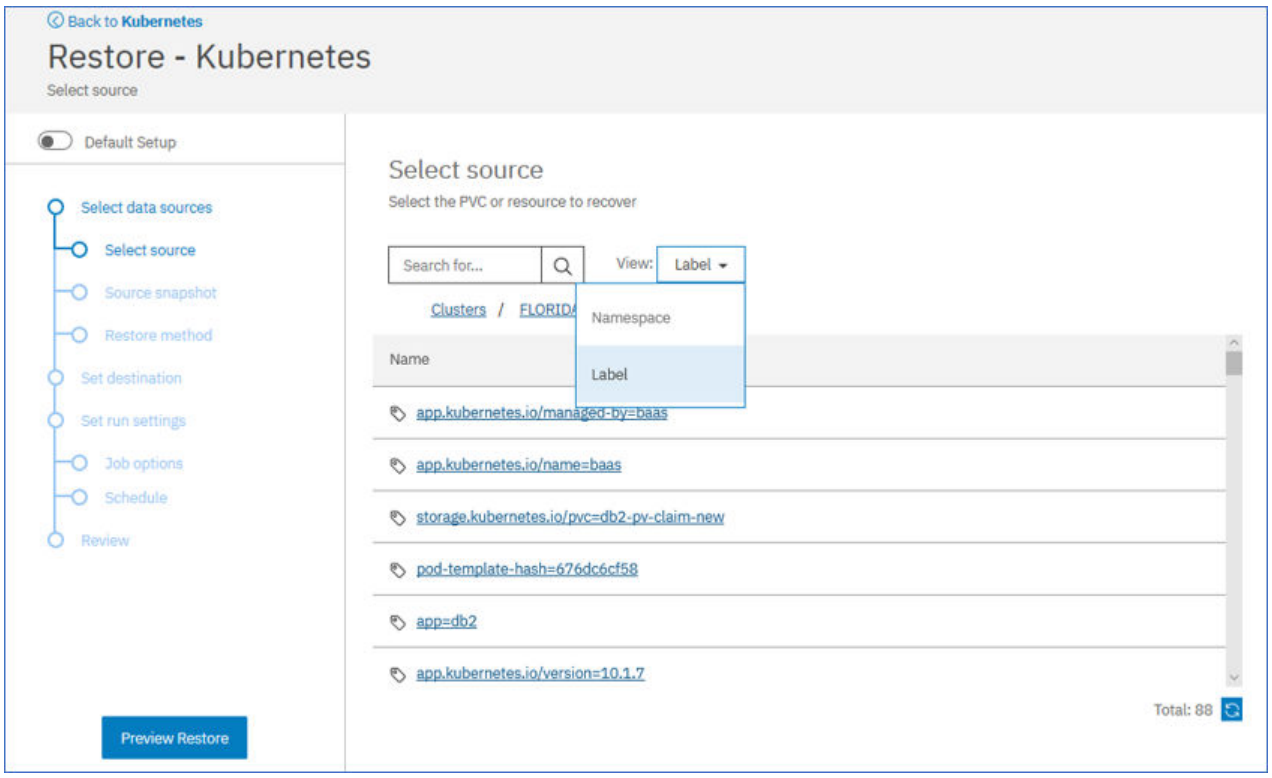
When you are setting up an SLA policy, you can select the labeled resources that you want to add to any predefined policy. When you assign the SLA policy and click save, the labeled resources are added to the policy for the next time it is run.

Ad hoc backup job

When you create an ad hoc backup job, you first select the SLA policy that you want to assign. Then, you can select the labeled resources that you want to add to the job.

Restoring labeled resources

When you are creating a restore job, you can change the view to view your resources by their assigned labels. From within the restore wizard, change the **View** setting to `Label`. When you expand the cluster after you change the view, the labeled resources in that cluster are listed in the table as shown.



Job log files

When you run a backup or restore operation for your labeled resources, the messages that are collected for the job log file list the labels in that operation.

Testing the connection to a Kubernetes cluster

You can test the connection to a Kubernetes cluster that you added to IBM Spectrum Protect Plus. The test function verifies communication with the cluster and tests domain name server (DNS) settings between the IBM Spectrum Protect Plus server and the cluster.

Procedure

To test the connection to a cluster, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **Kubernetes**.
2. Click **Manage clusters**.
   The list of available clusters are displayed.
3. Scroll through the list and locate the cluster that you want to test.
4. Click the **Actions** menu that is associated with the cluster and select **Test**.

   The test report shows you a list of the tests that ran and the status.

Backing up Kubernetes container data

You can use the IBM Spectrum Protect Plus user interface to define service level agreement backup jobs to protect container data such as Kubernetes persistent volumes, namespace-scoped resources, and cluster-scoped resources.

Backing up persistent volumes in a Kubernetes cluster

You can use IBM Spectrum Protect Plus to define backup jobs that run according to a service level agreement (SLA) policy. The SLA policy specifies how often backup operations are run, and how long snapshot or copy backups are retained.

Before you begin

Take the following actions:

- Ensure that persistent volume claims (PVCs) for the volumes that you want to protect are formatted. Backup requests are directed to PVCs. Backup operations of raw block volumes are not supported.
- If you do not plan to use the default SLA policy for containers, ensure that you configure an SLA policy. For instructions, see "Creating an SLA policy for containers" on page 329.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the Accounts pane. The user must be assigned the Containers Admin role. For instructions, see Chapter 19, "Managing user access," on page 643.
- If a PVC is associated with multiple SLA policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

About this task

To start protecting your PVCs on a regular schedule, you must apply an SLA policy to your PVC. The SLA policy also defines the backup target locations for your PVCs.

Procedure

To define an SLA backup job for one or more PVCs, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **Kubernetes**.
2. In the **Kubernetes Backup** pane, select the PVCs that you want to back up. You can use one of the following methods:

Method	Steps
To back up all PVCs in a cluster	Select the checkbox for a cluster name. A cluster is identified by the cluster icon  .
To back up PVCs that are associated with a namespace	a. Click **View** > **Namespace**. b. Click the name of a cluster that contains the PVCs that you want to back up. The list of namespaces within the cluster is displayed. A namespace is identified by the namespace icon  . c. To back up all PVCs in the namespace, select the checkbox for the namespace. To back up individual PVCs, click the namespace link and select the checkbox for each PVC that you want to back up. A PVC is identified by the PVC icon  .

Method	Steps
To back up PVCs that are associated with a label	a. Click **View** > **Label**. b. Click the name of a cluster that contains the PVCs that you want to back up. The list of labels within the cluster is displayed. A label is shown as a key-value pair and identified by the label icon . c. To back up all PVCs that are assigned to a label, select the checkbox for a label. To back up individual PVCs, click the label name and select the checkbox for each PVC that you want to back up.
To use the search function to filter the list of PVCs by SLA	a. Enter your search criteria in the **Search for** field. You can enter all or part of a PVC name. Alternatively, you can leave the **Search for** field empty to show all PVCs in an SLA. b. Select an item from the **All PVCs** menu to filter the results that match the search criteria. You can filter the results to show all PVCs, PVCs that are not in any SLA, and PVCs that are in a specific SLA. c. Select the checkbox for each PVC that you want to back up.

3. Click **Select an SLA Policy** and select one or more policies from the **SLA Policy** table. You can choose the default **Container** policy, or choose custom SLA policies that you defined.

   This action assigns the SLA policy to the selected PVCs. If you assign an SLA policy at the label or namespace level, any new PVCs that you create with the label or in the namespace will be automatically assigned to the SLA.

4. To create the job definition, click **Save**. The job will run at its scheduled time, as defined by the SLA policies that you selected.

   To start a job for a selected SLA policy immediately, use one of the following methods:

   - Scroll to the **SLA Policy Status** pane and locate the SLA policy in the table. Then, click **Actions** > **Start**.

     If you cannot find the SLA policy in the **Policy** column of the table, click **Auto Refresh** or the Refresh  icon.

   - Click **Jobs and Operations** > **Schedule**. Select the job and click **Actions** > **Start**. The job name is identified by the k8s_*SLA_name* format.

   When the job for the selected SLA policy runs, all PVCs that are associated with that SLA policy are included in the backup operation.

   Running on-demand backup jobs:

   To back up only selected PVCs, you can run an on-demand job. An on-demand job is a snapshot-only backup job that runs immediately.

   - For a single PVC, select the PVC and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is disabled.
   - For one or more PVCs, click **Create job**, select **Ad hoc backup**, and follow the instructions in “Running an ad hoc backup job” on page 627.

What to do next

To meet your organization's backup requirements, you can configure additional options for the SLA. For instructions, see “Specifying SLA options for Kubernetes backup jobs” on page 419.

Discontinuing SLA backups for a PVC: If you no longer want a PVC to participate in SLA backup jobs, remove the SLA policy assignment from the PVC by taking the following actions:

1. In the **Kubernetes Backup** pane, browse the clusters table, select the PVC for which you want to discontinue backup operations, and click **Select an SLA Policy**.

2. In the **SLA Policy** table, identify the SLA policies that are assigned to the PVC. The checkboxes for the assigned SLAs are selected.
3. Clear the checkbox for the SLA policy that you want to remove.
4. Click **Save**. The SLA policy is no longer assigned to the PVC.

Related concepts

"Backup and restore types" on page 403
Container Backup Support provides multiple types of backup and restore functions for your PVC and other cluster resources. You can use the IBM Spectrum Protect Plus user interface or the Kubernetes or OpenShift command line to initiate backup and restore operations.

"SLA policies" on page 405
Service level agreement (SLA) policies define how often snapshot backup and copy backup operations are run, and how long snapshots and copy backups are retained. You can set up custom SLAs that meet your operational requirements.

Backing up namespace-scoped resources

You can define backup jobs for your namespace-scoped resources in a Kubernetes environment and create container service level agreement (SLA) policies to run regular scheduled jobs. A container SLA policy specifies how often backup operations run, the target location for the backup operation, and how long snapshot or copy backups are retained.

Before you begin

Take the following actions:

- If you do not plan to use the default SLA policy for containers, configure a new SLA policy. For instructions, see "Creating an SLA policy for containers" on page 329.
- Ensure that appropriate roles and resource groups are assigned to the user who will run the backup job. The user must be assigned the Containers Admin role. For instructions, see Chapter 19, "Managing user access," on page 643.
- If a resource is associated with multiple SLA policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

About this task

To start protecting your namespace-scoped resources on a regular schedule, you must apply an SLA policy to those resources.

Procedure

To define a backup job for a namespace or namespace-scoped resources, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **Kubernetes**.
2. In the **Kubernetes Backup** pane, expand the cluster that contains the namespaces to be backed up.

   Tip: Toggle between `Namespace` or `Label` views by using the options from the **View** menu.

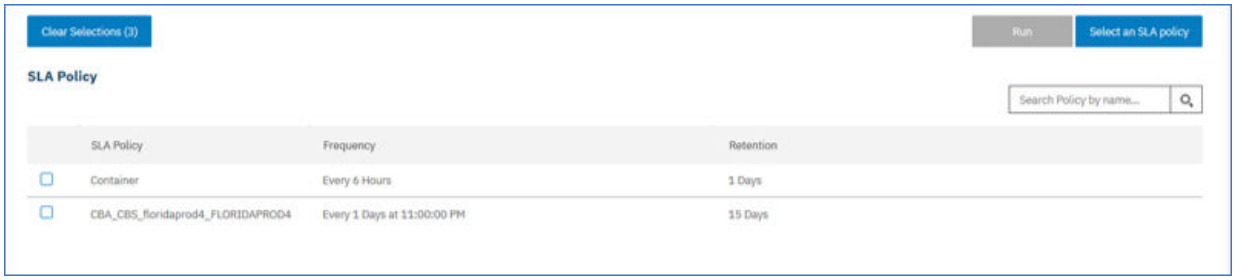


Figure 39. Viewing namespace resources, or labeled resources

3. Select the namespace-scoped resources to be included in the backup job. Alternatively, if you select a namespace, the associated PVC is also added to the SLA policy.
   - To back up all resources in a cluster, select the cluster.
   - To select namespaces in a cluster, expand the cluster by clicking the cluster name and then select the namespaces.
   - To select resources based on their assigned label, expand the cluster by clicking the cluster name, view the labels, and then select the label to be added to the SLA policy. Labels are indicated by the label icon

   For more information about resource icons, see "Backing up and restoring Kubernetes clusters" on page 408.
4. Click **Select an SLA Policy** to open the **SLA Policy** pane.

   You can select the default container SLA policy or any defined custom SLA policy that is listed in the table. Available policies are listed with details about the backup frequency and retention settings.



5. Click **Save** to save the SLA policy.

   To run the job immediately, scroll to the **SLA Policy Status** pane. Click **Actions** > **Start** for the policy that you want to run.

   If you assign an SLA policy at the resource level, any new PVCs that you create with that label or in the namespace is automatically assigned to the SLA policy.

What to do next

You can configure more options for the SLA policy. For instructions, see "Specifying SLA options for Kubernetes backup jobs" on page 419.

Note: To remove a namespace from an SLA backup job, complete these steps:

1. In the **Kubernetes Backup** pane, expand the cluster that contains the namespace resource that you want to remove from the policy.
2. Select the resources that you want to remove from the policy. Click **Select an SLA Policy**.
3. Clear the checkbox next to the namespace that you want to remove from the policy.
4. Click **Save**.

Related concepts

"Backup and restore types" on page 403
Container Backup Support provides multiple types of backup and restore functions for your PVC and other cluster resources. You can use the IBM Spectrum Protect Plus user interface or the Kubernetes or OpenShift command line to initiate backup and restore operations.

"SLA policies" on page 405

Service level agreement (SLA) policies define how often snapshot backup and copy backup operations are run, and how long snapshots and copy backups are retained. You can set up custom SLAs that meet your operational requirements.

Specifying SLA options for Kubernetes backup jobs

After you select a service level agreement (SLA) for your backup job, you can configure more options for that SLA. Additional SLA options include running scripts, excluding resources from the backup operation, and forcing a full base backup copy if required.

About this task

You can set up a regular backup job by creating an SLA policy. Container SLA policy definitions can back up any combination of persistent volumes, cluster-scoped and namespace-scoped resources. Copies can be stored locally on the cluster or backed up to the vSnap server.

Procedure

1. In the navigation panel, click **Manage Protection** > **Containers** > **Kubernetes**.
2. In the **Policy Options** column of the **SLA Policy Status** table, click the clipboard icon  for an SLA policy and set the following options:

   Pre-Script
   Select this checkbox to run a script before a job runs. Windows-based machines support batch and PowerShell scripts while Linux-based machines support shell scripts. Take one of the following actions:

   - To use a script server, select **Use Script Server** and choose an uploaded script from the **Script** or **Script Server** list.
   - To run a script on an application server, clear the **Use Script Server** checkbox, and choose an application server from the **Application Server** list.

   Scripts and script servers are configured by using the **System Configuration** > **Script** page.

   Post-Script
   Select this checkbox to run a script after a job runs. Windows-based machines support batch and PowerShell scripts while Linux-based machines support shell scripts. Take one of the following actions:

   - To use a script server, select **Use Script Server** and choose an uploaded script from the **Script** or **Script Server** list.
   - To run a script on an application server, clear the **Use Script Server** checkbox, and choose an application server from the **Application Server** list.

   Scripts and script servers are configured by using the **System Configuration** > **Script** page.

   Continue job/task on script error
   Select this checkbox to continue running the job when the script that is associated with the job fails.

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

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

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

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

Separate multiple filters with a semicolon.

3. Click **Save**.

Restoring Kubernetes container data

You can use the IBM Spectrum Protect Plus user interface to restore persistence volume claims (PVCs) from a snapshot or a copy backup. A snapshot restore operation is generally the faster method for restoring a PVC or resources.

Before you begin

Review the following restrictions:

- You cannot restore a snapshot backup to a different cluster or namespace.
- To help prevent issues with restore operations, do not manually delete snapshots of volumes that are protected by Container Backup Support.
- A PVC can be restored only in production mode.

About this task

To create a restore job, use the **Restore** wizard. You can create on-demand jobs that run once or you can create recurring restore jobs that run on a schedule.

The following table summarizes the destinations that are supported for snapshot and copy restore operations.

Table 85. Restore scenarios for PVCs

	Destination: original cluster		**Destination: alternative cluster**
Restore by specifying	**Snapshot restore**	**Copy restore**	**Copy restore**[1]
Original PVC and same namespace	Supported	Supported	Supported
Original PVC and another existing namespace	Does not apply	Supported	Supported
Original PVC and new namespace	Does not apply	Not supported	Not supported
New PVC and same namespace	Supported	Supported	Supported
New PVC and another existing namespace	Does not apply	Supported	Supported
New PVC and new namespace	Does not apply	Not supported	Not supported

You can also restore snapshot backups and copy backups to another storage class of the same storage provisioner. For example, if a PVC is defined to a storage class whose provisioner is `csi-rbd`, you can restore the snapshot backup of the PVC to another storage class whose provisioner is also `csi-rbd`.

[1] Snapshot restore jobs are not available when an alternative cluster is specified as the target.

Procedure

To restore PVCs from snapshot or copy backups, define a restore job by completing the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **Kubernetes**.
2. Click **Create job** to open the **Create job** page.
3. In the **Restore** pane, click **Select** to open the **Restore** wizard.

   You can also open the **Restore** wizard by clicking **Jobs and Operations** > **Create job**. Then, click **Select** in the **Restore** pane, and click **Kubernetes**.

   Tips:

   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, switch **Default setup** to **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
4. On the **Select source** page, browse the table and select the PVC that you want to restore by clicking the plus icon next to the PVC.

   The selected PVCs are displayed in the **Item** list. To remove an item from the list, click the minus icon next to the item.

   Alternatively, you can search for a PVC by specifying all or part of the PVC name in the **Search for** field and click the search icon.

   You can also toggle between the namespace and label views by clicking the options in the **View** menu.
5. On the **Source snapshot** page, use one of the following methods to select the source that you want to restore from:

Table 86. Methods for restoring a PVC

Source type	Steps
From a snapshot	a. Click **Origin** > **From Snapshot**. b. In the **Type of Restore** field, select the type of restore job that you want to create: **On-Demand** Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard. **Recurring** Creates a repeating restore job that runs on a schedule. This option is useful if you want to test a restore operation regularly. c. For on-demand restore only: Click the date range field and specify a range of dates. The available snapshot backups within that date range are shown. d. For on-demand restore only: If you are restoring a single PVC, select a snapshot from the list of available snapshots. If you are restoring more than one item, select a restore point for each item that you selected. e. Click **Next** to continue.

Table 86. Methods for restoring a PVC (continued)

Source type	Steps
From a copy backup	a. Click **Origin** > **From Copy**. b. In the **Type of Restore** field, select the type of restore job that you want to create: **On-Demand** Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard. **Recurring** Creates a repeating restore job that runs on a schedule. This option is useful if you want to test a restore operation regularly. c. In the **Restore Location Type** menu, select a type of location to restore data from: **Site** The site to which data was backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which data was copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which data was copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud service archive to which data was copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which data was copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. d. In the **Select a Location** menu, take one of the following actions: • If you are restoring data from a site, select one of the following restore locations: **Demo** The demonstration site from which to restore copy backups. This menu item is available only if you upgraded the product from IBM Spectrum Protect Plus Version 10.1.6 or earlier. **Primary** The primary site from which to restore copy backups. **Secondary** The secondary site from which to restore copy backups. • If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu. e. For on-demand restore only: Click the date range field and specify a range of dates. The available copy backups within that date range are shown. f. For on-demand restore only: If you are restoring a single PVC, select a backup from the list of available items. If you are restoring more than one PVC, select a restore point for each PVC that is listed. g. Click **Next** to continue.

6. Optional: On the **Restore method** page, enter a new name for the restored PVC.

   To designate the new PVC, you can enter up to 221 characters for the PVC name and a prefix of 32 characters. You can include alphanumeric characters, periods (.), and hyphens (-). The new PVC

name must not contain uppercase letters and must not end with a hyphen or a period. For example, `restored-pvc1` is a valid PVC name, where `restored-` is the prefix, and `pvc1` is the original PVC name.

If you do not enter a new name, the original PVC name is used.

Requirement: If you are restoring a PVC from a snapshot to the same PVC and namespace, ensure that you delete the original PVC before the restore job runs because the restore operation does not overwrite existing PVCs. Before you delete the original PVC, take a snapshot of the PVC to ensure that the latest changes are saved.

Click **Next** to continue.

7. On the **Set destination** page, configure the destination for the restored item.
   - For restore operations from snapshot backups, select from the following options:

     Original storage class
     Click this option to restore a snapshot backup of a PVC to the original storage class.

     Choose from available storage class
     Click this option to restore a snapshot backup of a PVC to a different storage class. Then, from the list of available storage classes, select a storage class of the same provisioner as the PVC.

     Original namespace
     Click this option to restore a snapshot backup to the original namespace.
   - For restore operations from copy backups, select from the following options:

     Restore to original cluster
     Click this option to restore a copy backup to the original cluster.

     Restore to alternate cluster
     Click this option to restore a copy backup to a different cluster. Then, select a cluster from the list of available clusters.

     Original storage class
     Click this option to restore a copy backup of a PVC to the original storage class.

     Choose from available storage class
     Click this option to restore a copy backup of a PVC to a different storage class. Then, from the list of available storage classes, select a storage class of the same provisioner as the PVC.

     Original namespace
     Click this option to restore a copy backup to the original namespace.

     Choose from available namespaces
     Click this option to restore a copy backup to a different namespace. Then, select a namespace from the list of available namespaces.
8. Optional: On the **Job options** page, configure additional options for the restore job:

   Run cleanup immediately on job failure
   This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

   Allow session overwrite
   Enable this option to allow a scheduled session of a recovery job to force a pending session to clean up associated resources so that the new session can run.

   Continue with restores of other selected PVCs even if one fails
   If one PVC is not successfully restored, the restore job continues for all other PVCs that are being restored. If this option is not enabled, the restore job stops when the recovery of a PVC fails.

   Click **Next** to continue.
9. Optional: If you are running the wizard in advanced setup mode, on the **Apply scripts** page, specify scripts to run before or after an operation runs at the job level. Batch and PowerShell scripts are supported.

Pre-Script
Select this checkbox to choose an uploaded script and an application or script server where the pre-script will run. To select an application server where the pre-script will run, clear the **Use Script Server** checkbox. Scripts and script servers are configured on the **System Configuration** > **Script** page.

Post-Script
Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server where the post-script will run, clear the **Use Script Server** checkbox. Scripts and script servers are configured on the **System Configuration** > **Script** page.

Continue job/task on script error
Select this checkbox to continue running the job when the script that is associated with the job fails.

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

If you clear this checkbox, the restore operation is not attempted, and the pre-script or post-script task status is reported as FAILED.

10. If you specified a recurring job in Step "5" on page 421, on the **Schedule** page, enter a name for the job, specify the frequency and start time for the job, and click **Next** to continue.

    The **Schedule** page is not displayed for on-demand restore jobs.
11. On the **Review** page, review your restore job settings and click **Submit** to create the job.

Results

For on-demand jobs, the job begins after you click **Submit**, and the **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view the progress of the restore operation, expand the job. You can also download the log file by clicking **Download .zip**.

For recurring jobs, the job begins at the scheduled start time when you start the schedule on the **Jobs and Operations** > **Schedule** page.

You can view the status of running jobs on the **Jobs and Operations** > **Running Jobs** page.

What to do next

To verify whether a PVC is restored, issue the following **kubectl** command:

```
kubectl get pvc restored_pvc –n namespace
```

where *restored_pvc* specifies the name of the restored PVC, and *namespace* specifies the namespace of the restored PVC.

Restoring Kubernetes cluster-scoped and namespace-scoped resources

You can restore Kubernetes cluster-scoped and namespace-scoped resources from a snapshot or a copy backup. A snapshot restore operation is generally the faster method for restoring these resources. You can also choose to restore a copy backup to another namespace or cluster, or restore backups to another storage class of the same storage provisioner.

Before you begin

Review the following restrictions:

- You cannot restore a snapshot backup to a different cluster or namespace.
- To help prevent issues with restore operations, do not manually delete snapshots of volumes that are protected by Container Backup Support.

About this task

To create the restore job, use the **Restore** wizard. You can create on-demand jobs that run one time or you can create recurring restore jobs that run on a schedule.

The following table summarizes the destinations that are supported for cluster-scoped resource snapshot and copy restore operations.

Table 87. Restore scenarios for cluster-scoped resources

Restoring data to this location	Destination: Original cluster		Destination: Alternate cluster
	Snapshot restore	**Copy restore**	**Copy restores only**
Original namespace or cluster	Supported	Supported	Supported
Original namespace or to an alternative cluster	Supported	Supported	Supported
Alternative namespace on the original cluster. This option is for copy backups only.	Does not apply	Supported	Not Supported
Alternative namespace on an alternative cluster. This option is for copy backups only.	Does not apply	Supported	Supported

Procedure

To restore your cluster-scoped and namespace-scoped resources from snapshots or copy backups, define a restore job by completing the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **Kubernetes**.
2. Click **Create job** to open the **Create job** page.
3. In the **Restore** pane, click **Select** to open the **Restore** wizard.

   Tips:

   - For a summary of your wizard selections, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, set the mode to **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
4. On the **Select source** page, browse the table and expand the cluster to show the resources that are available for the restore operation. Click the plus icon next to the resources you want to add to the job.

   The resources you select are displayed in the **Item** list. If you need to remove an item from the list, click the minus icon next to the item.

   Alternatively, you can search for a resource by specifying all or part of the resource name in the **Search for** field.

   You can also toggle between the namespace and label views by clicking the options in the **View** menu.
5. On the **Source snapshot** page, use one of the following methods to select the source that you want to restore from:

Table 88. Methods for restoring a PVC or namespace resource

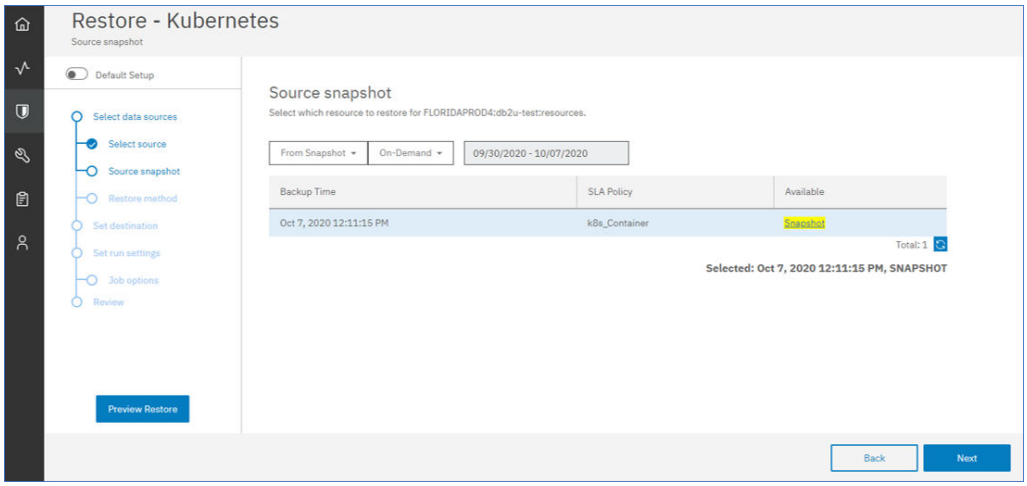
Source type	Steps
From a snapshot	a. Choose **From Snapshot** or **From Copy**. b. Click **Type of Restore**, and select the type of restore job that you want to create: **On-Demand** Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard. **Recurring** Creates a repeating restore job that runs on a schedule. This option is useful if you want to test a restore operation regularly. When you select the backup for the restore job, summary information shows the timestamp for that backup as shown.  c. For on-demand restore jobs, you can specify the date range to show the available snapshots backups within that date range. d. For on-demand restore jobs when you are restoring a single resource, you can select a snapshot from the list of available items. If you are restoring more than one item, you can select a restore point for each item that you selected. e. Click **Next** to continue.

Table 88. Methods for restoring a PVC or namespace resource (continued)

Source type	Steps
From a copy backup	a. Click **Origin** > **From Copy**. b. In the **Type of Restore** field, select the type of restore job that you want to create: **On-Demand** Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard. **Recurring** Creates a repeating restore job that runs on a schedule. This option is useful if you want to test a restore operation regularly. c. In the **Restore Location Type** menu, select a type of location to restore data from: **Site** The site where data was backed up to. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service where data was copied to. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage**. **Repository server** The repository server where data was copied to. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server**. **Cloud service archive** The cloud archive service where data was copied to. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server where data was copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. d. In the **Select a Location** menu, take one of the following actions: • If you are restoring data from a site, select one of the following restore locations: **Demo** The demonstration site from which to restore copy backups. This menu item is available only if you previously backed up data to a demo site in IBM Spectrum Protect Plus Version 10.1.6 or earlier. **Primary** The primary site from which to restore copy backups. **Secondary** The secondary site from which to restore copy backups. • If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu. e. For on-demand restore only: Click the date range field and specify a range of dates to show the available copy backups within that date range. f. For on-demand restore only: If you are restoring a single resource, select a backup from the list of available items. If you are restoring more than one resource for example, select a restore point for each resource that is listed. g. Click **Next** to continue.

6. Skip the **Restore method** page by clicking **Next**.
   You cannot add a new PVC name for restoring cluster-scoped and namespace-scoped resources.
7. On the **Set destination** page, configure the destination for the restored item.

- Snapshot backups can only be restored to their original storage classes and namespaces. For restore operations from Snapshot backups, select the following options:

  Storage class
  This value is not applicable.

  Original namespace
  Choose to restore the resource from the snapshot to the `Original Namespace`.

- Copy backups can be restored to their original cluster, storage classes, or namespace. They can also be restored to a different cluster, storage class or namespace. For restore operations from copy backups, select from the following options:

  Restore to original cluster
  Select this option to restore a copy backup to the original cluster.

  Restore to alternate cluster
  Select this option to restore a copy backup to a different cluster which you can pick from a list of available clusters.

  Original storage class
  This value is not applicable.

  Choose from available storage class
  Select this option to restore the copy backup to a different storage class, which you can pick from the list of available storage classes.

  Restriction: If you are restoring multiple resources and the associated storage classes are from different storage provisioners, you cannot select a different storage class.

  Original namespace
  Select this option to restore the copy backup to the original namespace.

  Choose from available namespaces
  Select this option to restore the copy backup to a different namespace which you can pick from the list of available namespaces.

8. On the **Job options** page, configure more options for the restore job:

   Run cleanup immediately on job failure
   This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

   Allow session overwrite
   Enable this option to allow a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run.

   Continue with restores of other selected resources even if one fails
   If one resource is not successfully restored, the restore job continues for all other resources that are being restored. If this option is not enabled, the restore job stops when the recovery of a resource fails.

   Click **Next** to continue.

9. Optional: If you are running the wizard in advanced setup mode, on the **Apply scripts** page, specify scripts to run before or after an operation runs at the job level. Batch and PowerShell scripts are supported.

   Pre-Script
   Select this checkbox to choose an uploaded script and an application or script server where the pre-script will run. To select an application server where the pre-script will run, clear the **Use Script Server** checkbox. Scripts and script servers are configured on the **System Configuration** > **Script** page.

   Post-Script
   Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server where the post-script will run, clear the **Use**

Script Server checkbox. Scripts and script servers are configured on the **System Configuration** > **Script** page.

Continue job/task on script error
Select this checkbox to continue running the job when the script that is associated with the job fails.

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

If you clear this check box, the restore operation is not attempted, and the pre-script or post-script task status is reported as FAILED.

10. If you specified a recurring job in Step "5" on page 425, on the **Schedule** page, enter a name for the job, specify the frequency and start time for the job, and click **Next** to continue.

    The **Schedule** page is not displayed for on-demand restore jobs.
11. On the **Review** page, review your restore job settings and click **Submit** to create the job.

Results

For on-demand jobs, a job begins after you click **Submit**, and the **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view progress of the restore operation, expand the job. You can also download the log file by clicking **Download .zip**.

For recurring jobs, a job begins at the scheduled start time when you start the schedule on the **Jobs and Operations** > **Schedule** page.

All running jobs are viewable on the **Jobs and Operations** > **Running Jobs** page.

What to do next

To verify the contents of the restore job, you can check the Velero logs by issuing the following **kubectl** command: v

```
velero restore logs <restorename> -n spp-velero --insecure-skip-tls-verify
```

where *restorename* specifies the name of the restored resource.

Expiring Kubernetes job sessions

You can expire a Kubernetes backup job session to override the retention settings that were assigned when a snapshot or copy backup was created. When a job session is expired, the restore point (the snapshot or copy backup) will be removed during the next maintenance job.

About this task

Complete this task if you do not want to wait for a job session to automatically expire according to the retention setting of the assigned service level agreement policy.

Procedure

To expire a Kubernetes job session, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **IBM Spectrum Protect Plus** > **Restore Point Retention**.
2. On the **Backup Sessions** tab, search for a job session or restore point. Alternatively, on the **Virtual Machines / Databases** tab, select **Applications**, and search for a catalog entry by entering the name.

   Names can be searched by entering partial text, by using the asterisk (*) as a wildcard character, or by using the question mark (?) for pattern matching. For more information about using the search function, see Appendix A, "Search guidelines," on page 681.

3. Optional: If you are searching from the **Backup Sessions** tab, use filters to narrow your search for snapshots or copy backups. You can also specify the date range when the associated backup job started.
   a) In the **Type** field, select **Application**.
   b) In the **Subpolicy Type** field, select **Snapshot** to search for snapshot backups or select **Backup** to search for copy backups.
   c) If needed, click the **Backup Time Range** field and select a date range to search on.
4. Click the search icon .
5. In the search results, select the job session that you want to expire.
6. If you are on the **Backup Sessions** tab, from the **Actions** menu, select one of the following options:
   - To expire a single job session, click **Expire**.
   - To expire all unexpired job sessions for the selected job, click **Expire All Job Sessions**.

   If you are on the **Virtual Machines / Databases** tab, select **Containers** in drop-down menu and search for a catalog item by name. Then, click the deletion icon for the resource that you want to expire.
7. Follow the instructions in the confirmation window and click **OK**.

Related tasks

"Managing IBM Spectrum Protect Plus restore points" on page 616
You can use the **Restore Point Retention** pane to search for restore points in the IBM Spectrum Protect Plus catalog by backup job name, view their creation and expiration dates, and override the assigned retention.

Backing up and restoring OpenShift clusters

To protect persistent volumes and other OpenShift resources that are attached to a cluster, create service level agreement (SLA) polices and create jobs for backup and restore operations in the IBM Spectrum Protect Plus user interface.

Ensure that your OpenShift environment meets the system requirements in "Container Backup Support requirements" on page 75.

To protect cluster-scoped and namespace-scoped resources in your OpenShift container workloads, ensure that the Velero tool is installed. For instructions, see "Installing Velero to protect resources" on page 198.

The following table shows the icon and naming convention for each type of resource that is displayed in the **OpenShift Backup** pane, as well as a description of each type of resource:

Table 89. OpenShift resource types

Naming convention	Resource type	Description
clustername	Cluster	An OpenShift cluster.
clustername:resources	Cluster resources	Cluster-scoped resources such as PersistentVolumes, ClusterRoles, StorageClasses, CSIDrivers, VolumeSnapshotClasses, and CustomResourceDefinitions.
project	Project	A project in the cluster.
clustername:*project*:resources	Project resources	Project-scoped resources such as PersistentVolumeClaims, pods, containers, ConfigMaps, secrets, services, and deployments.

Table 89. OpenShift resource types (continued)		
Naming convention	**Resource type**	**Description**
clustername:*project*:*pvcname*	Persistent volume claim (PVC)	PVCs in a project.

Related concepts

"Overview of Container Backup Support" on page 401
IBM Spectrum Protect Plus Container Backup Support protects data of persistent volumes, namespace-scoped resources, and cluster-scoped resources that are associated with containers in a Kubernetes or Red Hat OpenShift environment. You can run snapshot backup operations to create locally stored backup copies in the cluster, or you can run backup copies to the vSnap server for longer-term retention.

"Protecting containers by using the command line" on page 466
As an application developer in a Kubernetes or OpenShift environment, you can use the command-line interface to back up and restore container data, and to query the status of Container Backup Support requests.

Registering an OpenShift cluster

If necessary, you can use the IBM Spectrum Protect Plus user interface to manually register an OpenShift cluster or to modify the properties of a registered OpenShift cluster.

About this task

After Container Backup Support is installed, the application host for the Container Backup Support container is automatically registered upon startup of the cluster host in OpenShift. When a cluster is registered with IBM Spectrum Protect Plus, an inventory of the resources in the cluster is automatically captured, enabling you to complete backup and restore jobs, as well as run reports.

However, if the automatic registration was unsuccessful or if a registered cluster was accidentally unregistered, you can manually register the cluster by using the IBM Spectrum Protect Plus user interface.

You can also modify the properties of the registered cluster, such as changing the SSH port that is used to connect to the Container Backup Support container agent service.

For example, if you use a load balancer to distribute the workload in your cluster, you can edit the load balancer to use the port number for the Container Backup Support agent container service. You can then register the load balancer and port number with IBM Spectrum Protect Plus so that you do not have to configure the port number again.

Procedure

To manually register a cluster or to modify cluster properties, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **OpenShift**.
2. In the **OpenShift** page, click **Manage clusters**.
3. Take one of the following actions:
   - To manually register a cluster, click **Add cluster**.
   - To update existing cluster properties, in the list of host addresses, click the edit icon for the cluster host that you want to update.
4. Update the fields in the **Application Properties** section:

   Cluster Name
   The name of the cluster host or load balancer for the Container Backup Support container. You can enter a host name or IP address.

The cluster name must match the value that is used for the **clusterName** parameter in the `baas-values.yaml` file.

Host Address
The host address for the cluster host. The cluster can be registered using any one of the IP addresses in the cluster. Enter any IP address that is listed in the **clusterAPIServerips** parameter in the `baas-values.yaml` file. You can enter an IP address or a fully qualified domain name.

Port Number
The SSH port for the connection to the Container Backup Support agent container service.

By default, the port is automatically assigned by OpenShift during installation of Container Backup Support. To obtain this port number, issue the following command at the **oc** command line:

```
oc get service -n baas | grep baas-spp-agent
```

The output is similar to the following example:

```
baas-spp-agent            NodePort    10.110.235.90    <none>        22:31299/TCP   111m
```

The port number is the numerical string that follows `22:`. In the example, the port number is 31299.

Use existing user
Select this checkbox to use a previously entered username and password for the cluster host. Select a username from the **Select user** list.

User ID
Enter the username for the application host. The username must match the value that you specified for the DATAMOVER_USERNAME installation variable in the `baas-options.sh` file. For more information, see "Setting up the installation variables" on page 206.

The credentials will be added to the list of existing users. This field is not available if you are using an existing user.

Password
Enter the password for the application host. The password must match the value that you specified for the DATAMOVER_PASSWORD installation variable in the `baas-options.sh` file. For more information, see "Setting up the installation variables" on page 206.

The credentials will be added to the list of existing users. This field is not available if you are using an existing user.

5. Optional: Populate the field in the **Options** section:

Maximum concurrent PVCs
Set the maximum number of PVC snapshots or copy backups to create concurrently. Cluster performance is impacted when you back up many PVCs concurrently, as each PVC uses multiple threads and consumes bandwidth when copying data. Use this option to control the impact on cluster resources and minimize the impact on production operations.

The default value is 10.

6. Click **Save**. IBM Spectrum Protect Plus confirms a network connection, adds the cluster to the IBM Spectrum Protect Plus database, and then catalogs the cluster resources, including projects and PVCs.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a network administrator to review the connection.

What to do next

To verify that the clusters are updated, review the job log. In the navigation panel, click **Jobs and Operations**. Click the **Running Jobs** tab, and look for the most recent `Application Server Inventory` log entry. You can specify a filter to show only inventory jobs by clicking the filter icon, selecting **Inventory**, and clicking **Apply**.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name. If the status of inventory job status is `Partial`, click **Job Log** and review the log entries to find the error.

Clusters must be detected to ensure that their resources can be backed up. You can run a manual inventory at any time to detect updates in cluster resources. For instructions about running a manual inventory, see "Detecting OpenShift cluster resources" on page 433. For instructions about scheduling OpenShift backup jobs, see "Backing up OpenShift container data" on page 436.

Detecting OpenShift cluster resources

OpenShift cluster resources are automatically detected after the cluster is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the cluster was added.

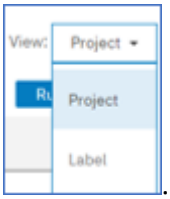
About this task

Run an inventory job periodically to help to ensure that all cluster resources are detected and that they are available for backup operations. To verify that projects in a cluster are detected, you can click the name of the cluster to expand it.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **OpenShift**.
2. Optional: In the list of clusters, click a cluster name to show the resources in that cluster.
3. To view the resources that are detected, click a hyperlink-enabled name to navigate to the different levels of a cluster. Click a breadcrumb to move to a previous level in the cluster.

   You can view the cluster resources by project or by label by changing the View options, 

.

   For more information about working with labels, see "Working with OpenShift labels" on page 434.

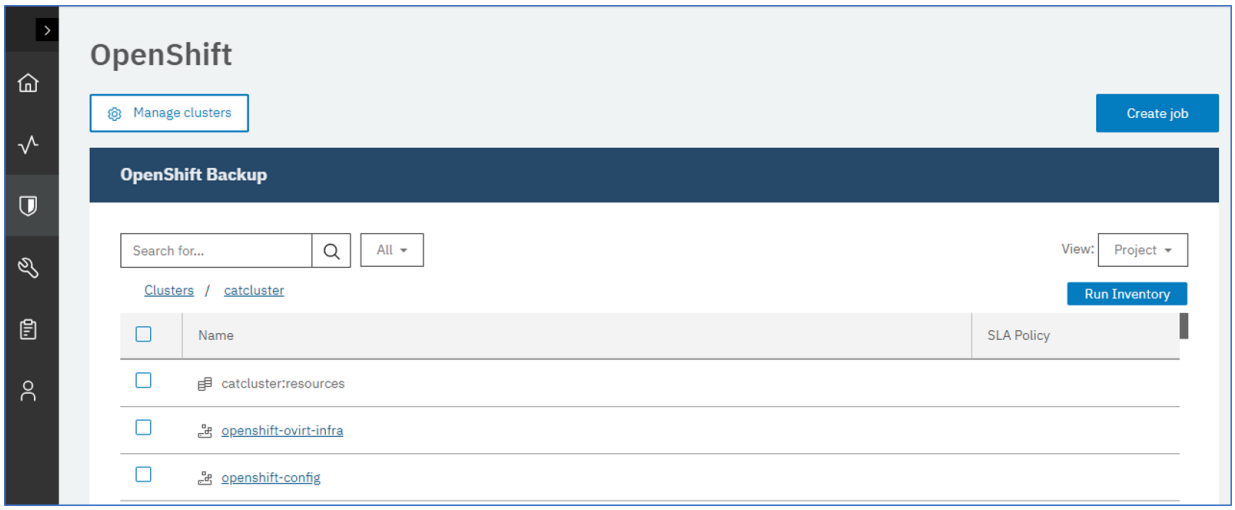


Figure 40. RedHat OpenShift cluster resources

The following table shows the icon and naming convention for each type of resource that is displayed in the **OpenShift Backup** pane, as well as a description of each type of resource:

Table 90. OpenShift resource types		
Naming convention	**Resource type**	**Description**
clustername	Cluster	An OpenShift cluster.
clustername:resources	Cluster resources	Cluster-scoped resources such as PersistentVolumes, ClusterRoles, StorageClasses, CSIDrivers, VolumeSnapshotClasses, and CustomResourceDefinitions.
project	Project	A project in the cluster.
clustername:*project*:resources	Project resources	Project-scoped resources such as PersistentVolumeClaims, pods, containers, ConfigMaps, secrets, services, and deployments.
clustername:*project*:*pvcname*	Persistent volume claim (PVC)	PVCs in a project.

An example might include a cluster, a project in that cluster, and a PVC in that project such as *CLUSTERPROD1*:*db2-project*:*pvc-1*.

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

If you do not select a cluster in the list of clusters and click **Run Inventory**, an inventory job is started for all clusters.

What to do next

To monitor the inventory job, in the navigation panel, click **Jobs and Operations**. Click the **Running Jobs** tab, and look for the most recent `Application Server Inventory` log entry. You can specify a filter to show only inventory jobs by clicking the filter icon, selecting **Inventory**, and clicking **Apply**.

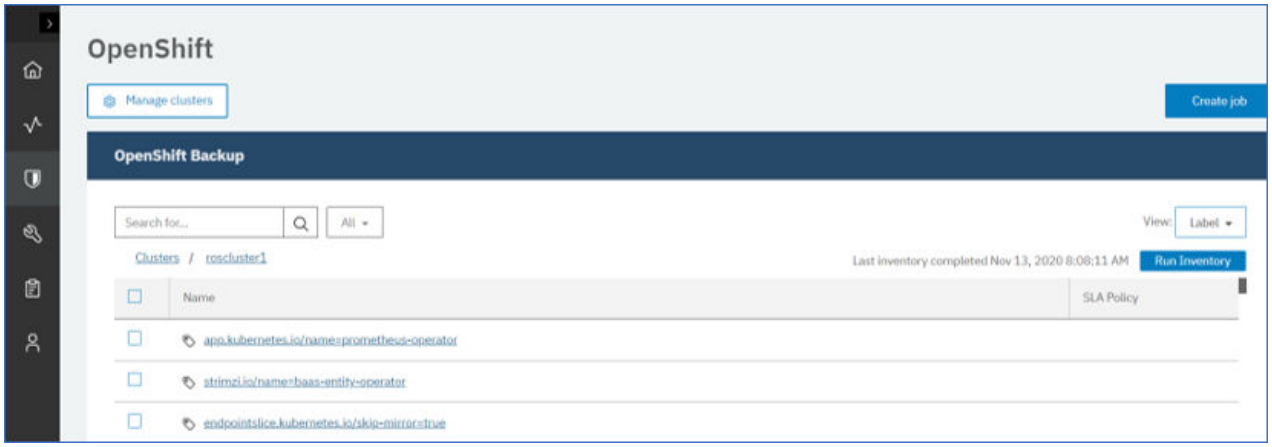
Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name. If the status of an inventory job is `Partial`, click **Job Log** and review the log entries to find the error.

Working with OpenShift labels

Labels are used in OpenShift environments to label key pairs, such as pods. Labels are used to specify identifying attributes of objects that are meaningful and relevant within the environment, and are generally used to organize subsets of objects. With IBM Spectrum Protect Plus, you can view and manage your OpenShift resources by using their labels.

Inventory

To view your resources by their assigned labels, change the **View** setting to `Label`. When you expand the cluster after you changed the view, the labeled resources in that cluster are listed in the table as shown.



Backing up and SLA policy

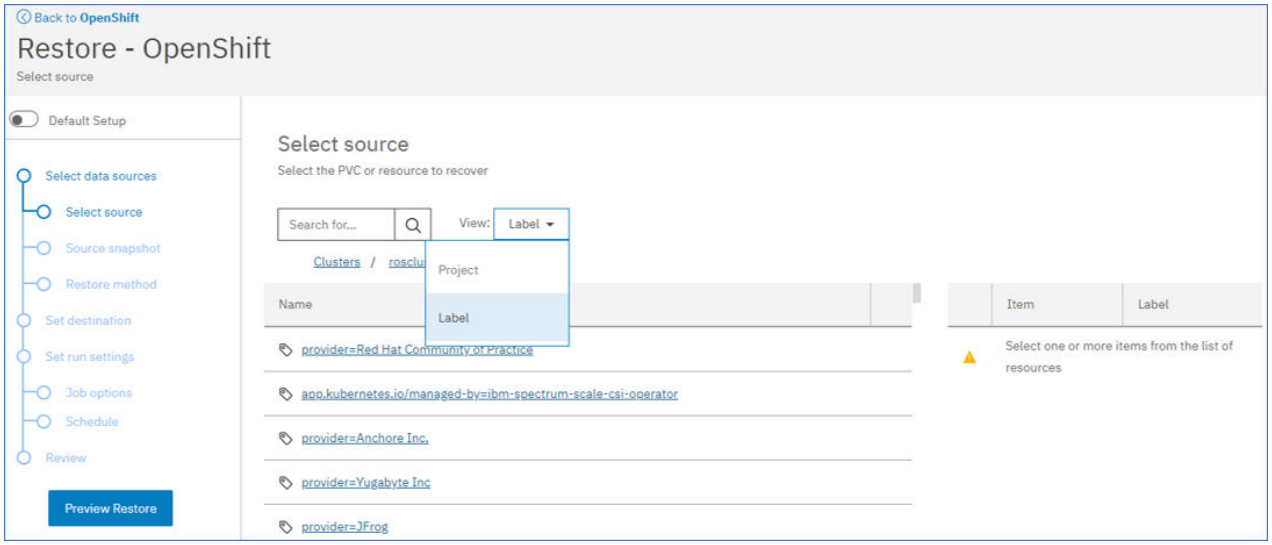
When you are setting up an SLA policy, you can select the labeled resources that you want to add to any predefined policy. When you assign the SLA policy and click save, the labeled resources are added to the policy for the next time it is run.

Ad hoc backup job

When you create an ad hoc backup job, you first select the SLA policy that you want to assign. Then, you can select the labeled resources that you want to add to the job.

Restoring labeled resources

When you are creating a restore job, you can change the view to view your resources by their assigned labels. From within the restore wizard, change the **View** setting to `Label`. When you expand the cluster after you change the view, the labeled resources in that cluster are listed in the table as shown.



Job log files

When you run a backup or restore operation for your labeled resources, the messages that are collected for the job log file list the labels in that operation.

Testing the connection to an OpenShift cluster

You can test the connection to an OpenShift cluster that you added to IBM Spectrum Protect Plus. The test function verifies communication with the cluster and tests domain name server (DNS) settings between the IBM Spectrum Protect Plus server and the cluster.

Procedure

To test the connection to a cluster, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **OpenShift**.
2. Click **Manage clusters**.

   The list of available clusters are displayed.
3. Scroll through the list and locate the cluster that you want to test.
4. Click the **Actions** menu that is associated with the cluster and select **Test**.

   The test report shows you a list of the tests that ran and the status.

Backing up OpenShift container data

You can use the IBM Spectrum Protect Plus user interface to define service level agreement backup jobs to protect container data such as persistent volumes, project-scoped resources, and cluster-scoped resources. You can also create application-consistent backups of container data.

Backing up persistent volumes in an OpenShift cluster

You can use the IBM Spectrum Protect Plus user interface to define backup jobs that run according to a service level agreement (SLA) policy. The SLA policy specifies how often backup operations are run and how long snapshot or copy backups are retained.

Before you begin

Take the following actions:

- Ensure that persistent volume claims (PVCs) for the volumes that you want to protect are formatted. Backup requests are directed to PVCs. Backup operations of raw block volumes are not supported.
- If you do not plan to use the default SLA policy for containers, ensure that you configure an SLA policy. For instructions, see "Creating an SLA policy for containers" on page 329.
- Assign appropriate roles and resource groups to users who will be running backup and restore operations. Grant users access to resources and roles by using the Accounts pane. The user must be assigned the Containers Admin role. For instructions, see Chapter 19, "Managing user access," on page 643.
- If a PVC is associated with multiple SLA policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.
- For PVCs that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the VMs are protected during a PVC backup job. For the data to be consistent, ensure that you power off the VMs before running the job.

About this task

To start protecting your PVCs on a regular schedule, you must apply an SLA policy to your PVC. The SLA policy also defines the backup target locations for your PVCs.

Procedure

To define an SLA backup job for one or more PVCs, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **OpenShift**.

2. In the **OpenShift Backup** pane, select the PVCs that you want to back up. You can use one of the following methods:

Method	Steps
To back up all PVCs in a cluster	Select the checkbox for a cluster name. A cluster is identified by the cluster icon .
To back up PVCs that are associated with a project	a. Click **View** > **Project**. b. Click the name of a cluster that contains the PVCs that you want to back up. The list of projects within the cluster is displayed. A project is identified by the project icon . c. To back up all PVCs in the project, select the checkbox for the project. To back up individual PVCs, click the project link and select the checkbox for each PVC that you want to back up. A PVC is identified by the PVC icon .
To back up PVCs that are associated with a label	a. Click **View** > **Label**. b. Click the name of a cluster that contains the PVCs that you want to back up. The list of labels within the cluster is displayed. A label is shown as a key-value pair and identified by the label icon . c. To back up all PVCs that are assigned to a label, select the checkbox for a label. To back up individual PVCs, click the label name and select the checkbox for each PVC that you want to back up.
To use the search function to filter the list of PVCs by SLA	a. Enter your search criteria in the **Search for** field. You can enter all or part of a PVC name. Alternatively, you can leave the **Search for** field empty to show all PVCs in an SLA. b. Select an item from the **All PVCs** menu to filter the results that match the search criteria. You can filter the results to show all PVCs, PVCs that are not in any SLA, and PVCs that are in a specific SLA. c. Select the checkbox for each PVC that you want to back up.

3. Click **Select an SLA Policy** and select one or more policies from the **SLA Policy** table. You can choose the default **Container** policy, or choose custom SLA policies that you defined.

   This action assigns the selected SLA policy to the selected PVCs. If you assign an SLA policy at the label or project level, any new PVCs that you create with the label or in the project will be automatically assigned to the SLA.

4. To create the job definition, click **Save**. The job will run at its scheduled time, as defined by the SLA policies that you selected.

   To start a job for a selected SLA policy immediately, use one of the following methods:

   - Scroll to the **SLA Policy Status** pane and locate the SLA policy in the table. Then, click **Actions** > **Start**.

     If you cannot find the SLA policy in the **Policy** column of the table, click **Auto Refresh** or the Refresh icon.

   - Click **Jobs and Operations** > **Schedule**. Select the job and click **Actions** > **Start**. The job name is identified by the `openshift_`*`SLA_name`* format.

   When the job for the selected SLA policy runs, all PVCs that are associated with that SLA policy are included in the backup operation.

   Running on-demand backup jobs:

   To back up only selected PVCs, you can run an on-demand job. An on-demand job is a snapshot-only backup job that runs immediately.

- For a single PVC, select the PVC and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is disabled.
- For one or more PVCs, click **Create job**, select **Ad hoc backup**, and follow the instructions in "Running an ad hoc backup job" on page 627.

What to do next

If necessary, you can configure additional options for the SLA. For instructions, see "Specifying SLA options for OpenShift backup jobs" on page 442.

Discontinuing SLA backups for a PVC: If you no longer want a PVC to participate in SLA backup jobs, remove the SLA policy assignment from the PVC by taking the following actions:

1. In the **OpenShift Backup** pane, browse the clusters table, select the PVC for which you want to discontinue backup operations, and click **Select an SLA Policy**.
2. In the **SLA Policy** table, identify the SLA policies that are assigned to the PVC. The checkboxes for the assigned SLAs are selected.
3. Clear the checkbox for the SLA policy that you want to remove.
4. Click **Save**. The SLA policy is no longer assigned to the PVC.

Related concepts

"Backup and restore types" on page 403
Container Backup Support provides multiple types of backup and restore functions for your PVC and other cluster resources. You can use the IBM Spectrum Protect Plus user interface or the Kubernetes or OpenShift command line to initiate backup and restore operations.

"SLA policies" on page 405
Service level agreement (SLA) policies define how often snapshot backup and copy backup operations are run, and how long snapshots and copy backups are retained. You can set up custom SLAs that meet your operational requirements.

Backing up namespace-scoped resources

You can define backup jobs for your namespace resources in an OpenShift environment and create container service level agreement (SLA) policies to run regular jobs. The container SLA policy specifies how often backup operations run, and how long snapshot or copy backups are retained.

Before you begin

Take the following actions:

- If you do not plan to use the default SLA policy for containers, ensure that you configure an SLA policy. For instructions, see "Creating an SLA policy for containers" on page 329.
- Ensure that appropriate roles and resource groups are assigned to the user who will run the backup job. Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. The user must be assigned the Containers Admin role. For instructions, see Chapter 19, "Managing user access," on page 643.
- If a resource is associated with multiple SLA policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.
- For OpenShift clusters that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the custom resource data for the VMs are protected when resources are backed up. For the data to be consistent, ensure that you power off the VMs before running a backup job.

About this task

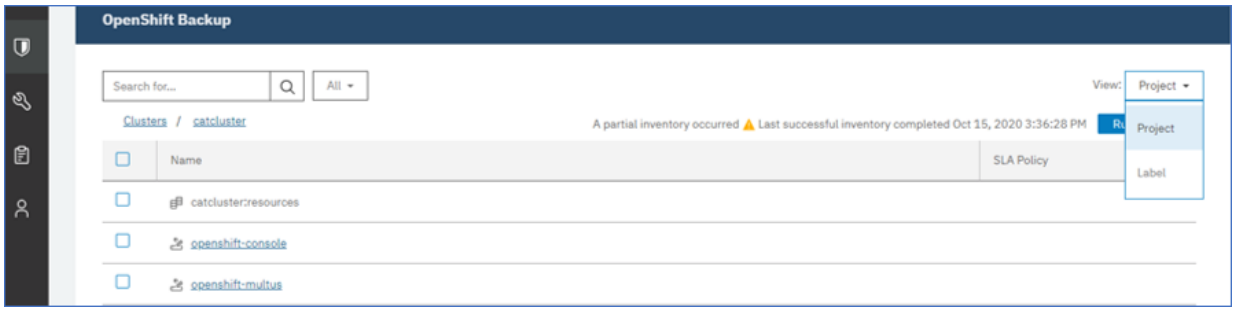
To start protecting your namespace resources on a regular schedule, you must apply an SLA policy to those resources. In the SLA policy definition, the backup target locations for your resources are set.

Procedure

To define a backup job for a namespace or namespace-scoped resources, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **OpenShift**.
2. In the **OpenShift Backup** pane, expand the cluster that contains the namespaces to be backed up.

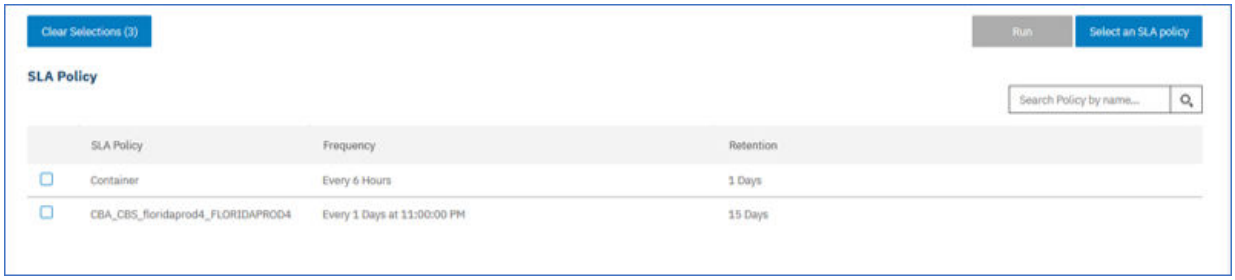
   Tip: Toggle between `Project` or `Label` views by using the options from the **View** menu.



3. Select the required resources to be included in the backup operation.
   - Select the entire cluster to backup all resources in that cluster.
   - Expand the cluster by clicking the cluster name to show the namespaces in that cluster. Select the namespaces to be backed up.
   - Expand the cluster by clicking the cluster name to show the labels in that cluster. Select the labeled resources to be backed up. Labels are indicated by the label icon .

   For more information about the different resource icons, see "Backing up and restoring OpenShift clusters" on page 430.
4. Click **Select an SLA Policy** to open the **SLA Policy** pane.

   You can choose the default container SLA policy. Alternatively, you can choose any defined custom SLA policy that is listed in the table. The policies that are available are listed with details about the policy frequency and retention settings.



5. Click **Save** to save the SLA policy.

   To run the job immediately, scroll to the **SLA Policy Status** pane. Click **Actions** > **Start** for the policy you want to run.

   If you assign an SLA policy at the label or namespace level, any new PVCs that you create with the label or in the namespace will be automatically be assigned to the SLA policy.

What to do next

You can configure extra options for the SLA policy. For instructions, see "Specifying SLA options for Kubernetes backup jobs" on page 419

Note: If you want to remove a namespace from an SLA backup job, follow these steps actions:

1. In the **OpenShift Backup** pane, expand the cluster that contains the namespace resource that you want to remove from the policy.
2. Select the resources you want to remove from the policy. Click **Select an SLA Policy**.

3. Deselect the checkbox next to the namespace you want to remove from the policy.
4. Click **Save**.

Related concepts

"Backup and restore types" on page 403
Container Backup Support provides multiple types of backup and restore functions for your PVC and other cluster resources. You can use the IBM Spectrum Protect Plus user interface or the Kubernetes or OpenShift command line to initiate backup and restore operations.

"SLA policies" on page 405
Service level agreement (SLA) policies define how often snapshot backup and copy backup operations are run, and how long snapshots and copy backups are retained. You can set up custom SLAs that meet your operational requirements.

Backing up application-consistent data

You can define backup jobs for your application namespaces in an OpenShift environment and create container service level agreement (SLA) policies to run scheduled backup jobs. By defining Velero pre-backup and post-backup hooks as annotations on the application pod, you can create application-consistent backups of your data.

Before you begin

Take the following actions:

- If you do not plan to use the default SLA policy for containers, configure an SLA policy. For instructions, see "Creating an SLA policy for containers" on page 329.
- Ensure that appropriate roles and resource groups are assigned to the user who runs the backup jobs. Before an IBM Spectrum Protect Plus user can run backup and restore jobs, roles and resource groups must be assigned to the user. The user must be assigned the Containers Admin role. For instructions, see Chapter 19, "Managing user access," on page 643.
- If a resource is associated with multiple SLA policies, ensure that the policies are not scheduled to run concurrently. Either schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.
- For OpenShift clusters that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, ensure that you power off the VMs before running a backup job.

About this task

In an application-consistent backup operation, Velero pre-backup and post-backup hooks are used to quiesce an application before a snapshot is created, and to resume the application when the snapshot is completed.

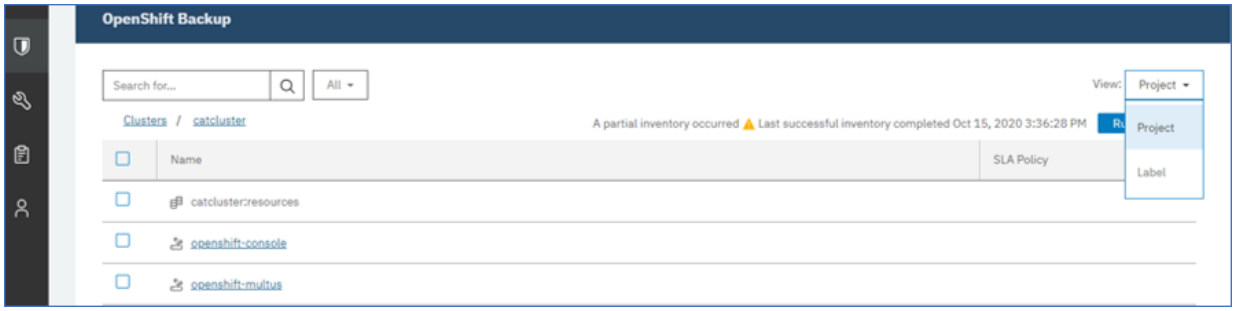
To start protecting your application in a consistent manner on a regular schedule, you must apply an SLA policy to those namespaces and persistent volume claims (PVCs) where the application is running. In addition, you must specify pre and post hooks as annotations on the application pods. For instructions about specifying these application hooks, see Backup Hooks in the Velero documentation. The backup target locations for your application namespaces and PVCs are set in the SLA policy definition.

Procedure

To define an SLA backup job for one or more application namespaces within the corresponding PVCs, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **OpenShift**.
2. In the **OpenShift Backup** pane, expand the cluster that contains the namespaces to be backed up.

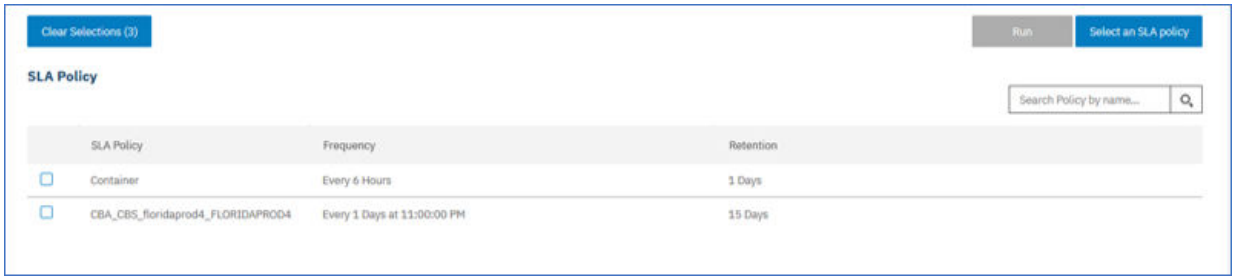
   Tip: Toggle between **Project** or **Label** views by using the options from the **View** menu.



3. Select the namespaces that you want to include in the application-consistent backup operation. Use one of the following methods:
   - Expand the cluster by clicking the cluster name to show the namespaces in that cluster. Select the namespaces to be backed up.
   - If resources and PVCs are labeled uniquely, you can use the labels to select the applications to back up. Expand the cluster by clicking the cluster name to show the labels in that cluster. Select the labeled resources to be backed up. Labels are indicated by the label icon .

   For more information about the resource icons, see "Backing up and restoring OpenShift clusters" on page 430.
4. Click **Select an SLA Policy** to open the **SLA Policy** pane.

   Select the default container SLA policy or any custom SLA policy that is listed in the table. For each SLA policy in the table, the backup frequency and retention period are shown.



5. Click **Save** to save the SLA policy.

   To run the job immediately, scroll to the **SLA Policy Status** pane. Click **Actions** > **Start** for the policy that you want to run.

   If you assign an SLA policy at the label or namespace level, any new PVCs that you create with the label or in the namespace are automatically assigned to the SLA policy.

What to do next

You can configure extra options for the SLA policy. For instructions, see "Specifying SLA options for OpenShift backup jobs" on page 442.

Deleting a namespace: To remove a namespace from an SLA backup job, complete the following steps:

1. In the **OpenShift Backup** pane, expand the cluster that contains the namespace resource that you want to remove from the policy.
2. Select the resources that you want to remove from the policy. Click **Select an SLA Policy**.
3. Clear the checkbox next to the namespace that you want to remove from the policy.
4. Click **Save**.

Specifying SLA options for OpenShift backup jobs

After you select a service level agreement (SLA) for your backup job, you can configure more options for that SLA. Additional SLA options include running scripts, excluding resources from the backup operation, and forcing a full base backup copy if required.

About this task

You can set up a regular backup job by creating an SLA policy. Container SLA policy definitions can back up any combination of persistent volumes, cluster-scoped and namespace-scoped resources. Copies can be stored locally on the cluster or backed up to the vSnap server.

Procedure

1. In the navigation panel, click **Manage Protection** > **Containers** > **OpenShift**.
2. In the **Policy Options** column of the **SLA Policy Status** table, click the clipboard icon  for an SLA policy and set the following options:

   Pre-Script
   Select this checkbox to run a script before a job runs. Windows-based machines support batch and PowerShell scripts while Linux-based machines support shell scripts. Take one of the following actions:

   - To use a script server, select **Use Script Server** and choose an uploaded script from the **Script** or **Script Server** list.
   - To run a script on an application server, clear the **Use Script Server** checkbox, and choose an application server from the **Application Server** list.

   Scripts and script servers are configured by using the **System Configuration** > **Script** page.

   Post-Script
   Select this checkbox to run a script after a job runs. Windows-based machines support batch and PowerShell scripts while Linux-based machines support shell scripts. Take one of the following actions:

   - To use a script server, select **Use Script Server** and choose an uploaded script from the **Script** or **Script Server** list.
   - To run a script on an application server, clear the **Use Script Server** checkbox, and choose an application server from the **Application Server** list.

   Scripts and script servers are configured by using the **System Configuration** > **Script** page.

   Continue job/task on script error
   Select this checkbox to continue running the job when the script that is associated with the job fails.

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

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

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

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

   Separate multiple filters with a semicolon.
3. Click **Save**.

Restoring OpenShift container data

You can use the IBM Spectrum Protect Plus user interface to restore container data such as persistent volumes, project-scoped resources, and cluster-scoped resources. You can also restore data from application-consistent container backups.

Preparing to restore persistent volumes and resources

Before you restore persistent volume claims (PVCs) and resources to a different namespace, you must prepare your OpenShift environment.

About this task

The security features of OpenShift Container Platform enforce the separation of namespaces (projects) more strictly than in Kubernetes clusters. If an application, consisting of volumes and resources, is restored to a different namespace, the target namespace must have the same settings as the source namespace. These settings depend on the application.

The procedure that you follow depends on whether an application uses Security Context Constraints (SCCs). Also, depending on the application that is being restored, the actions might vary slightly.

Procedure

- If an application to be restored does not use SCCs, you must configure the target namespace to use the same user ID (UID) range as the source namespace.
  a) Determine the group and UID range of the namespace in which the application runs by issuing the following commands:

     ```
     $ oc describe ns source-ns | grep scc
     ```

     where *source-ns* is the name of the source namespace.

     The output is similar to the following example:

     ```
     openshift.io/sa.scc.mcs: s0:c26,c15
     openshift.io/sa.scc.supplemental-groups: 1000670000/10000
     openshift.io/sa.scc.uid-range: 1000670000/10000
     ```

     In the previous example, the group range is shown in the `openshift.io/sa.scc.supplemental-groups` field and the UID range is shown in the `openshift.io/sa.scc.uid-range` field.
  b) Adjust the values of the group and UID range in the target namespace.

     Because the adjustments affect all applications that run in the target namespace, ensure that the target namespace is empty. Then, issue the following command:

     ```
     $ oc edit ns target-ns
     ```

     where *target-ns* is the name of the target namespace.

     In the editor, adjust the values for **`supplemental-groups`** and **`uid-range`** to match the values in the source namespace.
- If the application to be restored uses SCCs, ensure that the SCCs for both source and target namespaces are identical. The following example illustrates how to set the SCC settings for Db2 on Db2U (Db2 Universal Container):
  a) Determine the user account that is used by the stateful set in the source namespace:

     ```
     $ oc describe -n db2u statefulset.apps/c-db2ucluster1-db2u |grep -i account
     ```

     The output is similar to the following example:

```
Service Account: account-db2u-db2ucluster1
```

b) Determine the SCC that is used by the associated pod in the source namespace:

```
$ oc describe -n db2u pod/c-db2ucluster1-db2u-0 |grep -i scc
```

The output is similar to the following example:

```
openshift.io/scc: c-db2ucluster1-scc
```

c) Switch to the target namespace (`db2-clone`) by issuing the following command:

```
$ oc project db2-clone
```

d) Add the same SCC (`c-db2ucluster1-scc`) to the user account (`account-db2u-db2ucluster1`) in the target namespace by issuing the following command:

```
$ oc adm policy add-scc-to-user c-db2ucluster1-scc -z account-db2u-db2ucluster1
```

Preparing to restore persistent volumes that have VMs on OpenShift clusters

In an OpenShift cluster that contains virtual machines (VMs) that run as part of the OpenShift Virtualization feature, you can restore VMs that were backed up during a persistent volume claim (PVC) backup job. When a PVC is restored, any VMs that are within the PVC are restored as well. Before you can restore the PVCs, you must prepare the VMs.

For data to be consistent, ensure that all VMs that are running within a PVC are powered off before you begin a restore job.

Depending on the target for the restored PVC, you must complete specific actions before and after a restore job runs.

Table 91. Actions to take before and after a restore job

Restore target	**Actions**
Restoring a PVC to the same namespace and same PVC name	Before you start the restore job, rename or delete the existing VM. After the restore job is completed, use the restored PVC as the boot source to re-create the VM with original VM name.
Restoring a PVC to the same namespace and a different PVC name	After the restore job is completed, create the VM with a new VM name by using the restored PVC as the boot source.
Restoring a PVC to an alternative namespace and same PVC name	After the restore job is completed, use the restored PVC as the boot source to re-create the VM with the original VM name.
Restoring a PVC to an alternative namespace and different PVC name	After the restore job is completed, create the VM with a new VM name by using the restored PVC as the boot source.
Restoring PVC to an alternative namespace and creating the restored VM in a different namespace	See Enabling user permissions to clone data volumes across namespaces.

After you prepared the VMs, begin a restore job. For instructions, see one of the following topics:

- “Restoring persistent volumes in an OpenShift cluster” on page 445
- “Restoring OpenShift cluster-scoped and namespace-scoped resources” on page 449
- “Restoring application-consistent data on OpenShift clusters” on page 454

After the restore job is completed, re-create the VM by using the restore PVC as the boot source. For more information, see Table 91 on page 444.

Restoring persistent volumes in an OpenShift cluster

You can use the IBM Spectrum Protect Plus user interface to restore persistence volume claims (PVCs) from a snapshot or a copy backup. A snapshot restore operation is generally the faster method for restoring a PVC or resources.

Before you begin

Review the following restrictions:

- You cannot restore a snapshot backup to a different cluster or namespace.
- To help prevent issues with restore operations, do not manually delete snapshots of volumes that are protected by Container Backup Support.
- A PVC can be restored only in production mode.

For PVCs that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the VMs are restored as part of a PVC restore job. Before you begin the restore job, ensure that you prepare the VMs. For instructions, see "Preparing to restore persistent volumes that have VMs on OpenShift clusters" on page 444.

About this task

To create a restore job, use the **Restore** wizard. You can create on-demand jobs that run once or you can create recurring restore jobs that run on a schedule.

The following table summarizes the destinations that are supported for snapshot and copy restore operations.

Table 92. Restore scenarios

	Destination: original cluster		**Destination: alternative cluster**
Restore by specifying	**Snapshot restore**	**Copy restore**	**Copy restore**[2]
Original PVC and same project	Supported	Supported	Supported
Original PVC and another existing project	Does not apply	Supported	Supported
Original PVC and new project	Does not apply	Not supported	Not supported
New PVC and same project	Supported	Supported	Supported
New PVC and another existing project	Does not apply	Supported	Supported
New PVC and new project	Does not apply	Not supported	Not supported

You can also restore snapshot backups and copy backups to another storage class of the same storage provisioner. For example, if a PVC is defined to a storage class whose provisioner is `csi-rbd`, you can restore the snapshot backup of the PVC to another storage class whose provisioner is also `csi-rbd`.

[2] Snapshot restore jobs are not available when an alternative cluster is specified as the target.

Procedure

To restore PVCs from snapshot or copy backups, define a restore job by completing the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **OpenShift**.
2. Click **Create job** to open the **Create job** page.
3. In the **Restore** pane, click **Select** to open the **Restore** wizard.

   You can also open the **Restore** wizard by clicking **Jobs and Operations** > **Create job**. Then, click **Select** in the **Restore** pane, and click **OpenShift**.

   Tips:

   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, switch **Default setup** to **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
4. On the **Select source** page, browse the table and select the PVC that you want to restore by clicking the plus icon next to the PVC.

   The selected PVCs are displayed in the **Item** list. To remove an item from the list, click the minus icon next to the item.

   Alternatively, you can search for a PVC by specifying all or part of the PVC name in the **Search for** field and click the search icon.

   You can also toggle between the namespace and label views by clicking the options in the **View** menu.
5. On the **Source snapshot** page, use one of the following methods to select the source that you want to restore from:

Table 93. Methods for restoring a PVC

Source type	Steps
From a snapshot	a. Click **Origin** > **From Snapshot**. b. In the **Type of Restore** field, select the type of restore job that you want to create: **On-Demand** Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard. **Recurring** Creates a repeating restore job that runs on a schedule. This option is useful if you want to test a restore operation regularly. c. For on-demand restore only: Click the date range field and specify a range of dates. The available snapshot backups within that date range are shown. d. For on-demand restore only: If you are restoring a single PVC, select a snapshot from the list of available snapshots. If you are restoring more than one item, select a restore point for each item that you selected. e. Click **Next** to continue.

Table 93. Methods for restoring a PVC (continued)

Source type	Steps
From a copy backup	a. Click **Origin** > **From Copy**. b. In the **Type of Restore** field, select the type of restore job that you want to create: **On-Demand** Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard. **Recurring** Creates a repeating restore job that runs on a schedule. This option is useful if you want to test a restore operation regularly. c. In the **Restore Location Type** menu, select a type of location to restore data from: **Site** The site to which data was backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which data was copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which data was copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud service archive to which data was copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which data was copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. d. In the **Select a Location** menu, take one of the following actions: • If you are restoring data from a site, select one of the following restore locations: **Demo** The demonstration site from which to restore copy backups. This menu item is available only if you upgraded the product from IBM Spectrum Protect Plus Version 10.1.6 or earlier. **Primary** The primary site from which to restore copy backups. **Secondary** The secondary site from which to restore copy backups. • If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu. e. For on-demand restore only: Click the date range field and specify a range of dates. The available copy backups within that date range are shown. f. For on-demand restore only: If you are restoring a single PVC, select a backup from the list of available items. If you are restoring more than one PVC, select a restore point for each PVC that is listed. g. Click **Next** to continue.

6. On the **Restore method** page, enter a name for the restored PVC. The PVC name can be the original PVC name or a new PVC name.

To designate the new PVC, you can enter up to 221 characters for the PVC name and a prefix of 32 characters. You can include alphanumeric characters, periods (.), and hyphens (-). The new PVC name must not contain uppercase letters and must not end with a hyphen or a period. For example, `restored-pvc1` is a valid PVC name, where `restored-` is the prefix, and `pvc1` is the original PVC name.

Requirement: If you are restoring a PVC from a snapshot to the same PVC and project, ensure that you delete the original PVC before the restore job runs because the restore operation does not overwrite existing PVCs. Before you delete the original PVC, take a snapshot of the PVC to ensure that the latest changes are saved.

Click **Next** to continue.

7. On the **Set destination** page, configure the destination for the restored item.
   - For restore operations from snapshot backups, select from the following options:

     Original storage class
     Click this option to restore a snapshot backup of a PVC to the original storage class.

     Choose from available storage class
     Click this option to restore a snapshot backup of a PVC to a different storage class. Then, from the list of available storage classes, select a storage class of the same provisioner as the PVC.

     Original namespace
     Click this option to restore a snapshot backup to the original project.
   - For restore operations from copy backups, select from the following options:

     Restore to original cluster
     Click this option to restore a copy backup to the original cluster.

     Restore to alternate cluster
     Click this option to restore a copy backup to a different cluster. Then, select a cluster from the list of available clusters.

     Original storage class
     Click this option to restore a copy backup of a PVC to the original storage class.

     Choose from available storage class
     Click this option to restore a copy backup of a PVC to a different storage class. Then, from the list of available storage classes, select a storage class of the same provisioner as the PVC.

     Original namespace
     Click this option to restore a copy backup to the original project.

     Choose from available namespaces
     Click this option to restore a copy backup to a different project. Then, select a project from the list of available projects.
8. Optional: On the **Job options** page, configure additional options for the restore job:

   Run cleanup immediately on job failure
   This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

   Allow session overwrite
   Enable this option to allow a scheduled session of a recovery job to force a pending session to clean up associated resources so that the new session can run.

   Continue with restores of other selected PVCs even if one fails
   If one PVC is not successfully restored, the restore job continues for all other PVCs that are being restored. If this option is not enabled, the restore job stops when the recovery of a PVC fails.

   Click **Next** to continue.
9. Optional: If you are running the wizard in advanced setup mode, on the **Apply scripts** page, specify scripts to run before or after an operation runs at the job level. Batch and PowerShell scripts are supported.

Pre-Script
Select this checkbox to choose an uploaded script and an application or script server where the pre-script will run. To select an application server where the pre-script will run, clear the **Use Script Server** checkbox. Scripts and script servers are configured on the **System Configuration** > **Script** page.

Post-Script
Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server where the post-script will run, clear the **Use Script Server** checkbox. Scripts and script servers are configured on the **System Configuration** > **Script** page.

Continue job/task on script error
Select this checkbox to continue running the job when the script that is associated with the job fails.

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

If you clear this checkbox, the restore operation is not attempted, and the pre-script or post-script task status is reported as FAILED.

10. If you specified a recurring job in Step "5" on page 446, on the **Schedule** page, enter a name for the job, specify the frequency and start time for the job, and click **Next** to continue.

    The **Schedule** page is not displayed for on-demand restore jobs.
11. On the **Review** page, review your restore job settings and click **Submit** to create the job.

Results

For on-demand jobs, the job begins after you click **Submit**, and the **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view the progress of the restore operation, expand the job. You can also download the log file by clicking **Download .zip**.

For recurring jobs, the job begins at the scheduled start time when you start the schedule on the **Jobs and Operations** > **Schedule** page.

You can view the status of running jobs on the **Jobs and Operations** > **Running Jobs** page.

What to do next

To verify whether the PVC is restored, issue the following **oc** command:

```
oc get pvc restored_pvc –n namespace
```

where *restored_pvc* specifies the name of the restored PVC, and *namespace* specifies the project of the restored PVC.

For restored PVCs that contain VMs, you must re-create the VMs by using the restore PVC as the boot source. For more information, see "Preparing to restore persistent volumes that have VMs on OpenShift clusters" on page 444.

Restoring OpenShift cluster-scoped and namespace-scoped resources

You can restore OpenShift cluster-scoped and namespace-scoped resources from a snapshot or a copy backup. A snapshot restore operation is generally the faster method for restoring these resources. You can also choose to restore a copy backup to another namespace or cluster. You can also restore snapshot backups and copy backups to another storage class of the same storage provisioner.

Before you begin

Review the following restrictions:

- You cannot restore a snapshot backup to a different cluster or namespace.
- To help prevent issues with restore operations, do not manually delete snapshots of volumes that are protected by Container Backup Support.

For OpenShift clusters that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the custom resource data for the VMs is restored along with the resources. Before you begin the restore job, ensure that you prepare the VMs. For instructions, see "Preparing to restore persistent volumes that have VMs on OpenShift clusters" on page 444.

About this task

To create the restore job, use the **Restore** wizard. You can create on-demand jobs that run one time or you can create recurring restore jobs that run on a schedule.

The following table summarizes the destinations that are supported for snapshot and copy restore operations.

Table 94. Restore scenarios

	Destination: Original cluster		**Destination: Alternate cluster**
Restore by specifying	**Snapshot restore**	**Copy restore**	**Copy restores only**
Original PVC and same namespace	Supported	Supported	Supported
Original PVC and another existing namespace	Does not apply	Supported	Supported
Original PVC and new namespace	Does not apply	Not Supported	Not Supported
New PVC and same namespace	Supported	Supported	Supported
New PVC and another existing namespace	Does not apply	Supported	Supported
New PVC and new namespace	Does not apply	Not Supported	Not Supported

Procedure

To restore your cluster-scoped and namespace-scoped resources from snapshots or copy backups, define a restore job by completing the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **OpenShift**.
2. Click **Create job** to open the **Create job** page.
3. In the **Restore** pane, click **Select** to open the **Restore** wizard.

   Tips:

   - For a summary of your wizard selections, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, set the mode to **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
4. On the **Select source** page, browse the table and expand the cluster to show the resources that are available for the restore operation. Click the plus icon  next to the resources you want to add to the job.

The resources you select are displayed in the **Item** list. If you need to remove an item from the list, click the minus icon next to the item.

Alternatively, you can search for a resource by specifying all or part of the resource name in the **Search for** field.

5. On the **Source snapshot** page, use one of the following methods to select the source that you want to restore from:

Table 95. Methods for restoring a PVC or namespace resource

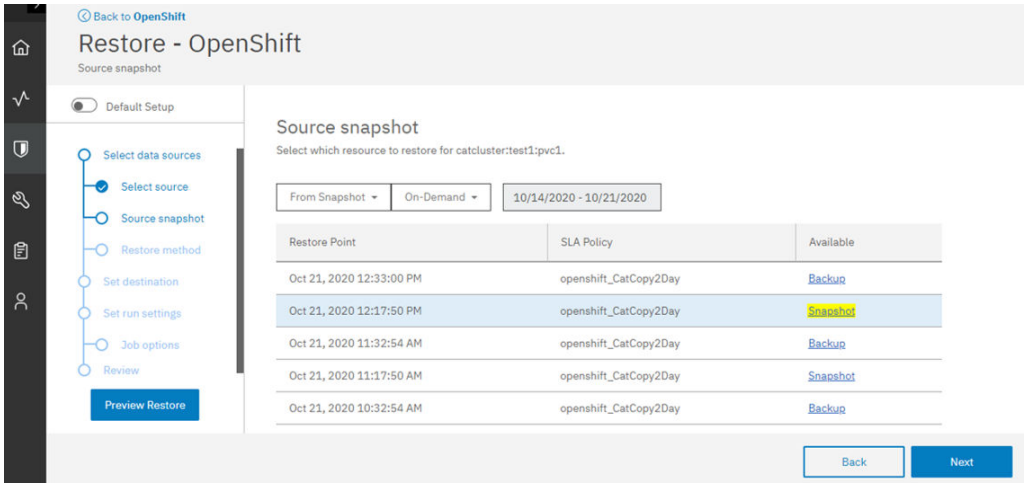
Source type	Steps
From a snapshot	a. Choose **From Snapshot** or **From Copy**. b. Click **Type of Restore**, and select the type of restore job that you want to create: **On-Demand** Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard. **Recurring** Creates a repeating restore job that runs on a schedule. This option is useful if you want to test a restore operation regularly. When you select the backup for the restore job, summary information shows the timestamp for that backup as shown.  c. For on-demand restore jobs, you can specify the date range to show the available snapshots backups within that date range. d. For on-demand restore jobs when you are restoring a single resource, you can select a snapshot from the list of available items. If you are restoring more than one item, you can select a restore point for each item that you selected. e. Click **Next** to continue.

Table 95. Methods for restoring a PVC or namespace resource (continued)

Source type	Steps
From a copy backup	a. Click **Origin** > **From Copy**. b. In the **Type of Restore** field, select the type of restore job that you want to create: **On-Demand** Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard. **Recurring** Creates a repeating restore job that runs on a schedule. This option is useful if you want to test a restore operation regularly. c. In the **Restore Location Type** menu, select a type of location to restore data from: **Site** The site where data was backed up to. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service where data was copied to. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage**. **Repository server** The repository server where data was copied to. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server**. **Cloud service archive** The cloud archive service where data was copied to. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server where data was copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. d. In the **Select a Location** menu, take one of the following actions: • If you are restoring data from a site, select one of the following restore locations: **Demo** The demonstration site from which to restore copy backups. This menu item is available only if you previously backed up data to a demo site in IBM Spectrum Protect Plus Version 10.1.6 or earlier. **Primary** The primary site from which to restore copy backups. **Secondary** The secondary site from which to restore copy backups. • If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu. e. For on-demand restore only: Click the date range field and specify a range of dates to show the available copy backups within that date range. f. For on-demand restore only: If you are restoring a single resource, select a backup from the list of available items. If you are restoring more than one PVC for example, select a restore point for each PVC that is listed. g. Click **Next** to continue.

6. Skip the **Restore method** page by clicking **Next**.
   You cannot add a new PVC name for restoring cluster-scoped and namespace-scoped resources.
7. On the **Set destination** page, configure the destination for the restored item.

- Snapshot backups can only be restored to their original storage classes and namespaces. For restore operations from Snapshot backups, select the following options:

  Storage class
  Choose to restore the resource from the snapshot to the `Original Storage Class`.

  Original namespace
  Choose to restore the resource from the snapshot to the `Original Namespace`.
- Copy backups can be restored to their original cluster, storage classes, or namespace. They can also be restored to a different cluster, storage class or namespace. For restore operations from copy backups, select from the following options:

  Restore to original cluster
  Select this option to restore a copy backup to the original cluster.

  Restore to alternate cluster
  Select this option to restore a copy backup to a different cluster which you can pick from a list of available clusters.

  Original storage class
  Select this option to restore the copy backup to the original storage class.

  Choose from available storage class
  Select this option to restore the copy backup to a different storage class, which you can pick from the list of available storage classes.

  Restriction: If you are restoring multiple resources and the associated storage classes are from different storage provisioners, you cannot select a different storage class.

  Original namespace
  Select this option to restore the copy backup to the original namespace.

  Choose from available namespaces
  Select this option to restore the copy backup to a different namespace which you can pick from the list of available namespaces.

8. On the **Job options** page, configure more options for the restore job:

   Run cleanup immediately on job failure
   This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

   Allow session overwrite
   Enable this option to allow a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run.

   Continue with restores of other selected resources even if one fails
   If one resource is not successfully restored, the restore job continues for all other resources that are being restored. If this option is not enabled, the restore job stops when the recovery of a resource fails.

   Click **Next** to continue.
9. Optional: If you are running the wizard in advanced setup mode, on the **Apply scripts** page, specify scripts to run before or after an operation runs at the job level. Batch and PowerShell scripts are supported.

   Pre-Script
   Select this checkbox to choose an uploaded script and an application or script server where the pre-script will run. To select an application server where the pre-script will run, clear the **Use Script Server** checkbox. Scripts and script servers are configured on the **System Configuration** > **Script** page.

   Post-Script
   Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server where the post-script will run, clear the **Use**

Script Server checkbox. Scripts and script servers are configured on the **System Configuration** > **Script** page.

Continue job/task on script error
Select this checkbox to continue running the job when the script that is associated with the job fails.

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

If you clear this check box, the restore operation is not attempted, and the pre-script or post-script task status is reported as FAILED.

10. If you specified a recurring job on the **Schedule** page, enter a name for the job, specify the frequency and start time for the job, and click **Next** to continue.

    The **Schedule** page is not displayed for on-demand restore jobs.
11. On the **Review** page, review your restore job settings and click **Submit** to create the job.

Results

For on-demand jobs, a job begins after you click **Submit**, and the **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view progress of the restore operation, expand the job. You can also download the log file by clicking **Download .zip**.

For recurring jobs, a job begins at the scheduled start time when you start the schedule on the **Jobs and Operations** > **Schedule** page.

All running jobs are viewable on the **Jobs and Operations** > **Running Jobs** page.

What to do next

To verify the contents of the restore job, you can check the Velero logs by issuing the following **kubectl** command: v

```
velero restore logs <restorename> -n spp-velero --insecure-skip-tls-verify
```

where *restorename* specifies the name of the restored resource.

For restored PVCs that contain VMs, you must re-create the VMs by using the restore PVC as the boot source. For more information, see “Preparing to restore persistent volumes that have VMs on OpenShift clusters” on page 444.

Restoring application-consistent data on OpenShift clusters

In an OpenShift environment, you can restore application-consistent data, including persistent volume claims (PVCs) and namespace-scoped resources, from a snapshot or copy backup.

The data that you are restoring must come from an application-consistent backup in which a pre-backup hook was run to quiesce an application before the snapshot was created, and a post-backup hook was run to resume the application after the snapshot was completed. You can also restore an application-consistent copy backup to another namespace.

You can restore data from an application-consistent backup to the following targets:

- Restore PVCs and namespace-scoped data to the original namespace.
- Restore PVCs and namespace-scoped data to an alternative namespace.
- Restore only PVCs to the original or alternative namespace.

Before you begin

Review the following restrictions:

- You cannot restore a snapshot backup to a different cluster or namespace.
- To help prevent issues with restore operations, do not manually delete snapshots of volumes that are protected by Container Backup Support.

Take the following actions:

- If you are restoring PVCs and resources to a different namespace, prepare the OpenShift environment. For instructions, see "Preparing to restore persistent volumes and resources" on page 443.
- To restore the data of an intact application to an earlier point in time, it is sufficient to restore only the persistent volumes. You do not have to restore the resources. Ensure that all activities on the volumes are stopped, and shut down the pods that access the volumes. You can shut down the pods by scaling down the associated deployments. For example, for the MongoDB application, issue the following command to shut down the pods:

  ```
  oc -n mongodb scale --replicas=0 deployment.apps/mongodb
  ```

  After the pods are shut down, you can restore the associated PVCs.
- To restore a PVC with its original PVC name, you must remove the existing PVC before beginning the restore job. The PVC restore operation does not overwrite existing PVCs. For example, issue the following command to remove a PVC named `mongo_data` from the namespace `mongodb`:

  ```
  oc -n mongodb delete pvc mongo_data
  ```

- For OpenShift clusters that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the custom resource data for the VMs is restored along with the resources. Before you begin the restore job, ensure that you prepare the VMs. For instructions, see "Preparing to restore persistent volumes that have VMs on OpenShift clusters" on page 444.

About this task

To create the restore job, use the **Restore** wizard. You can create on-demand jobs that run once or recurring restore jobs that run on a schedule.

The following table summarizes the destinations that are supported for snapshot and copy restore operations.

Table 96. Restore scenarios

	Destination: Original cluster		**Destination: Alternative cluster**
Restore by specifying	**Snapshot restore job**	**Copy restore job**	**Only copy restore jobs**
Original PVC and same namespace	Supported	Supported	Supported
Original PVC and another existing namespace	Does not apply	Supported	Supported
Original PVC and new namespace	Does not apply	Not supported	Not supported
New PVC and same namespace	Supported	Supported	Supported
New PVC and another existing namespace	Does not apply	Supported	Supported
New PVC and new namespace	Does not apply	Not supported	Not supported

Procedure

To restore application-consistent data from snapshot or copy backups, define a restore job by completing the following steps:

1. In the navigation panel, click **Manage Protection** > **Containers** > **OpenShift**.
2. Click **Create job** to open the **Create job** page.
3. In the **Restore** pane, click **Select** to open the Restore wizard.

   Tips:

   - For a summary of your wizard selections, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, switch **Default setup** to **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
4. On the **Select source** page, take the following actions:
   a) Browse the table and expand the cluster to show the application projects (namespaces) that are available for the restore job.

      Alternatively, you can search for a resource or a PVC by specifying all or part of the resource name in the **Search for** field .
   b) Click the project that you want to restore.
   c) If you want to include the resources, click the plus icon next to both the resources and the PVCs.

      The resources and PVCs that you select are displayed in the **Item** list.
   d) To remove an item from the list, click the minus icon next to the item.
5. On the **Source snapshot** page, use one of the following methods to select the source from which to restore the PVC or namespace resource:

Table 97. Methods for selecting a source

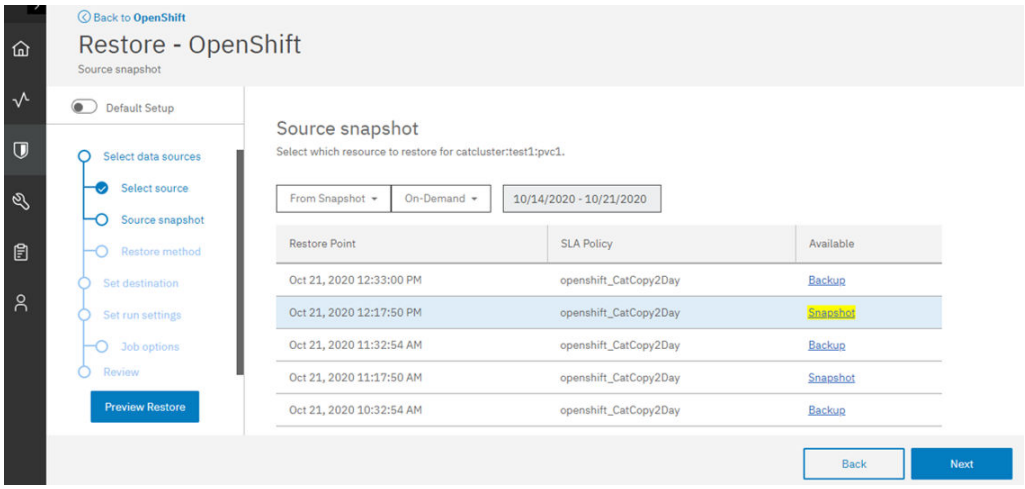
Source type	Steps
From a snapshot	a. Click **Origin** > **From Snashot**. b. Click **Type of Restore**, and select the type of restore job that you want to create: **On-Demand** Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard. **Recurring** Creates a repeating restore job that runs on a schedule. This option is useful if you want to test a restore operation regularly. When you select the backup for the restore job, the summary information includes the timestamp of the backup.  c. For on-demand restore jobs, you can specify the date range to show the available snapshots backups within that date range. d. For on-demand restore jobs when you are restoring a single resource, you can select a snapshot from the list of available items. If you are restoring more than one item, you can select a restore point for each item that you selected. e. Click **Next** to continue.

Table 97. Methods for selecting a source (continued)

Source type	Steps
From a copy backup	a. Click **Origin** > **From Copy**. b. In the **Type of Restore** field, select the type of restore job: **On-Demand** Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard. **Recurring** Creates a repeating restore job that runs on a schedule. This option is useful if you want to test a restore operation regularly. c. In the **Restore Location Type** menu, select a type of location to restore data from: **Site** The site to which data was backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which data was copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which data was copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which data was copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which data was copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. d. In the **Select a Location** menu, take one of the following actions: • If you are restoring data from a site, select one of the following restore locations: **Demo** The demonstration site from which to restore copy backups. This menu item is available only if you previously backed up data to a demo site in IBM Spectrum Protect Plus Version 10.1.6 or earlier. **Primary** The primary site from which to restore copy backups. **Secondary** The secondary site from which to restore copy backups. • If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu. e. For on-demand restore jobs: Click the date range field and specify a range of dates to show the available copy backups within that date range. f. For on-demand restore jobs: For example, if you are restoring a single resource, select a backup from the list of available items. If you are restoring more than one PVC, select a restore point for each PVC that is listed. g. Click **Next** to continue.

6. Skip the **Restore method** page by clicking **Next**.
   You cannot add a new PVC name for restoring cluster-scoped and namespace-scoped resources.
7. On the **Set destination** page, configure the destination for the restored item.

- Snapshot backups can be restored only to their original storage classes and namespaces. For restore operations from snapshot backups, select one of the following options:

  Original Storage Class
  Select this option to restore the resource from the snapshot to the original storage class.

  Original namespace
  Select this option to restore the resource from the snapshot to the `Original Namespace`.

- Copy backups can be restored to their original cluster, storage classes, or namespace. They can also be restored to a different cluster, storage class, or namespace. For restore operations from copy backups, select from one or more of the following options:

  Restore to original cluster
  Select this option to restore a copy backup to the original cluster.

  Restore to alternate cluster
  Select this option to restore a copy backup to a different cluster, which you can select from a list of available clusters.

  Original storage class
  Select this option to restore the copy backup to the original storage class.

  Choose from available storage class
  Select this option to restore the copy backup to a different storage class, which you can select from the list of available storage classes.

  Restriction: If you are restoring multiple resources and the associated storage classes are from different storage provisioners, you cannot select a different storage class.

  Original namespace
  Select this option to restore the copy backup to the original namespace.

  Choose from available namespaces
  Select this option to restore the copy backup to a different namespace, which you can select from the list of available namespaces.

8. Optional: On the **Job options** page, configure more options for the restore job:

   Run cleanup immediately on job failure
   This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

   Allow session overwrite
   Enable this option to allow a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run.

   Continue with restores of other selected resources even if one fails
   If one resource is not successfully restored, the restore job continues for all other resources that are being restored. If this option is not enabled, the restore job stops when the recovery of a resource fails.

   Click **Next** to continue.

9. Optional: If you are running the wizard in advanced setup mode, on the **Apply scripts** page, specify scripts to run before or after an operation runs at the job level. Batch and PowerShell scripts are supported.

   Pre-Script
   Select this checkbox to specify an uploaded script and an application or script server where the pre-script will run. To select an application server where the pre-script will run, clear the **Use Script Server** checkbox. Scripts and script servers are configured on the **System Configuration** > **Script** page.

   Post-Script
   Select this option to specify an uploaded script and an application or script server where the post-script will run. To select an application server where the post-script will run, clear the **Use**

Script Server checkbox. Scripts and script servers are configured on the **System Configuration** > **Script** page.

Continue job/task on script error
Select this checkbox to continue running the job when the script that is associated with the job fails.

When you select this checkbox, if a pre-script or post-script completes processing with a nonzero return code, the backup or restore job is attempted and the pre-script or post-script task status is reported as COMPLETED.

If you clear this check box, the restore job is not attempted, and the pre-script or post-script task status is reported as FAILED.

10. If you specified a recurring job on the **Schedule** page, enter a name for the job, specify the frequency and start time for the job, and click **Next** to continue.

    The **Schedule** page is not displayed for on-demand restore jobs.
11. On the **Review** page, review the specified restore job settings and click **Submit** to create the job.

Results

For on-demand jobs, a job begins after you click **Submit**, and the **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view progress of the restore operation, expand the job. You can also download the log file by clicking **Download .zip**.

For recurring jobs, a job begins at the scheduled start time when you start the schedule on the **Jobs and Operations** > **Schedule** page.

All running jobs are viewable on the **Jobs and Operations** > **Running Jobs** page.

What to do next

To verify the contents of the restore job, you can check the Velero logs by issuing the following **kubectl** command:

```
velero restore logs restorename -n spp-velero --insecure-skip-tls-verify
```

where *restorename* specifies the name of the restored resource.

Depending on the type of application that was restored, you might have to run manual recovery commands to get the application data out of the suspended state that it entered when it was backed up. You can create a temporary pod that is used only to run these actions. The template for this pod can be derived from the application and the pod can directly run the recovery action that is required.

In a more complex recovery, the action can be a sleep command. The recovery commands can then be run manually in a shell that is created within the temporary pod. After the recovery is completed, you can start the application by scaling up the deployment to the previous number of running replicas. For example, for the MongoDB application, issue the following command:

```
oc -n mongodb scale --replicas=1 deployment.apps/mongodb
```

For restored PVCs that contain VMs, you must re-create the VMs by using the restore PVC as the boot source. For more information, see “Preparing to restore persistent volumes that have VMs on OpenShift clusters” on page 444.

Expiring OpenShift job sessions

You can expire an OpenShift backup job session to override the retention settings that were assigned when a snapshot or copy backup was created. When a job session is expired, the restore point (the snapshot or copy backup) will be removed during the next maintenance job.

About this task

Complete this task if you do not want to wait for a job session to automatically expire according to the retention setting of the assigned service level agreement policy.

Procedure

To expire an OpenShift job session, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **IBM Spectrum Protect Plus** > **Restore Point Retention**.
2. On the **Backup Sessions** tab, search for a job session or restore point. Alternatively, on the **Virtual Machines / Databases** tab, select **Applications**, and search for a catalog entry by entering the name.

   Names can be searched by entering partial text, by using the asterisk (*) as a wildcard character, or by using the question mark (?) for pattern matching. For more information about using the search function, see Appendix A, “Search guidelines,” on page 681.
3. Optional: If you are searching from the **Backup Sessions** tab, use filters to narrow your search for snapshots or copy backups. You can also specify the date range when the associated backup job started.
   a) In the **Type** field, select **Application**.
   b) In the **Subpolicy Type** field, select **Snapshot** to search for snapshot backups or select **Backup** to search for copy backups.
   c) If needed, click the **Backup Time Range** field and select a date range to search on.
4. Click the search icon .
5. In the search results, select the job session that you want to expire.
6. If you are on the **Backup Sessions** tab, from the **Actions** menu, select one of the following options:
   - To expire a single job session, click **Expire**.
   - To expire all unexpired job sessions for the selected job, click **Expire All Job Sessions**.

   If you are on the **Virtual Machines / Databases** tab, select **Containers** in drop-down menu and search for a catalog item by name. Then, click the deletion icon for the resource that you want to expire.
7. Follow the instructions in the confirmation window and click **OK**.

Related tasks

“Managing IBM Spectrum Protect Plus restore points” on page 616

You can use the **Restore Point Retention** pane to search for restore points in the IBM Spectrum Protect Plus catalog by backup job name, view their creation and expiration dates, and override the assigned retention.

Monitoring Container Backup Support jobs and running reports

As the backup administrator, you can use the IBM Spectrum Protect Plus user interface to monitor Container Backup Support jobs and create reports that show the backup history of containers.

Viewing job logs

You can use the **Jobs and Operations** window to monitor Container Backup Support jobs, review job history, and view scheduled jobs.

About this task

You can identify the jobs in the **Running Jobs** and **Job History** tabs as follows:

- Inventory jobs are identified by the `Application Server Inventory` label.
- Maintenance jobs are identified by the `Maintenance` label.
- Backup job names are identified by the `k8s_`*`sla_name`* label on Kubernetes clusters and `openshift_`*`sla_name`* on OpenShift clusters.
- The job type is shown in the Type field. For example, a snapshot backup job is identified by type `Type: Backup - Snapshot`. A copy backup is identified by type `Type: Backup`.
- Restore job names are identified by the `onDemandRestore_`*`timestamp`* label. The job type is `Type: Restore`.

Procedure

1. In the IBM Spectrum Protect Plus navigation panel, click **Jobs and Operations**.
2. Click the appropriate tab:
   - To show the inventory, backup, and restore jobs that are running, click **Running Jobs**.
   - To show the jobs that ran successfully, completed processing with warnings, or jobs that failed, click **Job History**. You can download a job log from the page by selecting the job and clicking **Download.zip**.

     The downloaded file has the following naming convention: `JobLog_`*`job_name_timestamp`*`.zip`
   - To view the status of scheduled jobs, click **Schedule**.
   - To take a shortcut to create an ad hoc backup job or a restore job without going to the **Kubernetes** page in the **Manage Protection** section, click **Create job**.

Related concepts

“Creating jobs and job schedules” on page 620
The method for creating jobs and job schedules depends on the job type.

Related tasks

“Viewing jobs” on page 622
View information about the jobs that are running in your environment, the job history, the active resources that are associated with restore jobs, and scheduled jobs.

Creating reports for persistent volumes

You can run reports that show the backup history and backup storage utilization of your persistent volumes. By viewing the reports and taking any necessary action, you can help to ensure that your container data is protected by predefined service level agreement (SLA) policies.

- “Creating a backup history report” on page 463

- “Creating a backup utilization report” on page 464

Creating a backup history report

You can run a report to show the backup history of your protected persistent volumes. By viewing the backup history, you can determine whether your backup jobs are running as planned.

Before you begin

If you plan to schedule a report to run at specific times, ensure that you configure an SMTP server for email notifications. For instructions, see “Adding an SMTP server ” on page 295.

About this task

For each persistent volume claim (PVC), the backup history shows information about the Container Storage Interface (CSI) snapshots that were created in the Kubernetes or OpenShift environment and the backups that were copied to the IBM Spectrum Protect Plus vSnap server. You can view information such as the date and time of the backup operation, the size of the backup, and the duration of the copy operation. From this data, you can verify whether your scheduled backups are running according to the service level agreement (SLA) policy that you set for the PVC.

Procedure

1. In the IBM Spectrum Protect Plus navigation panel, click **Reports and Logs** > **Reports**.
2. In the **Name (job title)** column , locate the **Container Persistent Volume Backup History** row and take one of the following actions:

Action	Steps
To run a report immediately	a. Click the Run Report icon . b. In the **Run Report** window, modify the parameters as needed and click **Run**.
To schedule a report with the default parameters	a. Click the Schedule Report with default parameters icon . b. In the **Schedule Report with default parameters** window, specify the frequency, start time, and a recipient's email address. c. Click **Schedule**.
To create a custom report	a. Click the Create Custom Report icon . The **Create Custom Report** window is displayed. b. In the **Parameters** tab, enter a name and description for the custom report, and modify the report parameters as needed. The name of the report must not contain any spaces. c. To schedule the report to run at specific times, click the **Schedule** tab, and select **Define Schedule**. d. Specify the frequency, start time, and a recipient's email address. e. Click **Save Report**. The custom report is saved to the **Custom Reports** tab of the **Reports** window.

Action	Steps
To run a custom report	a. Click the **Custom Reports** tab. b. Identify the report that you want to run and click the Run Custom Report icon . c. In the **Run Custom Report** window, click **Run**.

Results

If you ran the report immediately, the backup history report is displayed in the **Container Persistent Volume Backup History** window. To download the report, click **Download** and select a report format. To return to the **Reports** window, click **Back to Reports**.

If you defined a schedule for the report, the backup history report is run at the scheduled time and sent to the recipient that you specified.

The descriptions of the reported data are shown in the following table:

Table 98. Details of the backup history report

Column	Description
SLA Policy	The SLA policy that is used to protect a PVC.
Protection Time	The date and time when each backup job was completed.
Status	The status of each backup job. If a backup job failed, a possible reason is provided.
Snapshot Backup?	An indication of whether the backup instance is a snapshot backup. A checkmark is displayed in the column to indicate that the instance is a snapshot backup. When a checkmark is displayed, no data is shown in the **Backup Size** and **Backup Speed** columns.
Backup Size	For copy backups, the amount of data that was backed up to the vSnap server. For snapshot backups that were created in the Kubernetes or OpenShift environment, or for backups that failed, no size is shown.
Backup Speed	The rate at which a copy backup was completed. For snapshot backups or backups that failed, no data is shown.

Creating a backup utilization report

You can run a report to show the utilization of your persistent volume backups on backup storage. Information such as where backups are stored, the size of each backup, and the number of restore points that are available are included.

Procedure

1. In the IBM Spectrum Protect Plus navigation panel, click **Reports and Logs** > **Reports**.
2. In the **Name (job title)** column , locate the **Container Persistent Volume Backup Utilization** row and take one of the following actions:

Action	Steps
To run a report immediately	a. Click the Run Report icon . b. In the **Run Report** window, modify the parameters as needed and click **Run**.

Action	Steps
To schedule a report with the default parameters	a. Click the Schedule Report with default parameters icon. b. In the **Schedule Report with default parameters** window, specify the frequency, start time, and a recipient's email address. c. Click **Schedule**.
To create a custom report	a. Click the Create Custom Report icon +. The **Create Custom Report** window is displayed. b. In the **Parameters** tab, enter a name and description for the custom report, and modify the report parameters as needed. The name of the report must not contain any spaces. c. To schedule the report to run at specific times, click the **Schedule** tab, and select **Define Schedule**. d. Specify the frequency, start time, and a recipient's email address. e. Click **Save Report**. The custom report is saved to the **Custom Reports** tab of the **Reports** window.
To run a custom report	a. Click the **Custom Reports** tab. b. Identify the report that you want to run and click the Run Custom Report icon. c. In the **Run Custom Report** window, click **Run**.

Results

The utilization report is shown in the **Container Persistent Volume Backup Utilization** section of the window. The descriptions of the reported data are shown in the following table:

Table 99. Details of the backup utilization report

Column	Description
PVC	The name of each persistent volume claim (PVC).
Namespace	The namespace in which the PVC exists.
SLA Policy	The service level agreement (SLA) policy that is used to protect each PVC.
Backup Storage	The location of the backup storage. The backup storage can be the hostname or IP address of a disk, the name of a cloud server, or the name of the repository server.
Backup Size	The size of each backup.
Recovery Points	The number of restore points that are available for each PVC.

Related concepts

“Managing reports and logs” on page 631

IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Protecting containers by using the command line

As an application developer in a Kubernetes or OpenShift environment, you can use the command-line interface to back up and restore container data, and to query the status of Container Backup Support requests.

Ensure that your environment meets the system requirements in "Container Backup Support requirements" on page 75.

Container Backup Support requests

To protect container data, you can submit Container Backup Support requests by using the Kubernetes or OpenShift command-line interface.

A Container Backup Support request is a Kubernetes custom resource that is of kind BaaSReq. The requests are specified in *YAML Ain't Markup Language* (YAML) configuration files. The request is then submitted by using the Kubernetes command-line tool (**kubectl**) or OpenShift command-line tool (**oc**).

Types of requests in Container Backup Support

The following table shows the available types of Container Backup Support requests. The request types are specified as values for the **requesttype** key in the YAML file. Links to instructions about creating and submitting the requests are also provided.

Table 100. Types of Container Backup Support requests for cluster resources

Request type	Description	Instructions
Backup	Schedule a backup operation for a persistent volume claim (PVC). This request is applicable for snapshot and copy backups.	"Scheduling backups of persistent volumes by using the command line" on page 468
BackupResources	Schedule a backup operation for namespace-scoped or cluster-scoped resources. This request is applicable for snapshot and copy backups.	"Scheduling resource backup jobs" on page 471
BackupLabel	Schedule a backup operation for a specific label. Resources that have been assigned that label will be added to the SLA policy whether that includes PVCs, cluster-scoped or namespace-scoped resource.	"Backing up persistent volumes and resources with labels using the command line" on page 475
BackupNamespace	Schedule a backup operation for a namespace. This namespace is added to the SLA policy whether that includes PVCs and namespace-scoped resources.	"Backing up persistent volumes and resources by namespace using the command line" on page 478
OnDemandBackup	Request an immediate snapshot backup of PVCs.	"Backing up a persistent volume on demand by using the command line" on page 472

Table 100. Types of Container Backup Support requests for cluster resources (continued)		
Request type	**Description**	**Instructions**
OnDemandBackupResources	Request an immediate snapshot backup of cluster-scoped or namespace-scoped resources.	"Backing up resources on demand" on page 474
Restore	Restore a PVC from a snapshot backup or a copy backup	"Restoring container data by using the command line" on page 481
Destroy	Delete all snapshot and copy backups and mark the scheduled job as **destroyed**	"Deleting container backups" on page 489
RestoreResources	Restore cluster-scoped or namespace-scoped resources from a snapshot backup or a copy backup.	"Restoring resources by using the command line" on page 484

Running a request

To initiate a request, create a YAML configuration file that specifies the request type and provides the required parameters. Then, submit the request by running the **kubectl create** command.

The following sample file (`baas-req.yaml`) shows the general format of a YAML file:

```
#------------------------
# Filename: baas-req.yaml
#------------------------

apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: request_name
  namespace: namespace
spec:
  requesttype: request_type
  scope: Cluster | Namespace
  sla: [sla_policy]
  volumesnapshotclass: snapshot_class_name
```

where:

request_name
Specifies the name of the request. For scheduled backup requests, the name of the request must match the PVC name.

namespace
Specifies the PCV or namespace in which the item to be protected exists. If you do not specify a namespace, the default namespace is used.

request_type
Specifies the type of request. For the list of available request types, see "Types of requests in Container Backup Support" on page 466.

scope
Specifies whether the resource for this request is a cluster-scoped or a namespace-scoped resource.

[*sla_policy*]
Specifies one or more service level agreement (SLA) policies that you assign to the request. For information about the specifications for the SLA policy, see "Scheduling backups of persistent volumes by using the command line" on page 468.

snapshot_class_name
Specifies the snapshot class for the volume. If you do not specify the snapshot class, the default snapshot class is used if the sidecar container `csi-snapshotter` in the default snapshot class matches the provisioner of the volume. Otherwise, the backup request is invalid.

snapshot_class_name applies to PVCs only.

To start the request that is specified in the `baas-req.yaml` sample file, issue the following command:

```
kubectl create -f baas-req.yaml
```

To check the status of a request, use one of the following methods:

- To list all Container Backup Support requests in all namespaces that you can access, issue the following command:

  ```
  kubectl get baasreq --all-namespaces
  ```

- To display the status of all Container Backup Support requests in a specified namespace, issue the following command:

  ```
  kubectl describe baasreq -n namespace
  ```

  where *namespace* is the namespace of the persistent volume.

- To display the status of a specific Container Backup Support request, issue the following command:

  ```
  kubectl describe baasreq request_name -n namespace
  ```

  where *request_name* is the name of the request, and *namespace* is the namespace of the persistent volume.

Backing up container data

To protect persistent volumes that are attached to a container, you can schedule backup operations to run as specified by predefined service level agreement (SLA) policies. You can also schedule backup operations to include namespace-scoped and cluster-scoped resources. To create snapshots of any combination of these resources immediately, you can run an on-demand backup request.

Scheduling backups of persistent volumes by using the command line

By using the Kubernetes or OpenShift command line, you can schedule backup requests based on service level agreement (SLA) policies. SLA policies specify how often backup operations are run and how long snapshot and copy backups are retained.

Before you begin

Backup requests are directed to persistent volume claims (PVCs) for the volumes that you want to protect. Before you schedule a backup job, take the following actions:

- Ensure that the PVC exists within the specified namespace.
- Ensure that the PVC is formatted. PVCs must be formatted before they can be backed up. For a PVC to be formatted correctly, it must be mounted and written to. Backup operations of raw block volumes are not supported.
- Determine which SLA policy to assign to PVCs. For instructions about viewing the available SLA policies, see “SLA policies” on page 405.
- For PVCs that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the VMs are protected during a PVC backup job. For the data to be consistent, ensure that you power off the VMs before running the job.

About this task

When a scheduled backup job runs, an inventory of the cluster resources is run automatically and a snapshot of the persistent volume is created at the frequency that is defined by the SLA. If the SLA specifies a copy backup policy, the snapshot of the volume is copied to an IBM Spectrum Protect Plus vSnap server.

All backup jobs are scheduled, except for on-demand backup jobs. To schedule backup jobs for a PVC, create a YAML configuration file with job specifications and apply the request on the command line in the Kubernetes or Red Hat OpenShift environment.

You can specify one or more SLA policies per PVC.

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

Procedure

1. Optional: Display a list of PVCs in your namespace by issuing the following command:

   ```
   kubectl get pvc -n namespace
   ```

   From the list of PVCs, identify the PVC that you want to back up.
2. Create a YAML file that defines the request for a scheduled backup. The YAML file must contain the following properties:

   ```
   #------------------------
   # Filename: filename.yaml
   #------------------------

   apiVersion: "baas.io/v1alpha1"
   kind: BaaSReq

   metadata:
     name: request_name
     namespace: namespace
   spec:
     requesttype: Backup
     sla: [sla_policy]
     volumesnapshotclass: snapshot_class_name
   ```

   where:

   filename
   Specifies the name of the YAML configuration file. The file type is `.yaml`.

   request_name
   Specifies the name of the backup request, which must match the name of the PVC for the volume that you want to back up. For example, to create a backup request for a PVC that is named `dbvol-01`, the name of the request must be `dbvol-01`.

   namespace
   Specifies the namespace in which the PVC exists.

   [*sla_policy*]

   Specifies the SLA policy that determines the schedule for backup operations. You can specify more than one SLA policy by using a comma-separated list within the brackets.

   For example, to assign the `daily` policy to a PVC, specify the following statement:

   ```
   sla: [daily]
   ```

   To assign the `every4hours`, `daily_midnight`, and `weekly` policies to the PVC, specify the following statement in the YAML file:

   ```
   sla: [every4hours,daily_midnight,weekly]
   ```

Alternatively, you can use the following format to specify a single SLA policy:

```
sla:
- daily
```

Or use the following format to specify multiple SLA policies:

```
sla:
- every4hours
- daily_midnight
- weekly
```

Ensure that you use the correct case when you specify the SLA policy name. Policy names are case-sensitive in YAML files.

To remove all SLA assignments from a PVC, delete the SLA policy names within the brackets, as shown in the following statement:

```
sla: []
```

Specifying the empty brackets is the only method that you can use to remove all SLA assignments from the PVC.

snapshot_class_name
Specifies the snapshot class for the volume. If you do not specify the snapshot class, the default snapshot class is used if the sidecar container `csi-snapshotter` in the default snapshot class matches the provisioner of the volume. Otherwise, the backup request is invalid.

3. Submit the backup request by issuing the following command:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

Results

After you submit the backup request, the first scheduled backup operation will start within the window that is defined by the SLA policy. The start time of the backup is recorded in the backup status.

What to do next

To view information about the backup operation, issue the **kubectl describe** command by using the request name or the PVC name. For instructions, see “Viewing the status of backup and restore jobs” on page 486.

Modifying parameters in a YAML file:

After scheduled backup jobs have started, you can modify the parameters in the YAML file and apply it to the same PVC if needed. For example:

- To assign a different SLA policy to the PVC or remove an SLA assignment, edit the values in the **sla** field in the YAML file. Then, apply the YAML file by using the **kubectl** command-line interface.
- If you no longer want the PVC to participate in any scheduled backup jobs, remove the SLA policy assignments by updating the **sla** field in the YAML file. To remove the PVC from all SLAs, modify the **sla** field as follows:

  ```
  sla: []
  ```

  Then, apply the YAML file by using the **kubectl** command-line interface.

Related concepts
“Backup and restore types” on page 403

Container Backup Support provides multiple types of backup and restore functions for your PVC and other cluster resources. You can use the IBM Spectrum Protect Plus user interface or the Kubernetes or OpenShift command line to initiate backup and restore operations.

"SLA policies" on page 405
Service level agreement (SLA) policies define how often snapshot backup and copy backup operations are run, and how long snapshots and copy backups are retained. You can set up custom SLAs that meet your operational requirements.

"Container Backup Support requests" on page 466
To protect container data, you can submit Container Backup Support requests by using the Kubernetes or OpenShift command-line interface.

"Troubleshooting Container Backup Support" on page 662
To help troubleshoot issues with Container Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Scheduling resource backup jobs

Use the **BackupResources** command to add namespace-scoped or cluster-scoped resources to a scheduled backup job in an existing SLA policy for Kubernetes or OpenShift.

Before you begin

For OpenShift clusters that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the custom resource data for the VMs are protected when resources are backed up. For the data to be consistent, ensure that you power off the VMs before running a backup job.

About this task

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

Procedure

1. Create a YAML file that defines the request for an on-demand backup operation. The YAML file must contain the following properties:

```
apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: request_name
  namespace: namespace
spec:
  requesttype: BackupResources
  scope: cluster | namespace
  sla: [sla_policy]
```

   where:

   name_of_request
   Specifies the name of the on-demand backup request. The name must be unique, and must not match the name of the namespace or cluster resource.

   A new on-demand backup for the cluster-scoped and namespace-scoped resources request must be created for each subsequent on-demand backup of the same resource. Create a new request and specify a different request name (*name_of_request*) in the YAML file for subsequent requests.

   namespace
   Specifies the namespace.

scope
Specifies whether the resource for this request is a cluster-scoped or a namespace-scoped resource.

[*sla_policy*]
Specifies the SLA policy that determines the schedule for backup operations. For example, to assign the `daily` policy to a PVC, specify the following statement:

```
sla: [daily]
```

Ensure that you use the correct case when you specify the SLA policy name. Policy names are case-sensitive in YAML files.

Any SLAs that are not in the corresponding scheduled backup request will be added to the list of SLAs in that request.

2. Start the backup operation for the selected resources by issuing the following command:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

Results

To view information about the backup, issue the **kubectl describe** command by using the request name or the cluster or namespace name. For instructions, see "Viewing the status of backup and restore jobs" on page 486.

Backing up a persistent volume on demand by using the command line

To create a snapshot immediately without waiting for a scheduled backup job to run, run an on-demand backup job in the Kubernetes or OpenShift command-line interface.

Before you begin

Backup requests are directed to persistent volume claims (PVCs) for the volumes that you want to protect. Before you schedule a backup job, take the following actions:

- Ensure that the PVC exists within the specified namespace.
- Ensure that the PVC is formatted. PVCs must be formatted before they can be backed up. For a PVC to be formatted correctly, it must be mounted and written to. Backup operations of raw block volumes are not supported.
- Determine which SLA policy to assign to PVCs. For instructions about viewing the available SLA policies, see "SLA policies" on page 405.
- For PVCs that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the VMs are protected during a PVC backup job. For the data to be consistent, ensure that you power off the VMs before running the job.

About this task

During an on-demand backup operation, only a snapshot is created. After the initial on-demand backup operation is completed, the volume will be protected according to the specified SLA policy.

Unlike a request for scheduled backups, the name of the on-demand request must be unique. In other words, the name of the request must not be the same as the name of the PVC.

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

Procedure

1. Optional: Display a list of PVCs in your namespace by issuing the following command:

```
kubectl get pvc -n namespace
```

From the list of PVCs, identify the PVC that you want to back up.

2. Create a YAML file that defines the request for an on-demand backup operation. The YAML file must contain the following properties:

```
#------------------------
# Filename: filename.yaml
#------------------------

apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: name_of_request
  namespace: namespace
spec:
  requesttype: OnDemandBackup
  pvcname: pvc_name
  sla: [sla_policy]
  volumesnapshotclass: snapshot_class_name
```

where:

filename
Specifies the name of the YAML configuration file. The file type is `.yaml`.

name_of_request
Specifies the name of the on-demand backup request. The name must be unique, and must not match the name of the PVC.

A new on-demand backup request must be created for each subsequent on-demand backup of the same PVC. In other words, to create a second on-demand backup of a PVC, create a new request and specify a different request name (*name_of_request*) in the YAML file.

namespace
Specifies the namespace in which the PVC exists.

pvc_name
Specifies the name of the PVC for the volume that you want to back up.

[*sla_policy*]
Specifies the SLA policy that determines the schedule for backup operations. For example, to assign the `daily` policy to a PVC, specify the following statement:

```
sla: [daily]
```

Ensure that you use the correct case when you specify the SLA policy name. Policy names are case-sensitive in YAML files.

Any SLAs that are not in the corresponding scheduled backup request for the PVC will be added to the list of SLAs in that request.

snapshot_class_name
Specifies the snapshot class for the volume. If you do not specify the snapshot class, the default snapshot class is used if the sidecar container `csi-snapshotter` in the default snapshot class matches the provisioner of the volume. Otherwise, the backup request is invalid.

3. Start the on-demand backup operation by issuing the following command:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

Results

To view information about the backup, issue the **`kubectl describe`** command by using the request name or the PVC name. For instructions, see “Viewing the status of backup and restore jobs” on page 486.

Related concepts

"Backup and restore types" on page 403
Container Backup Support provides multiple types of backup and restore functions for your PVC and other cluster resources. You can use the IBM Spectrum Protect Plus user interface or the Kubernetes or OpenShift command line to initiate backup and restore operations.

"Container Backup Support requests" on page 466
To protect container data, you can submit Container Backup Support requests by using the Kubernetes or OpenShift command-line interface.

"Troubleshooting Container Backup Support" on page 662
To help troubleshoot issues with Container Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Backing up resources on demand

You can use the **OnDemandBackupResources** command to create an immediate snapshot without waiting for a scheduled backup job to run, in the Kubernetes or OpenShift command-line interface.

Before you begin

For OpenShift clusters that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the custom resource data for the VMs are protected when resources are backed up. For the data to be consistent, ensure that you power off the VMs before running a backup job.

About this task

During an on-demand backup operation, only a snapshot is created. After the initial on-demand backup operation is completed, the namespace-scoped or cluster-scoped resource is protected according to the specified SLA policy.

Unlike a request for scheduled backups, the name of the on-demand request must be unique.

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

Procedure

1. Create a YAML file that defines the request for an on-demand backup operation. The YAML file must contain the following properties:

```
apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: request_name
  namespace: namespace
spec:
  requesttype: OnDemandBackupResources
  scope: cluster | namespace
  sla: [sla_policy]
```

   where:

   name_of_request
   Specifies the name of the on-demand backup request. The name must be unique, and must not match the name of the namespace or cluster resource.

   A new on-demand backup for the cluster-scoped and namespace-scoped resources request must be created for each subsequent on-demand backup of the same resource. Create a new request and specify a different request name (*name_of_request*) in the YAML file for subsequent requests.

   namespace
   Specifies the namespace for the BaaSReq and resources.

scope
Specifies whether the resource for this request is a cluster-scoped or a namespace-scoped resource.

[*sla_policy*]
Specifies the SLA policy that determines the schedule for backup operations. For example, to assign the `daily` policy to a resource, specify the following statement:

```
sla: [daily]
```

Ensure that you use the correct case when you specify the SLA policy name. Policy names are case-sensitive in YAML files.

Any SLAs that are not in the corresponding scheduled backup request will be added to the list of SLAs in that request.

2. Start the on-demand backup operation for the selected resources by issuing the following command:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

Results

To view information about the backup, issue the **kubectl describe** command by using the request name or the cluster or namespace name. For instructions, see "Viewing the status of backup and restore jobs" on page 486.

Backing up persistent volumes and resources with labels using the command line

You can create backup requests for persistent volumes and resources by specifying labels. Labels are key-value pairs that are attached to objects, such as pods or PVCs. By specifying one or more labels in a backup request, you can back up all PVCs or resources that are associated with those labels.

Before you begin

Backup requests are directed to persistent volume claims (PVCs) for the volumes that you want to protect. Before you schedule a backup job, take the following actions:

- Ensure that the PVC exists within the specified namespace.
- Ensure that the PVC is formatted. PVCs must be formatted before they can be backed up. For a PVC to be formatted correctly, it must be mounted and written to. Backup operations of raw block volumes are not supported.
- Determine which SLA policy to assign to PVCs. For instructions about viewing the available SLA policies, see "SLA policies" on page 405.
- For OpenShift clusters that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, ensure that you power off the VMs before running a backup job.

About this task

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

Procedure

1. Optional: Display a list of PVCs in a specified namespace by issuing the following command:

```
kubectl get pvc -n namespace --show-labels
```

From the list of PVCs, identify the label that is attached to the PVCs that you want to back up.

2. Create a YAML file that defines the request for the backup-by-label operation. The YAML file must contain the following properties:

```
#------------------------
# Filename: filename.yaml
#------------------------

apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: name_of_request
  namespace: namespace
spec:
  requesttype: BackupLabel
  sla: [sla_policy]
  volumesnapshotclass: snapshot_class_name
  backuplabels:
  - label_key: value
```

where:

filename
Specifies the name of the YAML configuration file. The file type is `.yaml`.

name_of_request
Specifies the name of the backup-by-label request. The name must be unique, and must not match the PVC name.

namespace
Specifies the namespace for the backup request.

[*sla_policy*]

Specifies the SLA policy that determines the schedule for backup operations. You can specify more than one SLA policy by using a comma-separated list within the brackets.

For example, to assign the `daily` policy to a PVC, specify the following statement:

```
sla: [daily]
```

To assign the `every4hours`, `daily_midnight`, and `weekly` policies to the PVC, specify the following statement in the YAML file:

```
sla: [every4hours,daily_midnight,weekly]
```

Alternatively, you can use the following format to specify a single SLA policy:

```
sla:
- daily
```

Or use the following format to specify multiple SLA policies:

```
sla:
- every4hours
- daily_midnight
- weekly
```

Ensure that you use the correct case when you specify the SLA policy name. Policy names are case-sensitive in YAML files.

To remove all SLA assignments from a label, delete the SLA policy names within the brackets, as shown in the following statement:

```
sla: []
```

snapshot_class_name
Specifies the snapshot class for the volume. If you do not specify the snapshot class, the default snapshot class is used if the sidecar container `csi-snapshotter` in the default snapshot class matches the provisioner of the volume. Otherwise, the backup request is invalid.

label_key: value
Specifies the key-value pair for the label that is attached to the PVCs that you want to back up. You can specify more than one label.

After you assign an SLA policy at the label level, any new PVCs that you create with that label will be automatically assigned to the SLA.

For example, to back up all PVCs that are associated with the `color: red` label and the `department: sales` label, specify the following statements:

```
backuplabels:
- color: red
- department: sales
```

Restrictions:

- PVC labels are key-value pairs. Any duplicate keys with different values are overwritten by the last key-value pair.
- The backup-by-label operation applies to all PVCs that have a specific label across the cluster. If any of the PVCs that have been backed up belong to a namespace that you do not have access to, you will not be able to restore those PVCs by using the command line. However, the PVCs can be restored by using the IBM Spectrum Protect Plus user interface, regardless of what namespace they belong to. For more information, see "Restoring Kubernetes container data" on page 420.

3. Submit the backup request by issuing the following command:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

Results

After you submit the backup request, the first scheduled backup operation will start within the window that is defined by the SLA policy. The start time of the backup is recorded in the backup status.

What to do next

To view information about the backup request, issue the **kubectl describe** command by using the request name. For example, to view information about a backup request that is named `backup-red-label` in the `baas` namespace, issue the following command:

```
kubectl describe baasreq backup-red-label -n baas
```

For instructions, see "Viewing the status of backup and restore jobs" on page 486.

Modifying parameters in a YAML file:

After scheduled backup-by-label jobs have started, you can modify the **sla** and the **backupLabels** parameters in the YAML file and apply the updates to the same request. For example:

- To assign a different SLA policy to the label or remove an SLA assignment, edit the values in the **sla** field in the YAML file. Then, apply the YAML file by using the **kubectl** command-line interface.
- If you no longer want the PVCs that are associated with a label to participate in any scheduled backup jobs, remove the SLA policy assignments by updating the **sla** field in the YAML file. To remove the label from all SLAs, modify the **sla** field as follows:

  ```
  sla: []
  ```

  Then, apply the YAML file by using the **kubectl** command-line interface.
- Update the **backupLabels**: *label_key: value* field to add, remove, or modify the labels.
- If you want to modify any other parameters, you must create a new request and specify a different request name (*name_of_request*) in the YAML file.

Related concepts

"Backup and restore types" on page 403
Container Backup Support provides multiple types of backup and restore functions for your PVC and other cluster resources. You can use the IBM Spectrum Protect Plus user interface or the Kubernetes or OpenShift command line to initiate backup and restore operations.

"SLA policies" on page 405
Service level agreement (SLA) policies define how often snapshot backup and copy backup operations are run, and how long snapshots and copy backups are retained. You can set up custom SLAs that meet your operational requirements.

"Container Backup Support requests" on page 466
To protect container data, you can submit Container Backup Support requests by using the Kubernetes or OpenShift command-line interface.

"Troubleshooting Container Backup Support" on page 662
To help troubleshoot issues with Container Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Backing up persistent volumes and resources by namespace using the command line

You can create backup requests for persistent volumes and resources by specifying a namespace. A physical cluster can be divided into virtual clusters that are called namespaces. By specifying a namespace in a backup request, you can back up all PVCs and resources in that namespace.

Before you begin

Backup requests are directed to persistent volume claims (PVCs) for the volumes that you want to protect. Before you schedule a backup job, take the following actions:

- Ensure that the PVC exists within the specified namespace.
- Ensure that the PVC is formatted. PVCs must be formatted before they can be backed up. For a PVC to be formatted correctly, it must be mounted and written to. Backup operations of raw block volumes are not supported.
- Determine which SLA policy to assign to PVCs. For instructions about viewing the available SLA policies, see "SLA policies" on page 405.
- For OpenShift clusters that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, ensure that you power off the VMs before running a backup job.

Creating application-consistent backups on OpenShift: If you want to create application-consistent backups of your data in an OpenShift environment, you can specify Velero pre and post backup hooks to quiesce an application before a snapshot is created, and to resume the application when the snapshot is completed. To start protecting your application in a consistent manner on a regular schedule, you must apply an SLA policy to those namespaces and PVCs where the application is running. In addition, you must specify pre and post hooks as annotations on the application pods. For instructions on how to specify these application hooks, see Backup Hooks in the Velero documentation.

About this task

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

Procedure

1. Optional: Display the list of PVCs in the namespace that you want to back up by issuing the following command:

```
kubectl get pvc -n namespace
```

2. Create a YAML file that defines the request for the backup-by-namespace operation. The YAML file must contain the following properties:

```
#------------------------
# Filename: filename.yaml
#------------------------

apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: name_of_request
  namespace: namespace
spec:
  requesttype: BackupNamespace
  sla: [sla_policy]
  volumesnapshotclass: snapshot_class_name
```

where:

filename
Specifies the name of the YAML configuration file. The file type is `.yaml`.

name_of_request
Specifies the name of the backup-by-namespace request. The name must be unique, and must not match the PVC name.

namespace
Specifies namespace that you want to assign a service level agreement (SLA) policy to.

After you assign the SLA at the namespace level, any new PVCs that you create in that namespace will automatically be assigned to the SLA.

[*sla_policy*]

Specifies the SLA policy that determines the schedule for backup operations. You can specify more than one SLA policy by using a comma-separated list within the brackets.

For example, to assign the `daily` policy to a PVC, specify the following statement:

```
sla: [daily]
```

To assign the `every4hours`, `daily_midnight`, and `weekly` policies to the PVC, specify the following statement in the YAML file:

```
sla: [every4hours,daily_midnight,weekly]
```

Alternatively, you can use the following format to specify a single SLA policy:

```
sla:
- daily
```

Or use the following format to specify multiple SLA policies:

```
sla:
- every4hours
- daily_midnight
- weekly
```

Ensure that you use the correct case when you specify the SLA policy name. Policy names are case-sensitive in YAML files.

To remove all SLA assignments from a namespace, delete the SLA policy names within the brackets, as shown in the following statement:

```
sla: []
```

snapshot_class_name
Specifies the snapshot class for the volume. If you do not specify the snapshot class, the default snapshot class is used if the sidecar container `csi-snapshotter` in the default snapshot class matches the provisioner of the volume. Otherwise, the backup request is invalid.

3. Submit the backup request by issuing the following command:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

Results

After you submit the backup request, the first scheduled backup operation will start within the window that is defined by the SLA policy. The start time of the backup is recorded in the backup status.

What to do next

To view information about the backup request, issue the **kubectl describe** command by using the request name. For example, to view information about a backup request that is named `backup-namepace1` in the `baas` namespace, issue the following command:

```
kubectl describe baasreq backup-namespace1 -n baas
```

For instructions, see "Viewing the status of backup and restore jobs" on page 486.

Modifying parameters in a YAML file:

After scheduled backup-by-namespace jobs have started, you can modify the SLA parameter in the YAML file and apply it to the same namespace if needed. For example:

- To assign a different SLA policy to the namespace or remove an SLA assignment, edit the values in the **sla** field in the YAML file. Then, apply the YAML file by using the **kubectl** command-line interface.
- If you no longer want the PVCs in a namespace to participate in any scheduled backup jobs, remove the SLA policy assignments by updating the **sla** field in the YAML file. To remove the namespace from all SLAs, modify the **sla** field as follows:

  ```
  sla: []
  ```

  Then, apply the YAML file by using the **kubectl** command-line interface.
- If you want to modify any other parameter, you must create a new request and specify a different request name (*name_of_request*) in the YAML file.

Related concepts

"Backup and restore types" on page 403
Container Backup Support provides multiple types of backup and restore functions for your PVC and other cluster resources. You can use the IBM Spectrum Protect Plus user interface or the Kubernetes or OpenShift command line to initiate backup and restore operations.

"SLA policies" on page 405
Service level agreement (SLA) policies define how often snapshot backup and copy backup operations are run, and how long snapshots and copy backups are retained. You can set up custom SLAs that meet your operational requirements.

"Container Backup Support requests" on page 466
To protect container data, you can submit Container Backup Support requests by using the Kubernetes or OpenShift command-line interface.

"Troubleshooting Container Backup Support" on page 662

To help troubleshoot issues with Container Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Restoring container data by using the command line

You can use the Kubernetes or OpenShift command line interface to restore a persistent volume from a snapshot or copy backup. A snapshot restore operation is generally faster than a copy restore operation.

Before you begin

To run **restoresource** requests for cluster-scoped and namespace-scoped resources, see "Restoring resources by using the command line" on page 484.

Review the following restrictions:

- For any type of restore operation, you cannot restore a volume to a different namespace or cluster.
- Snapshot or copy backups are restored to new persistent volumes. The persistent volume claim (PVC) for the new volume is automatically created when you restore the snapshot or copy backup. If you plan to use the same PVC name for the restored PVC, ensure that you delete the original PVC before the restore job runs. The restore operation does not overwrite existing PVCs. Before you delete the original PVC, take a snapshot of the PVC to ensure that the latest changes are saved.
- To help prevent issues with restore operations, do not manually delete snapshots of volumes that are protected by Container Backup Support.

For OpenShift clusters that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the custom resource data for the VMs is restored along with the resources. Before you begin the restore job, ensure that you prepare the VMs. For instructions, see "Preparing to restore persistent volumes that have VMs on OpenShift clusters" on page 444.

About this task

Depending on your recovery point objective and recovery time objective, you can run a `fast` restore or a `copy` restore operation:

- To restore a volume in the least amount of time, run a fast restore operation to restore a snapshot. If another operation is in progress on the same volume, the fast restore operation might take longer to complete.
- To restore a volume from a specified point in time from the IBM Spectrum Protect Plus vSnap server, run a copy restore operation.

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

Procedure

1. To view the restore points that are available for a PVC, query all the backups for the PVC by running the following command:

   ```
   kubectl describe BaaSReq pvc_name -n namespace
   ```

   Restore points are identified by the timestamp of the snapshot or copy backup.
2. In the status output that is displayed, identify the timestamp of the source snapshot or copy backup that you want to restore. The timestamps are shown in the `Status` section of the output before the type of backup.

   For example, the following output shows the timestamps for different types of backups:

```
Status:
 Timestamp:  2019-05-30 13:27:21
 Type:       FAST
 Timestamp:  2019-05-30 13:32:21
 Type:       COPY
```

where:

FAST
Denotes the backup type for a snapshot that is taken during a snapshot backup operation.

COPY
Denotes the backup type for a copy backup that is stored on an IBM Spectrum Protect Plus vSnap server.

3. To specify the restore request, create a YAML file with the following properties. Insert the timestamp for the source snapshot in the **restorepoint** parameter.

```
#------------------------
# Filename: filename.yaml
#------------------------

apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: name_of_restore_request
  namespace: namespace
spec:
  requesttype: restore
  pvcname: pvc_name
  targetvolume: target_volume_for_restore
  storageclass: storage_class_of_target_volume
  restorepoint: timestamp_of_backup
  restoretype: fast | copy
```

where:

filename
Specifies the name of the YAML configuration file.

name_of_restore_request
Specifies the name of the request for the restore job. The name must be unique, and must not match the name of the PVC.

A new restore request must be created for each subsequent restore of the same PVC. In other words, to restore a PVC again, create a new request and specify a different request name (*name_of_request*) in the YAML file.

namespace
Specifies the namespace for the request.

pvc_name
Specifies the name of the PVC that you want to restore.

target_volume_for_restore
Specifies the name of the PVC that you want to restore the volume to.

If you want to restore the volume to the original PVC, delete the original PVC and specify the same PVC name in this parameter. The restore operation does not overwrite existing PVCs. Before you delete the original PVC, take a snapshot of the PVC to ensure that the latest changes are saved.

storage_class_of_target_volume
Specifies the storage class that is defined for the target volume.

You can specify a storage class that is the same as the original PVC or specify a different storage class. If you do not specify the storage class, the storage class of the original PVC is used.

If you specify a storage class but do not specify the restore type with the **restoretype** parameter, a copy restore operation occurs.

timestamp_of_backup
Specifies the timestamp of the source snapshot or copy backup that you want to restore from. The timestamp is in Coordinated Universal Time (UTC) format.

If you do not specify a timestamp, the most recent snapshot or copy backup is restored.

restoretype: fast | copy
Specifies the type of restore operation to use.

fast
Restores a volume from a snapshot backup.

copy
Restores a volume from a copy backup.

This parameter is optional. If you do not specify a restore type, the type of restore is determined automatically. If a snapshot exists at the specified timestamp, a fast restore is run to restore the snapshot. If only a copy backup is available at the specified time, a copy restore is run to restore the copy backup.

4. Start the restore request by issuing the following command:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

What to do next

If you restored data to a new persistent volume, reconfigure the application container to mount the new volume after the snapshot or copy backup is restored.

To more efficiently manage your Container Backup Support requests, delete completed requests by issuing the following command:

```
kubectl delete baasreq name_of_restore_request -n namespace
```

By deleting completed requests, you gain the following benefits:

- The size of the `etcd` database is reduced and you can reuse the name of a request for another operation.
- The troubleshooting process is simplified.
- The tracking of backup and restore requests is simplified. At any time, you can obtain an accurate list of requests that are running in on your cluster when you issue the following command:

```
kubectl get baasreq -n namespace
```

Related concepts

“Backup and restore types” on page 403
Container Backup Support provides multiple types of backup and restore functions for your PVC and other cluster resources. You can use the IBM Spectrum Protect Plus user interface or the Kubernetes or OpenShift command line to initiate backup and restore operations.

“Container Backup Support requests” on page 466
To protect container data, you can submit Container Backup Support requests by using the Kubernetes or OpenShift command-line interface.

“Troubleshooting Container Backup Support” on page 662
To help troubleshoot issues with Container Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Related tasks

“Viewing the status of backup and restore jobs” on page 486

After you submit a backup or restore request, you can use the **kubectl get** and the **kubectl describe** commands to show information about your request.

Restoring resources by using the command line

You can use the Kubernetes or OpenShift command line interface to restore namespace-scoped resources and cluster-scoped resources from a snapshot or copy backup. A snapshot restore operation is generally faster than a copy restore operation, and is referred to as a *fast* restore on the command line.

Before you begin

Review the following restrictions:

- You cannot restore resources to an alternate namespace or cluster by using the command line. Both cluster-scoped and namespace-scoped resources must be restored to their original cluster or original namespace.
- Do not manually delete any cluster-scoped or namespace-scoped backups that are in use by Container Backup Support.

For OpenShift clusters that have virtual machines (VMs) running as part of the OpenShift Virtualization feature, the custom resource data for the VMs is restored along with the resources. Before you begin the restore job, ensure that you prepare the VMs. For instructions, see "Preparing to restore persistent volumes that have VMs on OpenShift clusters" on page 444.

About this task

Depending on your recovery point objective and recovery time objective, you can run a `fast` restore or a `copy` restore operation:

- To restore a resource from a specified point in time from the IBM Spectrum Protect Plus vSnap server, run a copy restore operation.

The following types of restore operation can be run from the command line:

1. On-demand restore operation of namespace-scoped resources
2. On-demand restore operation of cluster-scoped resources

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

Procedure

1. To view the restore points that are available for clusters or namespaces, query all the backups by running the following command:

   ```
   kubectl describe BaaSReq baasreq_name -n namespace
   ```

   Restore points are identified by the timestamp of the snapshot or copy backup.
2. In the status output that is displayed, identify the timestamp of the source snapshot or copy backup that you want to restore. The timestamps are shown in the `Status` section of the output before the type of backup.

   For example, the following output shows the timestamps for different types of backups:

   ```
   Status:
    Timestamp:  2020-10-30 13:27:21
    Type:       FAST
    Timestamp:  2020-10-30 13:32:21
    Type:       COPY
   ```

   where:

FAST
Denotes the backup type for a snapshot that is taken during a snapshot backup operation.

COPY
Denotes the backup type for a copy backup that is stored on an IBM Spectrum Protect Plus vSnap server.

3. To specify the restore request, create a YAML file with the following properties. Insert the timestamp for the source snapshot in the **restorepoint** parameter.

```
#------------------------
# Filename: filename.yaml
#------------------------
apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: request_name
  namespace: namespace
spec:
  backupbaasreq: request_used_for_resource_backup
  requesttype: RestoreResources
  scope: cluster | namespace
  restorepoint: timestamp_of_backup
  restoretype: fast | copy
```

where:

filename
Specifies the name of the YAML configuration file.

request_name
Specifies the name of the request for the restore job. The name must be unique, and must not match the name of the resource.

A new restore request must be created for each subsequent restore of the same resource. In other words, to restore a namespace-scoped or cluster-scoped resource again, create a new request and specify a different request name request_name in the YAML file.

namespace
Specifies the namespace for this restore request. For cluster-scoped resource restore jobs, this parameter specifies the namespace where the baasreq will be created.

backupbaasreq
Specifies the name of the request that was created to back up the resources. This baasreq is used to query the IBM Spectrum Protect Plus server for a list of available backups.

requesttype
The request type for the resources restore operation is specified as `RestoreResources` for the resource based restore request.

scope
Specifies whether this is a cluster-scoped or namespace-scoped restore request.

restorepoint
Is an optional parameter that specifies the *timestamp_of_backup* of the source snapshot or copy backup that you want to restore from. The timestamp is in Coordinated Universal Time (UTC) format.

restoretype: fast | copy
Specifies the type of restore operation to use.

fast
Restores a volume from a snapshot backup.

copy
Restores a volume from a copy backup.

This parameter is optional. If you do not specify a restore type, the type of restore is determined automatically. If a snapshot exists at the specified timestamp, a fast restore is run to restore the

snapshot. If only a copy backup is available at the specified time, a copy restore is run to restore the copy backup.

4. Start the restore request by issuing the following command:

```
kubectl create -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

What to do next

To more efficiently manage your Container Backup Support requests, delete completed requests by issuing the following command:

```
kubectl delete baasreq name_of_restore_request -n namespace
```

By deleting completed requests, you gain the following benefits:

- The size of the `etcd` database is reduced and you can reuse the name of a request for another operation.
- The troubleshooting process is simplified.
- The tracking of backup and restore requests is simplified. At any time, you can obtain an accurate list of requests that are running in on your cluster when you issue the following command:

```
kubectl get baasreq -n namespace
```

Managing container backup and restore jobs

You can query information about backup and restore jobs and delete snapshot and copy backups that are no longer needed.

Viewing the status of backup and restore jobs

After you submit a backup or restore request, you can use the **kubectl get** and the **kubectl describe** commands to show information about your request.

About this task

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

Procedure

1. To show a listing of all Container Backup Support requests in a namespace, issue the **kubectl get** command as follows:

```
kubectl get baasreq -n namespace
```

For example, to show all requests in the `production-01` namespace, issue the following command:

```
kubectl get baasreq -n production-01
```

The output is similar to the following example:

```
NAME         AGE
vol08-adhoc  17d
inv-adhoc2   17d
db-vol08     18d
db-vol09     17d
```

The request names are listed in the NAME column of the output.

2. Using the results from Step “1” on page 486, issue the **kubectl describe** command to show the status of a job. For example:

   - To show the list of all backups for any request, including backups from scheduled and on-demand backup requests, specify the name of the request and the namespace in the following command:

     ```
     kubectl describe baasreq request_name -n namespace
     ```

     where *request_name* is the name of the request. For on-demand backups, use the PVC name as the request name.

     For example, to show all backups for PVC `db-vol08` in the `production-01` namespace, issue the following command:

     ```
     kubectl describe baasreq db-vol08 -n production-01
     ```

     The output is similar to the following example:

     ```
     kubectl describe baasreq db-vol08 -n production-01
     Name:          db-vol08
     Namespace:     production-01
     Labels:        <none>
     Annotations:   <none>
     API Version:   baas.io/v1alpha1
     Backupstatus:  Ready
     Kind:          BaaSReq
     Metadata:
       Creation Timestamp:  2020-05-20T20:28:33Z
       Generation:          9
       Resource Version:    2955966
       Self Link:           /apis/baas.io/v1alpha1/namespaces/production-01/baasreqs/db-vol08
       UID:                 0e8d4412-522f-44b3-932c-1e6239f7bf8e
     Spec:
       Inprogress:   None
       Instanceid:   e05c400868ab9151e3c792d28edfbb18
       Origreqtype:  backup
       Requesttype:  backup
       Size:         1073741824
       Sla:
         joanne-copy2
       Spppvcname:           cluster01:production-01:db-vol08
       Volumesnapshotclass:  cirrus-csi-rbdplugin-snapclass
     Status:
       Snapshotname:  spp-1005-2161-172342eb32d
       Timestamp:     2020-05-20 22:24:25
       Type:          FAST
       Snapshotname:  2000.snapshot.824
       Timestamp:     2020-05-20 21:13:27
       Type:          COPY
       Snapshotname:  spp-1005-2161-17233c4e7a0
       Timestamp:     2020-05-20 20:28:14
       Type:          FAST
     ```

   - To show information about a restore job, issue the following command:

     ```
     kubectl describe baasreq request_name -n namespace
     ```

     where *request_name* is the request name of the restore job and *namespace* is the namespace of the PVC that was restored.

Results

In the command output, the **Backupstatus** field shows the status of a backup job. For restore jobs, the **Restorestatus** field shows the status of the restore job. For more information, see “Status of backup and restore jobs” on page 488.

The **instanceid** field contains a randomly generated string that uniquely identifies a volume in IBM Spectrum Protect Plus.

The **Spppvcname** field shows the name of the PVC that is reported in the IBM Spectrum Protect Plus **Jobs and Operations** window. The *namespace*:*pvc_name* format is used to identify the PVC. The values for the **instanceid** and **Spppvcname** fields uniquely identify a backup in IBM Spectrum Protect Plus.

In backup requests, the **Status** section shows the list of backups that were completed. For each backup, the timestamp of the backup is listed, followed by the type of backup that was run. The types of backups are defined as follows:

FAST
Denotes the backup type for a snapshot that is taken during a snapshot backup operation.

COPY
Denotes the backup type for a copy backup that is stored on an IBM Spectrum Protect Plus vSnap server.

Status of backup and restore jobs

When you use the **kubectl describe** command to show information about backup and restore jobs, the status of backup and restore jobs is displayed in the command output.

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

To display the status of a specific Container Backup Support request, enter the following command:

```
kubectl describe baasreq request_name -n namespace
```

where *request_name* is the name of the request, and *namespace* is the namespace in which the persistent volume exists. For more information, see “Viewing the status of backup and restore jobs” on page 486.

Reported backup status

The status of a backup job is shown in the `Backupstatus` field in the command output. The following table shows the possible statuses of a backup request:

Table 101. Status of backup jobs

Backup status	Description
None	No backup jobs were started for this schedule.
Requested	A backup job was started for this schedule.
Ready	At least one backup job was completed for this schedule.
Destroyed	All snapshot and copy backups of a persistent volume claim were deleted.
Invalid	An issue occurred with the request. A possible explanation is listed in the **Errmsg** field.

Reported restore status

The status of a restore job is shown in `Restorestatus` field in the command output. The following table shows the possible statuses of a restore job:

Table 102. Status of restore jobs

Restore status	Description
None	No restore jobs were requested.
Requested	A snapshot or a copy backup restore job is requested.

Table 102. Status of restore jobs (continued)	
Restore status	**Description**
Restored	A snapshot or a copy backup was successfully restored.
Invalid	An issue occurred with the request. A possible explanation is listed in the **Errmsg** field.

Deleting container backups

You can mark for deletion snapshot and copy backups of a persistent volume claim (PVC) by submitting a **destroy** request.

Before you begin

Before you submit a **destroy** request to delete container backups, consider the following consequences:

- All snapshots of the PVC are deleted when their expiration dates are reached as defined by the service level agreement (SLA) policy for the PVC.
- The snapshot and copy backups on the IBM Spectrum Protect Plus vSnap server will be marked for deletion. The deletion is managed by IBM Spectrum Protect Plus.
- The original backup request will not be deleted by the **destroy** request. You must run the **kubectl delete** command to delete it.
- The **destroy** request is not supported for on-demand backups. Use the **kubectl delete** command to delete an on-demand backup request. An on-demand snapshot is deleted when the snapshot expires or when the scheduled backup is destroyed.

Unless specified otherwise, you can use either the **kubectl** or **oc** command in the OpenShift environment.

Procedure

1. Create a YAML file for the **destroy** request that contains the following properties:

```
#------------------------
# Filename: filename.yaml
#------------------------

apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
 name: request_name
 namespace: namespace
spec:
 requesttype: Destroy
```

   where:

   filename
   The name of the YAML configuration file.

   request_name
   The name of the request, which must match the name of the PVC that was backed up. For example, if you want to delete all snapshots and copy backups for the PVC named `db-vol01`, the name of the request must also be `db-vol01`.

   namespace
   The namespace in which the PVC exists.

2. Submit the **destroy** request by entering the following command on the command line:

```
kubectl apply -f filename.yaml
```

where *filename* is the name of the YAML configuration file.

3. To check that the snapshots and copy backups for a PVC are deleted, issue the following command:

```
kubectl describe baasreq request_name -n namespace| grep Backupstatus
```

where *request_name* is the name of the PVC that was backed up.

In the command output, the following status shows that the backups were deleted:

```
Backupstatus:  Destroyed
```

What to do next

As a best practice, delete the completed request by issuing the following command:

```
kubectl delete baasreq request_name -n namespace
```

where *request_name* is the name of the PVC that was backed up.

By deleting completed requests, you gain the following benefits:

- The size of the `etcd` database is reduced and you can reuse the name of a request for another operation.
- The troubleshooting process is simplified.
- The tracking of backup and restore requests is simplified. At any time, you can obtain an accurate list of requests that are running in on your cluster when you issue the following command:

```
kubectl get baasreq -n namespace
```

If you delete the backup request without first destroying the backup, the backup request will continue to run and backups will be made according to the specified SLA policy until Container Backup Support is restarted.

Related information

“Types of requests in Container Backup Support” on page 466

Chapter 14. Protecting data on cloud systems

Cloud systems such as Microsoft 365 is a cloud-based subscription service that can be registered with IBM Spectrum Protect Plus so that you can start to protect your data. Register Microsoft 365 with IBM Spectrum Protect Plus so that you can set up backup jobs or regularly scheduled service level agreement (SLA) policies to protect your data..

If you choose to protect Microsoft 365 with IBM Spectrum Protect Plus, you need to purchase IBM Spectrum Protect Plus for Microsoft 365 Entity ID Monthly License, Part Number D25ZELL. For more information about this entitlement, see IBM Spectrum Protect Plus V10.1.5 announcement letter.

Microsoft 365

To protect Microsoft 365 email, calendars, contacts, and data on OneDrive cloud storage, you must register the Microsoft 365 application with Azure Active Directory.

Then, complete the following steps:

1. Deploy the Linux server that acts as the proxy host server.
2. Add a Microsoft 365 tenant and assign the cloud proxy server to the tenant.
3. Define a service level agreement (SLA) policy to create backup jobs.

If you choose to protect Microsoft 365 with IBM Spectrum Protect Plus, you need to purchase the IBM Spectrum Protect Plus for Microsoft Office 365 Entity ID Monthly License, Part Number D25ZELL. For more information about this entitlement, see the IBM Spectrum Protect Plus 10.1.5 announcement letter.

Registering with Azure Active Directory

To protect a Microsoft 365 application, you must register the application with Azure Active Directory and grant appropriate permissions. When you register a new application with Azure Active Directory, the application credentials such as application ID and application secret are made available on the Azure Active Directory portal.

Before you begin

Take the following actions:

- Ensure that you have an active Microsoft 365 subscription.
- Ensure that you have a Microsoft 365 administrative user ID and password.

  Tip: You can use an automated process to register and configure the Microsoft Azure Active Directory application. For instructions, see technote 6437493.

Procedure

1. Go to the Microsoft 365 welcome page and sign in to your account by using your Microsoft 365 administrative user ID and password.
2. To open the Azure Active Directory admin center, in the left pane, click the ellipsis to expand the **Show all** menu, and then click **Admin centers** > **Azure Active Directory**.
3. To open your tenant dashboard, in the left pane of the Azure Active Directory admin center, click **Azure Active Directory**.
4. In the tenant dashboard menu, click **App registrations** and then click **New registration**.
5. To specify a user-facing name for the Microsoft 365 application, on the "Register an application" page, enter a name in the **Name** field.
6. Use the default options for the remaining fields, and click **Register**. The app registration is set up with the user-facing name that you entered.

© Copyright IBM Corp. 2017, 2021

7. To obtain the application (client) ID and directory (tenant) ID string, click **Azure Active Directory** > **tenant - App registrations** > **App name**. Then, copy the application ID string and directory ID. These strings will be required later, when you register the Microsoft 365 application with IBM Spectrum Protect Plus.
8. To create a client secret for this application ID, click **Certificates & secrets** > **New client secret**.
9. In the "Add a client secret" pane, enter any username in the **Description** field, and click **Add**. A client secret is generated, and the value is displayed in the "Client secrets" pane.
10. Copy the client secret to the clipboard by using the copy facility next to the **Client secret value** field. This character string is also used for registration with IBM Spectrum Protect Plus.
11. To add permissions for this application ID, click **API permissions** > **Add permissions**.
12. Specify permissions for each API in the following table by taking the following actions:
    a) Select the API name, for example, Azure Active Directory Graph.
    b) For the permission name User.Read.All, select the **Delegated Permissions** type.
    c) For the remaining permissions, select the **Application Permissions** type for each permission name for the APIs that are listed in the table.

API	Permission name
Azure Active Directory Graph	Directory.Read.All
Azure Active Directory Graph	User.Read.All
Office 365 Exchange Online	full_access_as_app
Microsoft Graph	Calendars.ReadWrite
Microsoft Graph	Contacts.ReadWrite
Microsoft Graph	Files.ReadWrite.All
Microsoft Graph	Mail.Read
Microsoft Graph	Mail.ReadWrite
Microsoft Graph	Mail.Send
Microsoft Graph	Sites.Read.All
Microsoft Graph	User.Read
Microsoft Graph	User.Read.All
Microsoft Graph	User.ReadWrite.All

13. To save the selected permissions, click **Grant admin consent for *your organization name***, where *your organization name* specifies the name of your organization.

What to do next

Follow the instructions in “Registering the Microsoft 365 tenant with IBM Spectrum Protect Plus” on page 492.

Registering the Microsoft 365 tenant with IBM Spectrum Protect Plus

To ensure that the IBM Spectrum Protect Plus agent can connect to the Microsoft 365 tenant, you must register the Microsoft 365 tenant credentials, and the Linux server that acts as the proxy host server with

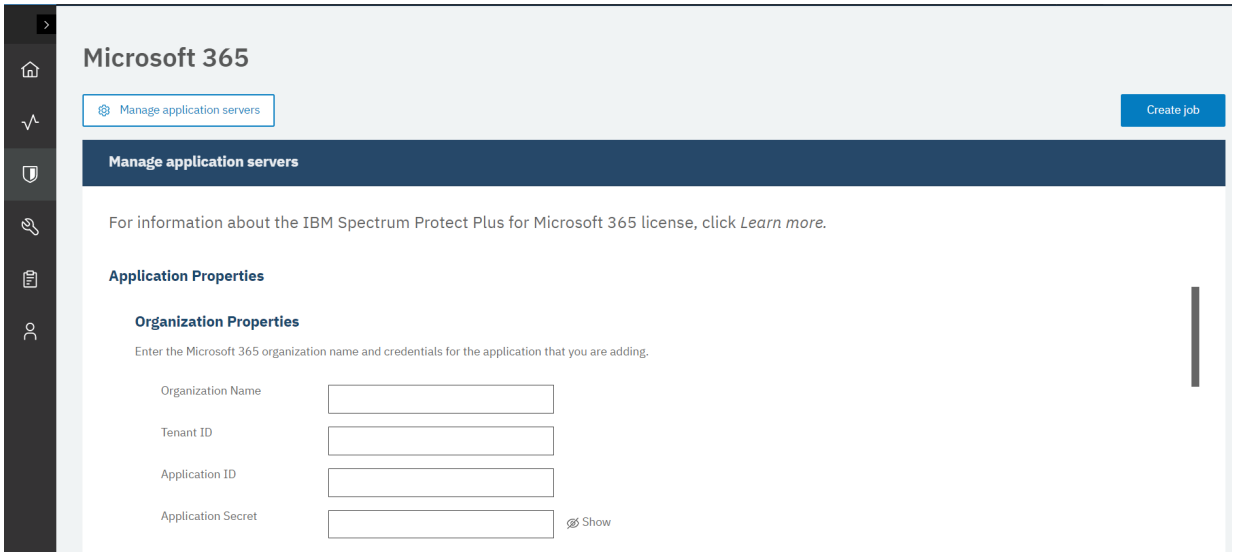
IBM Spectrum Protect Plus. This procedure is necessary to ensure that Microsoft 365 data can be backed up to IBM Spectrum Protect Plus.

Before you begin

- Ensure that you have a Linux system that can act as the cloud proxy server. IBM Spectrum Protect Plus deploys the backup agent on this server. For more information about the requirements, see Microsoft 365 requirements.
- Ensure that the Microsoft 365 application is registered with Azure Active Directory. For instructions, see “Registering with Azure Active Directory ” on page 491.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Cloud Management** > **Microsoft 365**.



2. On the Microsoft 365 page, click **Manage application servers**, and then click **Add application server**.
3. On the Organization Properties page, complete the following fields:
   a. In the **Organization Name** field, enter the name of the organization that you set up in the Azure Active Directory admin center.

      Note: Specify the Organization name in this format: *tenantname.onmicrosoft.com*, The names are not visible when you register the Azure application.

   b. In the **Tenant ID** field, enter the string from the **Directory (tenant) ID** field in the Azure Active Directory application registration.
   c. In the **Application ID** field, enter the string from the **Application (client) ID** field in the Azure Active Directory application registration.
   d. In the **Application Secret** field, enter the password string that was generated during the Azure Active Directory application registration.
4. On the Proxy Properties page, complete the following fields:
   a. In the **Host Address** field, enter the host name or IP of the Linux server that is being used as the proxy host.
   b. For host server authentication, select one of the following options:

   - **User**: Select an existing user, or enter a user ID and the associated password.
   - **SSH Key**: Select a Secure Shell (SSH) key from the drop-down list.
5. Click **Save**.

Results

When a proxy host is registered with IBM Spectrum Protect Plus, an inventory is run automatically on the Microsoft 365 organization, which returns the Microsoft 365 users in that resource.

Detailed process logs

The detailed process log is an additional Microsoft 365 process log file that can be useful for troubleshooting. This log is collected to track all backup and restore processes.

A detailed process log tracks the processes for each protected Microsoft 365 item. When you download the job log .zip file, you can view the detailed process log file along with standard diagnostic files.

Note: After you download the `joblog.zip`file, you can unzip the diag.tar.gz files to find the Audit.log file. This is the file with the Microsoft 365 processing information.

Detailed process log content and example

A detailed process log file includes the following information:

- Date and time of the operation.
- Operation type.
- Account that is associated with the operation.
- Indication of whether the event relates to OneDrive, a message, an calendar event, or a contact.
- Informational messages:
  - For OneDrive, the path and file name of the processed object is listed. If the operation is a redirected restore operation, that is indicated.
  - For messages, the date and time of the message is listed. If the operation is a redirected restore operation, any associated messages are listed.
  - For events, the subject of the event is listed.
  - For contacts, the name of the contact is listed.

Detailed process log example

The information in the detailed process log is provided in the following format:

```
[date time] [operation] [account] [relation] [message1] optional: [message2]
```

For example,

```
2020-02-13 19:15:27.805   Backup Completed  username@example.com  OneDrive
"my_new_document.pdf"
2020-02-13 19:13:46.754   Backup Completed  username@example.com  Message  "1/20/2020 10:52:01
PM +01:00"  "Welcome!"
2020-02-13 19:16:14.196   Backup Completed  username@example.com  Contact  "John Smith"
2020-02-13 19:14:48.847   Backup Completed  username@example.com  Event  "Monday meeting"
2020-02-13 19:18:22.544   Backup Failed  username@example.com  OneDrive
"my_folder\inventory.pdf"
2020-02-13 19:15:27.805   Restore Completed  username@example.com  OneDrive
"my_new_document.pdf" "my_new_document__2020-02-11_19_15.pdf"
2020-02-13 19:22:28.238   Backup Failed  username@example.com  OneDrive
"my_folder\inv\inventory.pdf"
```

Backing up Microsoft 365 data

After your Microsoft 365 organization is registered with IBM Spectrum Protect Plus, you can apply a service level agreement (SLA) policy to start protecting the Microsoft 365 data.

Procedure

1. In the IBM Spectrum Protect Plus navigation panel, expand **Manage Protection** > **Cloud Management** > **Microsoft 365**.
2. Select the checkbox for the organization.
3. Click **Select an SLA policy** and choose an SLA policy.

   For more information about SLA policies, see “Create backup policies” on page 259.
4. Save your choice. To define a new SLA or to edit an existing policy with custom retention periods or backup frequency rates, click **Manage Protection** > **Policy Overview**. In the "SLA policies" pane, click **Add SLA Policy**, and define policy preferences.

   Tip: Some options in the **Policy Options** field in the **SLA Policy Status** section differ in availability based on backup type.
5. To run the policy outside the scheduled job, take the following actions:
   a. To back up all organization data, select the checkbox for the organization.
   b. To back up data from an account, click **Organization** and select the checkbox for the user name that is associated with the account.
   c. To back up email, calendars, contacts, or OneDrive data for an account, click **Organization**, click area that you want to protect, and then click the user name and select the checkbox for the email, calendar, contacts, or OneDrive to back up.
6. Click **Run**. The status changes to **running** for the SLA and you can follow the progress of the job in the log.

Incremental forever backup for Microsoft 365

IBM Spectrum Protect Plus supports a backup strategy that is named *incremental forever*. Rather than scheduling periodic full backup jobs, this backup strategy requires only one initial full backup. Afterward, an ongoing sequence of incremental backup jobs occurs.

The incremental forever backup solution provides the following advantages:

- Reduces the amount of data that goes across the network
- Reduces data growth because all incremental backups contain only the objects that are new or changed since the previous backup job
- Reduces the duration of backup jobs

The IBM Spectrum Protect Plus incremental forever process includes the following steps:

1. The first backup job backs up all data from selected Microsoft 365 accounts.
2. All subsequent backup jobs back up only new or changed data from the selected accounts.

Restoring Microsoft 365 data

You can restore Microsoft 365 data from backup copies on vSnap servers or remote storage.

Before you begin

At least one Microsoft 365 backup job must have run successfully. For instructions about setting up a backup job, see “Backing up Microsoft 365 data” on page 495.

About this task

The following restore modes are supported:

- Restore data to the original account
- Restore data to another account
- Restore data to a specified path

Procedure

1. In the navigation panel, expand **Manage Protection** > **Cloud Management** > **Microsoft 365**.
2. Click **Create job**.
3. Select **Restore**.
4. In the **Select source** pane, complete the following steps:
   a) Click a source in the list to display the data that can be restored for the selected organization. You can also use the search function to search for available data and toggle the displayed data by using the **View** filter.
   b) To select data to restore, click the Add to restore list icon next to the data. You can select more than one item from the list. The selected items are added to the restore list. To remove an item from the source list, click the Remove from restore list icon next to the data.
   c) Click **Next** to continue.
5. On the "Source snapshot" page, select the restore type and the time when the data to be restored was backed up. Click **Next**.
6. On the "Select destination" page, complete the following fields, and click **Next** to continue.

Option	Description
Select a destination	Select the location to which data must be restored: **Restore to original account** Restores data to the original Microsoft 365 account **Restore to another account** Restores data to another Microsoft 365 account
Restore Path	Restores data to a selected directory path in the Microsoft 365 account

7. On the **"Job options"** page, if you want to run restore operations in parallel streams, specify a value in the **Max Parallel Streams** field. Click **Next**.
8. On the Review page, review your restore job settings.
9. To start the restore job, click **Submit**.

Results

A few moments after you click **Submit**, the on-demand restore job is added to the Running Jobs tab on the Jobs and Operations page. You can click the job record to display the details of the operation. You can also download the zipped log file by clicking **Download.zip**.

The account name for the restored data can be found in the log file for the restore operation. To locate the logs for a restore operation, in the navigation panel, click **Jobs and Operations** and then click the **Running Jobs** tab.

Chapter 15. Protecting databases

You must register the database applications that you want to protect in IBM Spectrum Protect Plus and then create jobs to back up and restore the databases and resources that are associated with the applications.

Restriction: IBM Spectrum Protect Plus might create folders on application servers when applications are registered with IBM Spectrum Protect Plus. Folders created by IBM Spectrum Protect Plus must remain for the product to function properly. However, if you must remove a folder that was created by IBM Spectrum Protect Plus, unregister the application and IBM Spectrum Protect Plus will clean up the folders that are associated with the registration.

Do not assign more than one application per machine as an application server to a resource group. For example, if Microsoft SQL Server and Microsoft Exchange Server occupy the same machine and both are registered with IBM Spectrum Protect Plus, only one of the applications can be added as an application server to a given resource group.

Db2

After you successfully add your IBM Db2 instances to IBM Spectrum Protect Plus, you can start to protect your Db2 data. Create service level agreements (SLA) policies to back up and maintain Db2 data.

Ensure that your Db2 environment meets the system requirements. For more information, see “Db2 requirements” on page 83.

Tip: If your Db2 data is stored in a multi-partitioned environment with multiple hosts, you can protect your Db2 data across each host. Each host in the multi-partitioned environment must be added to IBM Spectrum Protect Plus so that all instances and databases are detected for protection. For more information, see “Adding a Db2 application server” on page 500.

The IP address must be reachable from the IBM Spectrum Protect Plus server and from the vSnap server. Both must have a Windows Remote Management service that is listening on port 5985.

The fully qualified domain name must be resolvable and routable from the IBM Spectrum Protect Plus appliance server and from the vSnap server.

Prerequisites for Db2

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

Requirements for the IBM Spectrum Protect Plus Db2 application server are available here, Db2 requirements.

Space prerequisites

Ensure that you have enough space on the Db2 database management system, in the volume groups for the backup operation, and on the target volumes for copying files during the restore operation. For more information about space requirements, see Space requirements for Db2 protection. When you are restoring data to an alternative location, allocate extra dedicated volumes for the copy and restore processes. The data paths for table spaces and logs on the target host are the same as the paths on the original host. This setup is needed to allow copying of data from the mounted vSnap to the target host. Ensure that the protected Db2 database does not share the same local database directory with other Db2 databases.

© Copyright IBM Corp. 2017, 2021

Multi-partitioned Db2 environments

In order to protect Db2 multi-partitioned databases, the ACS backup mode must be set to parallel mode. To run parallel backup processing of partitions in your Db2 environment, ensure that one of the following prerequisites is met:

- The Db2 registry variable **DB2_PARALLEL_ACS** is set to YES, for example: **db2set DB2_PARALLEL_ACS=**YES.
- The Db2 registry variable **DB2_WORKLOAD** is set to SAP.

Restriction: The **DB2_PARALLEL_ACS** registry variable is available only in certain fix pack levels of Db2. If **DB2_PARALLEL_ACS** is not available in your version, you can choose to change **DB2_WORKLOAD** to SAP.

More configuration requirements

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

- Db2 archive logging is activated, and Db2 is in recoverable mode.
- Ensure that the effective file size **ulimit** `-f` for the IBM Spectrum Protect Plus agent user and the Db2 instance user, is set to `unlimited`. Alternatively, set the value to a sufficiently high value to allow copying of the largest database files in your backup and restore jobs. If you change the **ulimit** setting, restart the Db2 instance to finalize the configuration.
- If you are running IBM Spectrum Protect Plus in an AIX or Linux environment, ensure that the installed sudo version is at the recommended level. For more information, see technote 2013790. Then, set sudo privileges as described in "Setting sudo privileges for Db2" on page 500.
- In a Linux environment, ensure that the Linux utility package `util-linux-ng` or `util-linux` package is current.
- Unicode characters in file path names cannot be handled by IBM Spectrum Protect Plus. All names must be in ASCII.
- The database table spaces, online logs, and the local database directory can be on one or separate dedicated logical volumes that are managed by either LVM2 or JFS2. For layout two examples, see the following pictures. In the first picture, two types of volume groups shown. In the second picture, all volumes for data and logs are on one volume group.

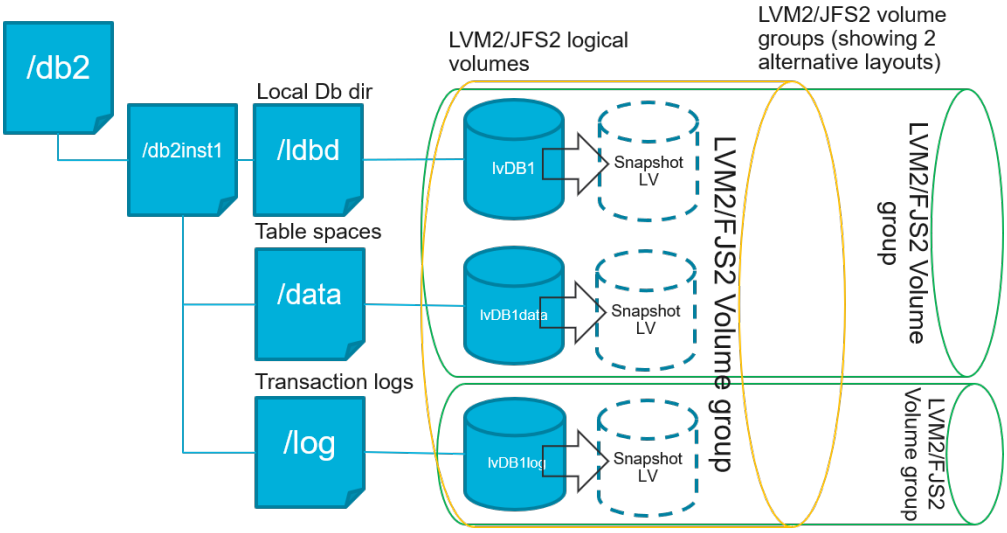


Figure 41. Logical volume layout examples

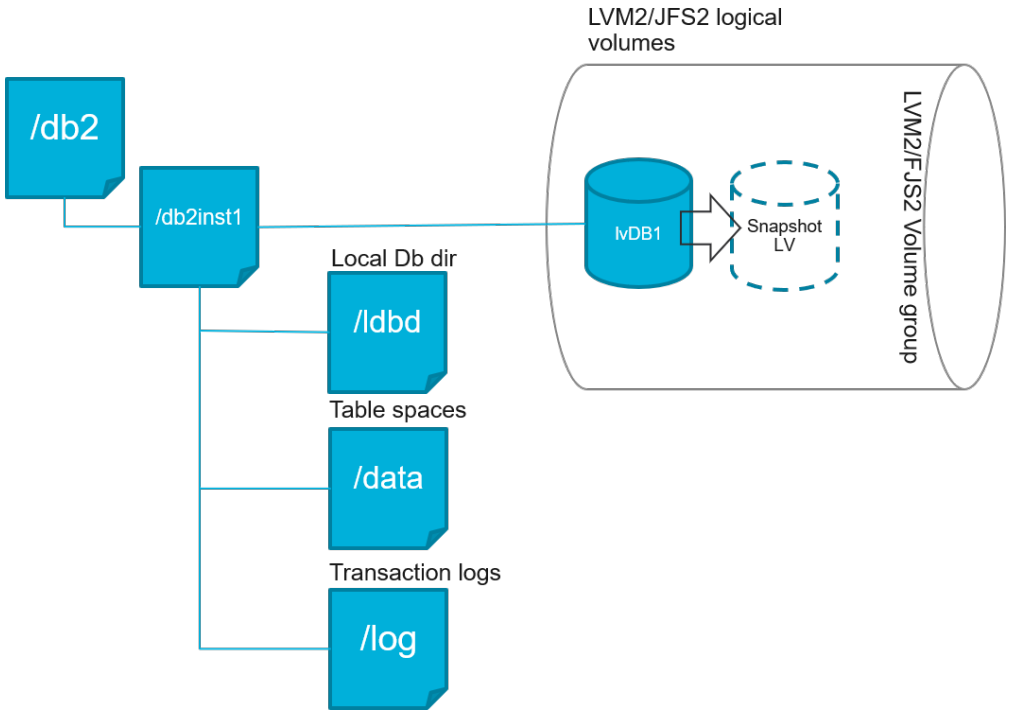


Figure 42. Single logical volume layout example

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

Space requirements for Db2 protection

Before you start backing up Db2 databases, ensure you have enough free disk space on the target and source hosts, and in the vSnap repository. Extra free disk space is required on the volume groups on the source host for creating temporary Logical Volume Manager (LVM) snapshots of the logical volumes that the Db2 database and log files are stored on. To create LVM snapshots of a protected Db2 database, ensure that the volume groups with Db2 data have sufficient free space.

LVM snapshots

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

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

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

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

LVM2 and JFS2

When you run a Db2 backup operation, Db2 requests a snapshot. This snapshot is created on a Logical Volume Management (LVM) system or a Journaled File System (JFS) for each logical volume with data or logs for the selected database. In Linux systems, the logical volumes are managed by LVM2 with `lvm2` commands. On AIX, the logical volumes are managed by JFS2 and created with the JFS2 snapshot command as external snapshots.

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

On the Linux operating system, the LVM2 volume manager stores the snapshot of a logical volume within the same volume group. On the AIX operating system, the JFS2 volume manager stores the snapshot of a logical volume within the same volume group. For both, there must be enough space on the machine to store the logical volume. The logical volume grows in size as data changes on the source volume while the snapshot exists. In multi-partitioned environments, when multiple partitions share the same volume, an extra snapshot of the volume is created for each partition. Ensure that the volume group has sufficient free space for the required snapshots.

Setting sudo privileges for Db2

To use IBM Spectrum Protect Plus to protect your data, you must install the required version of the sudo program. For the Db2 application server, you must set up sudo in a specific way that might be different from other application servers.

Before you begin

To determine the correct version of sudo to be installed, see technote 2013790.

About this task

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

Procedure

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

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

   where `agent` specifies the name of the IBM Spectrum Protect Plus agent user.
2. Set a password for the new user by issuing the following command:

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

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

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

Adding a Db2 application server

To start protecting your Db2 data, you must add the host address where your Db2 instances are located. You can repeat the procedure to add every host that you want to protect with IBM Spectrum Protect

Plus. If your Db2 environment is multi-partitioned with multiple hosts, you must add each host to IBM Spectrum Protect Plus.

About this task

To add a Db2 application server to IBM Spectrum Protect Plus, you must have the host address of the machine.

Procedure

1. In the navigation, expand **Manage Protection** > **Applications** > **Db2**.
2. In the **Db2** window, click **Manage application servers**, and click **Add application server** to add the host machine.

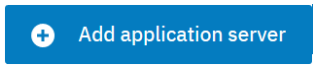


Figure 43. Adding a Db2 agent

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

   Note: The user must have sudo privileges set up.

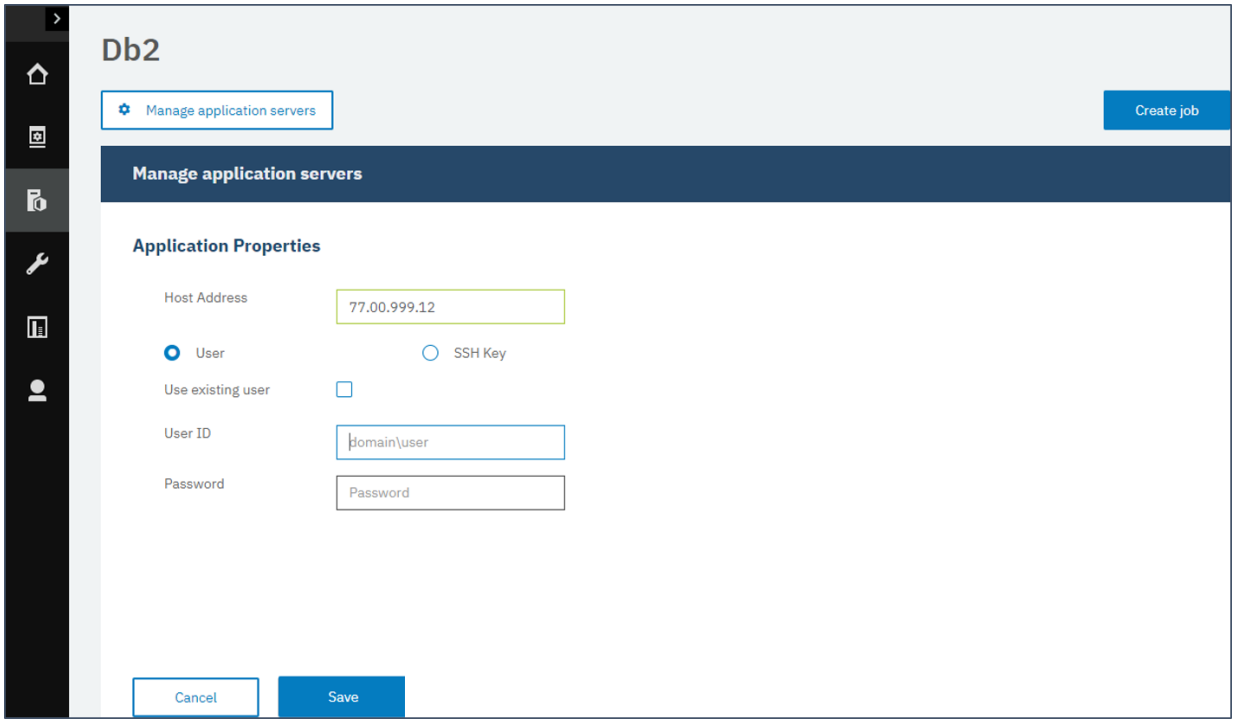


Figure 44. Managing agent users

Tip:

Db2 instances found are listed for each host. If your Db2 instance is partitioned, this information is listed with the host machine and the numbers of the partitions. For multi-host Database Partitioning Feature (DPF), the Db2 instance is displayed as a single unit.

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

If your Db2 data is in a multi-partitioned environment with multiple hosts, you must add each host. Repeat the procedure for each Db2 host.

What to do next

After you add your Db2 application servers to IBM Spectrum Protect Plus, an inventory is automatically run on each application server to detect the relevant databases in those instances.

To verify that the databases are added, review the job log. Go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.

Databases must be detected to ensure that they can be protected. For instructions about running an inventory, see Detecting Db2 resources.

Detecting Db2 resources

After you add IBM Db2 application servers to IBM Spectrum Protect Plus, an inventory to detect all Db2 instances and databases is run automatically. The inventory detects, lists, and stores all the Db2 databases for the selected host, and makes the databases available for protection with IBM Spectrum Protect Plus.

Before you begin

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

About this task

Any Db2 partitions that are found in the inventory are listed for the Db2 instance. Partitions are listed by their partition number for each host appended to the host name in the **Instances** table.

Procedure

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

   Tip: To add more Db2 instances to the **Instances** pane, follow the instructions in Adding a Db2 application server.
2. Click **Run Inventory**.

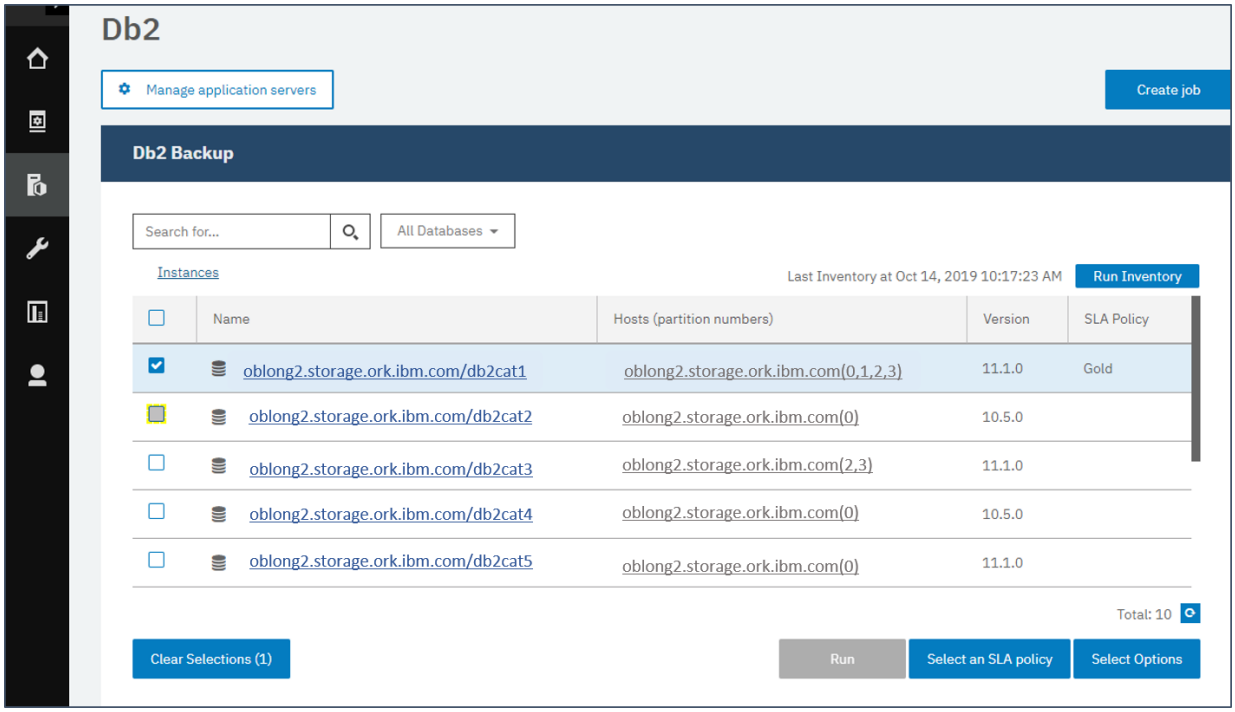


Figure 45. Detecting Db2 resources

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

To view the job log, go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.

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

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

What to do next

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

Testing the Db2 connection

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

Procedure

1. In the navigation panel, click **Manage Protection** > **Applications** > **Db2**.
2. In the **Db2** window, click **Manage Application Servers**, and select the **Host Address** you want to test.

   A list of the Db2 application servers that are available are shown.

3. Click **Actions** and choose **Test** to start the verification tests for physical, remote and operating system connections and settings.

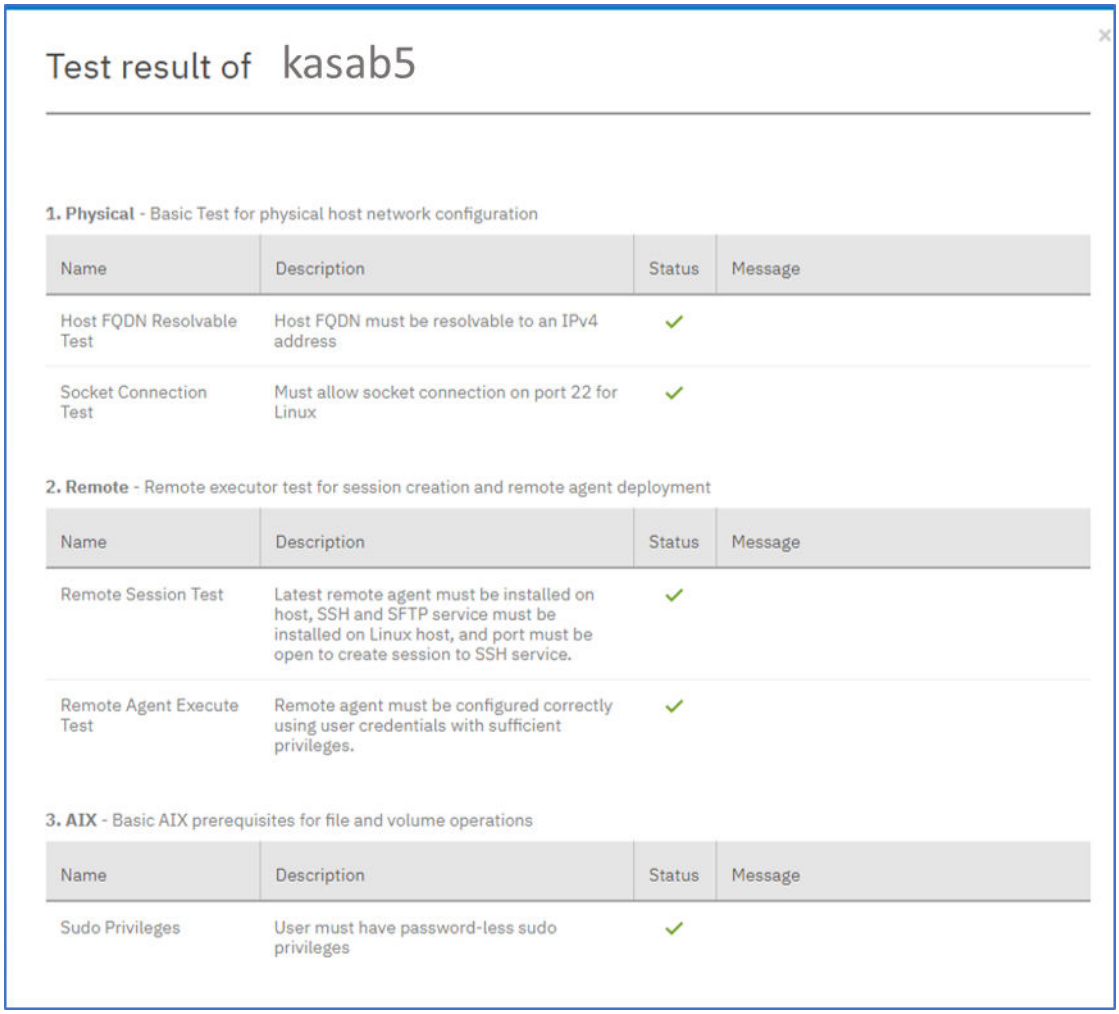


Figure 46. Testing the connection

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

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

Backing up Db2 data

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

Before you begin

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

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

- Add the application servers that you want to back up. For the procedure, see Adding a Db2 application server.

- Configure a Service Level Agreement (SLA) Policy. For the procedure, see Defining a Service Level Agreement backup job.
- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see Chapter 19, "Managing user access," on page 643.
- Inventory jobs should not be scheduled to run at the same time as backup jobs.
- Avoid configuring log backups for a single Db2 database with many backup jobs. If a single Db2 database is added to multiple job definitions with log backup enabled, a log backup from one job can truncate a log before it is backed up by the next job. This might cause point-in-time restore jobs to fail.

About this task

The following steps describe how to back up resources that are assigned to an SLA policy. To run an on-demand backup job for one or more resources regardless of whether those resources are already associated with an SLA policy, click **Create job**, select **Ad hoc backup**, and follow the instructions in "Running an ad hoc backup job" on page 627.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Applications** > **Db2**.
2. Select a resource to back up.
   - Select an entire instance in the **Instances** pane by clicking the instance name check-box. Any databases added to this instance are automatically assigned to the SLA policy that you choose.
   - Select a specific database in an instance by clicking the instance name, and choosing a database from the list of databases in that instance.
3. Click **Select Options** to enable or disable log backup, and to specify parallel streams to minimize time taken for large data movement in the backup operation. Click **Save** to commit the options.

   Tips for configuring options:

   Review the following tips to help you configure options for the backup job:

   - To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
   - If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
   - Options that are shown in yellow indicate that the option value has changed from the previously saved value.
   - To close the **Options** pane without saving changes, click **Select Options**.

   Select **Enable Log Backup** to back up archive logs, which allows point-in-time restore options and recovery options. For Db2 log backup settings information, see Log backups.

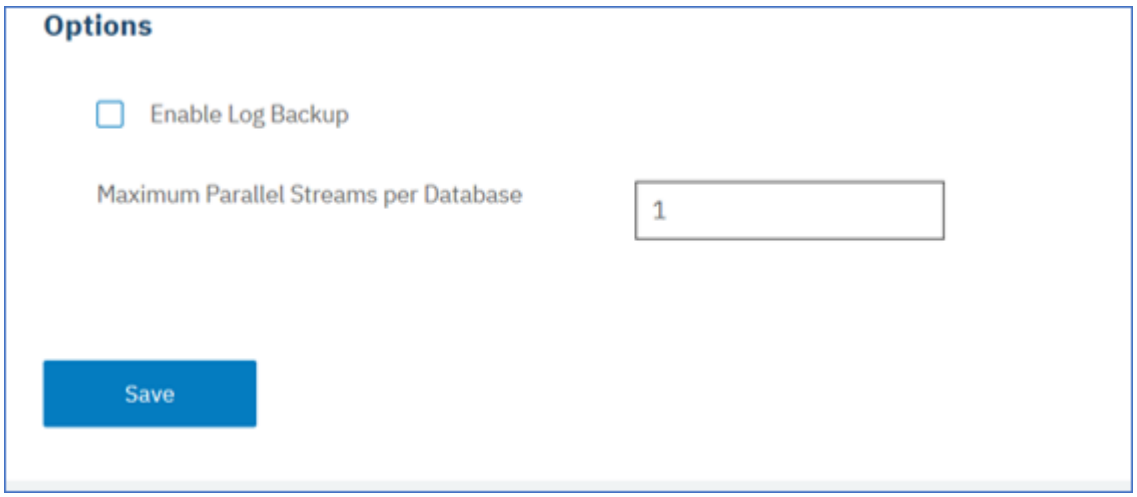


Figure 47. Backup pane with the Enable Log Backup option

If an on-demand job runs with the **Enable Log Backup** option enabled, log backup occurs. However, when the job runs again on a schedule, the option is disabled for that job run to prevent possible missing segments in the chain of backups.

When you save the options, those options are used for all backup jobs for this database or instance as selected.

4. Select the database or instance again, and click **Select SLA Policy** to choose an SLA policy for that database or instance.
5. Save the SLA options.

   To define a new SLA or to edit an existing policy with custom retention and frequency rates, select **Manage Protection** > **Policy Overview**. In the **SLA Policies** pane, click **Add SLA Policy**, and define your policy preferences.

What to do next

When the SLA policy is saved, you choose to run an on-demand backup any time by clicking **Actions** for that policy, and selecting **Start**. The status in the log changes to show that the backup is `Running`.

Defining a service level agreement backup job

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

Procedure

1. From the navigation menu, expand **Manage Protection** > **Applications** > **Db2**.
2. Select a Db2 instance to back up all the data in that instance, or click the instance name to view the databases available for backing up. You can then select individual databases in the Db2 instance that you want to back up.

   You can back up an entire instance with all of its associated data, or back up one or more databases.

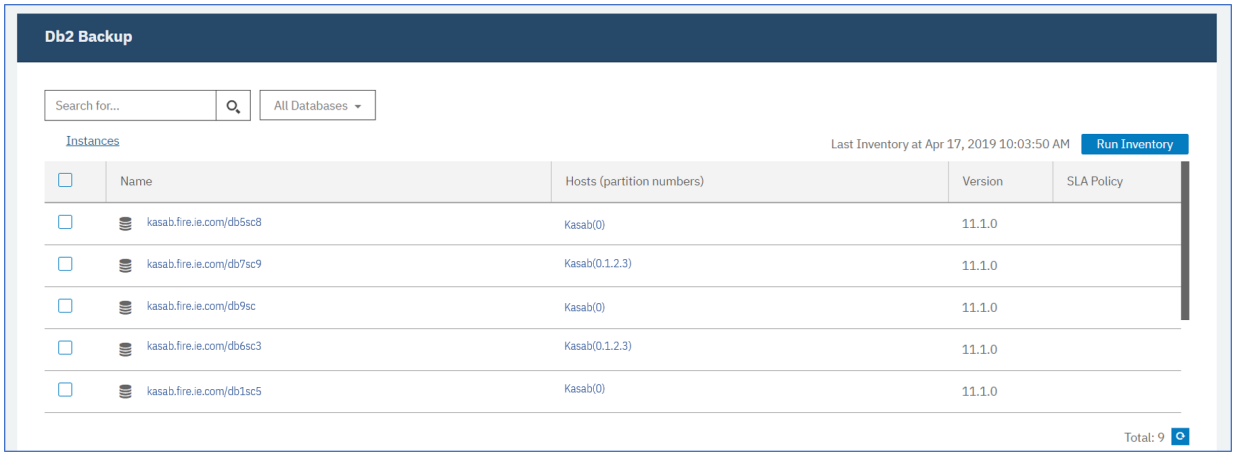


Figure 48. Db2 Backup pane showing instances

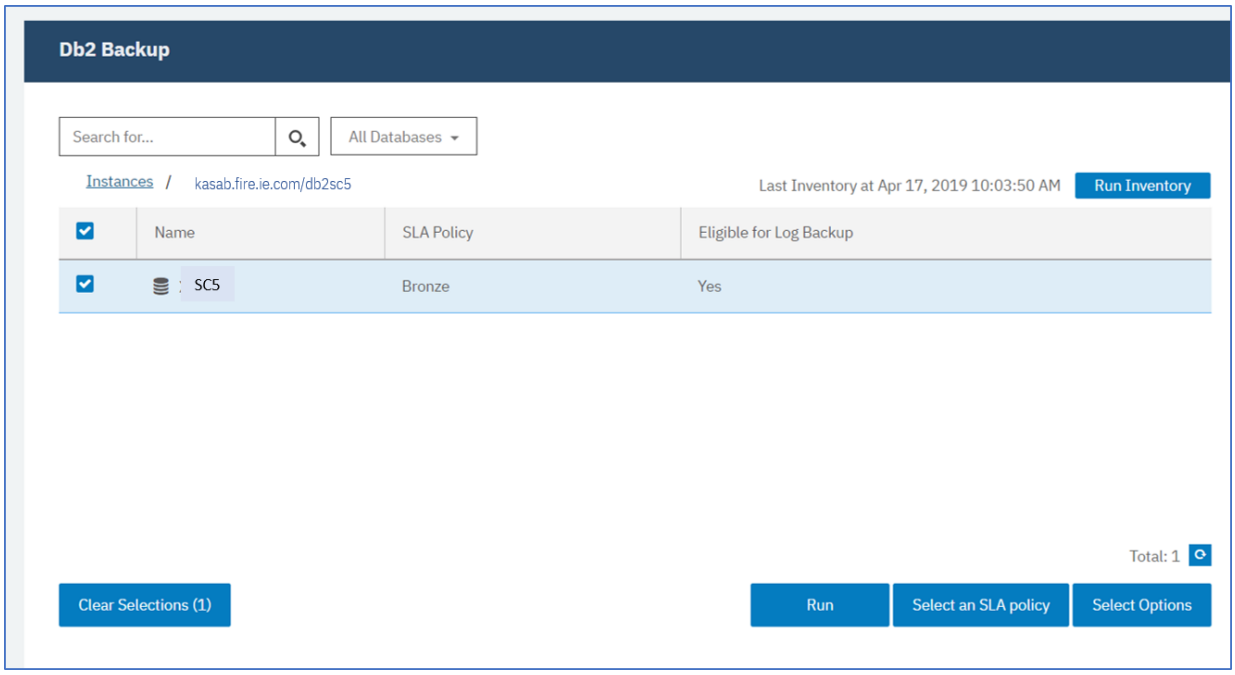


Figure 49. Db2 Backup pane showing databases in an instance

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

   The predefined Gold, Silver, and Bronze policies have different frequencies and retention rates. You can create a custom SLA policy or edit an existing policy by navigating to **Policy Overview** > **SLA Policies**.
4. Click **Select Options** to define options for your backup, such as enabling log backups for future recovery options, and specifying the parallel streams to reduce the time that is required to back up large databases. Save your changes.

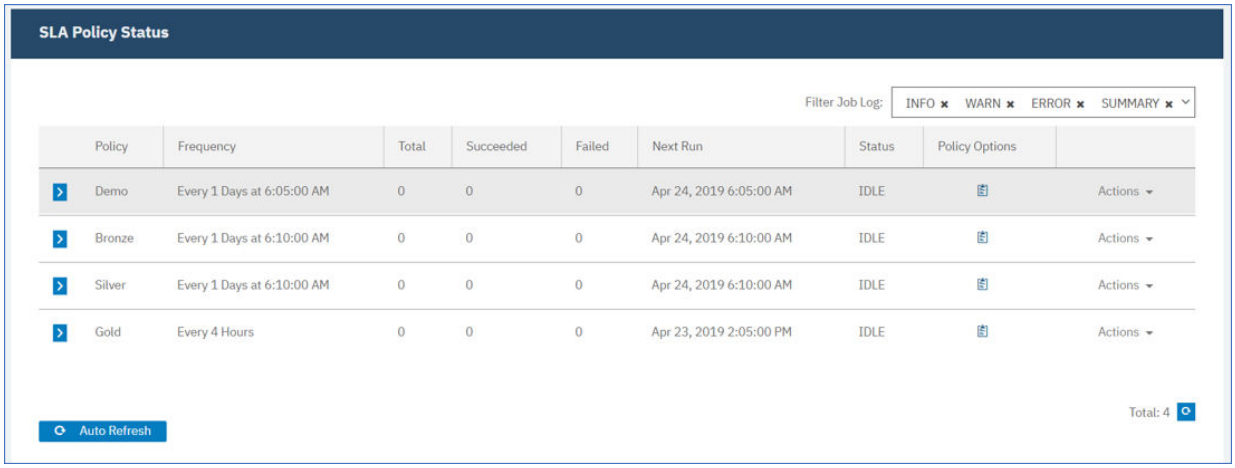


Figure 50. Backup options and SLA policies

5. Configure the SLA policy by clicking the icon in the **Policy Options** column of the **SLA Policy Status** table.

   To read about more SLA configuration options, see "Setting SLA configuration options for a backup job" on page 509.

6. To run the policy outside of the scheduled job, select the instance or database. Click **Actions** and select **Start**.

   The status changes to **Running** for your chosen SLA and you can follow the progress of the job in the job log shown.

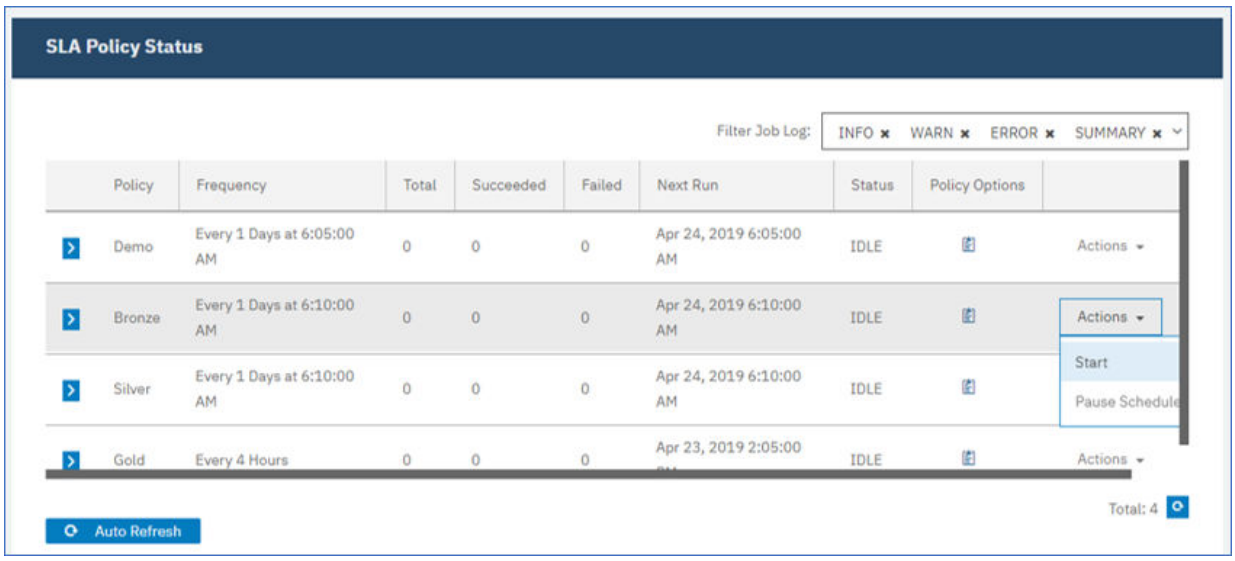


Figure 51. SLA policies

Tip: When the job for the selected SLA policy runs, all resources that are associated with that SLA policy are included in the backup operation. To back up only selected resources, you can run an on-demand job. An on-demand job runs the backup operation immediately.

- To run an on-demand backup job for a single resource, select the resource and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
- To run an on-demand backup job for one or more resources, click **Create job**, select **Ad hoc backup**, and follow the instructions in "Running an ad hoc backup job" on page 627.

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

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

Setting SLA configuration options for a backup job

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

Procedure

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

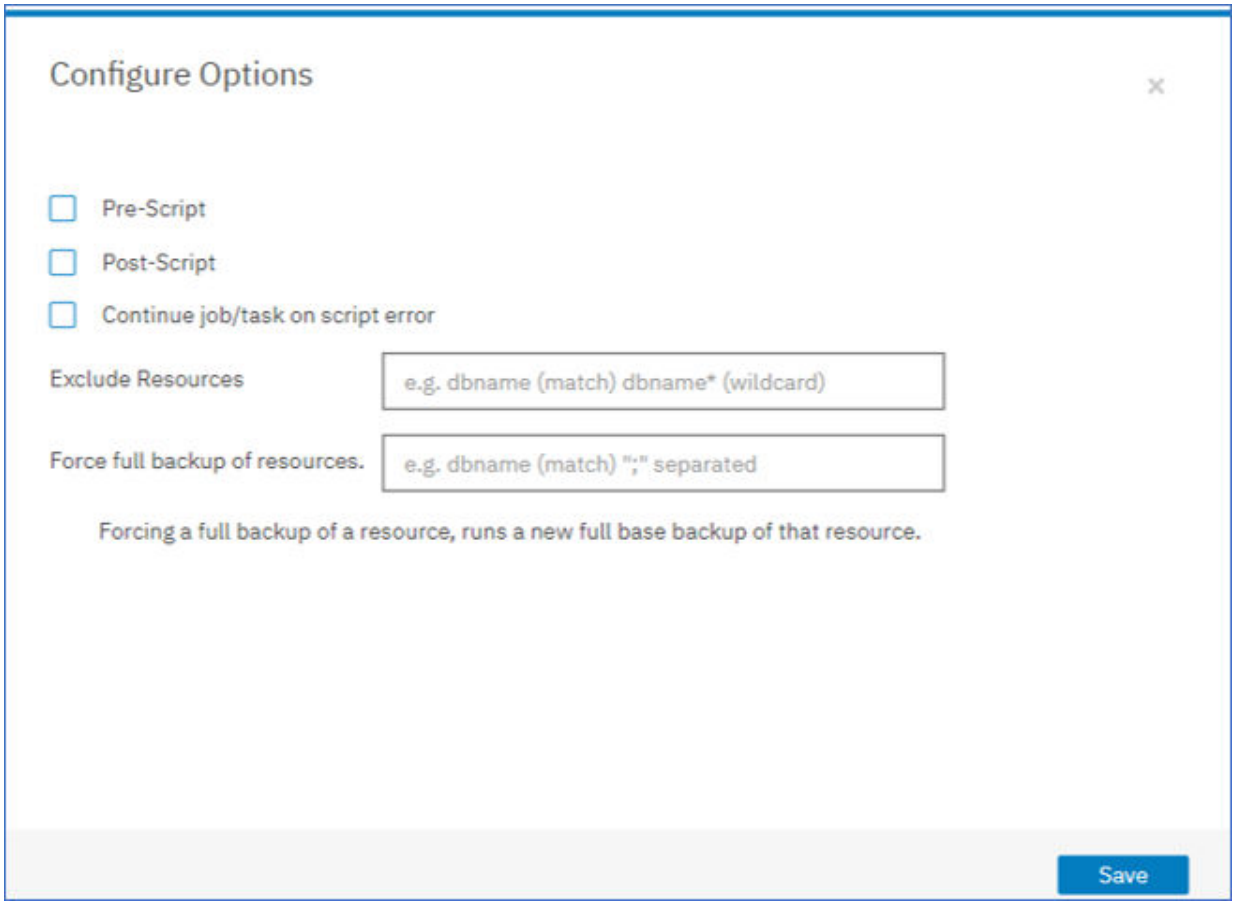


Figure 52. Specifying SLA configuration options

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

   Scripts and script servers are configured on the **System Configuration** > **Script** page. For more information about working with scripts, see Configuring scripts.
4. To continue running the job when the script that is associated with the job fails, select **Continue job/task on script error**.

If this option is selected, the backup or restore operation is reattempted and the script task status is reported as COMPLETED when the script completes processing with a nonzero return code. If this option is not selected, the backup or restore is not reattempted and the script task status is reported as FAILED.

5. To exclude resources from a backup job, specify the resources to exclude from the job. Enter an exact resource name in the **Exclude Resources** field. If you are unsure of a name, use wildcard asterisks that are specified before the pattern (**text*) or after the pattern (*text**). Multiple wildcards can be entered with standard alphanumeric characters and the following special characters: - _ and *. Separate entries with a semicolon.
6. To create a full new backup of a resource, enter the name of that resource in the **Force full backup of resources** field. Separate multiple resources with a semicolon.

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

Log backups

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

Ensure that you select the **Enable Log Backups** option to allow rollforward recovery when you set up a backup job or service level agreement (SLA) policy. When selected for the first time, you must run a backup job for the SLA policy to activate log archiving to IBM Spectrum Protect Plus on the database. This backup creates a separate volume on the vSnap repository, which is mounted persistently on the Db2 application server. The backup process updates either **LOGARCHMETH1** or **LOGARCHMETH2** parameters to point to that volume for log archiving purposes. The volume is kept mounted on the Db2 application server unless the **Enable Log Backup** option is cleared and a new backup job is run.

Restriction: In Db2 multi-partitioned environments, the **LOGARCHMETH** parameters across partitions must match.

When either **LOGARCHMETH1** or **LOGARCHMETH2** parameters are set with a value other than OFF, you can use archived logs for rollforward recovery. You can cancel log backup jobs at any time by clearing the **Enable Log Backups** option: go to **Manage Protection** > **Applications** > **Db2**, select the instance and click **Select Options**. This change takes effect after the next successful backup job completes, and the **LOGARCHMETH** parameter value is changed back to its original setting.

Important: IBM Spectrum Protect Plus can enable log backup jobs only when the **LOGARCHMETH1** parameter is set to LOGRETAIN or if one of the **LOGARCHMETH** parameters is set to OFF.

If the LOGARCHMETH1 parameter is set to LOGRETAIN.

IBM Spectrum Protect Plus changes the **LOGARCHMETH1** parameter value to enable log backups.

If either LOGARCHMETH1 or LOGARCHMETH2 parameters are set to OFF and the other is set to DISK, TSM, or VENDOR.

IBM Spectrum Protect Plus uses the **LOGARCHMETH** parameter that is set to off to enable log backups.

If both LOGARCHMETH parameters are set to DISK, TSM, or VENDOR.

This setting combination causes an error when IBM Spectrum Protect Plus attempts to enable log backups. To resolve the error, set one of the parameters to OFF, and run the backup job with the **Enable Log Backups** option selected.

Truncating archive log backups

IBM Spectrum Protect Plus automatically deletes older transactional logs after a successful database backup. This action ensures that the capacity of the log archive volume is not compromised by retention of older log files. These truncated log files are stored in the vSnap repository until the corresponding backup expires and is deleted. The retention of database backups is defined in the SLA policy that you

select. For more information about SLA policies, see “Defining a service level agreement backup job” on page 506.

IBM Spectrum Protect Plus does not manage the retention of other archived log locations.

For more information about Db2 settings, see IBM Db2 Welcome page.

Restoring Db2 data

To restore Db2 data from the vSnap repository, define a job that restores data from either the newest backup or an earlier backup copy. You can choose to restore data to the original instance or to an alternative instance on a different machine, and specify recovery options, and save the job.

Before you begin

Important: For all restore operations, Db2 must be at the same version level on the source and target hosts. In addition to that requirement, you must ensure that an instance with the same name as the instance that is being restored exists on each host. This requirement applies when the target instance has the same name, and when the names are different. In order for the restore operation to succeed, both instances must be provisioned, one with original name and the other with the new name.

If your Db2 environment includes partitioned databases, the data of all partitions is backed up during regular backup jobs. All instances are listed in the backup pane. Multi-partitioned instances are shown with partition numbers and host names.

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

- At least one Db2 backup job is set up and running successfully. For instructions about setting up a backup job, see “Backing up Db2 data” on page 504.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see Chapter 19, “Managing user access,” on page 643.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.
- Restore jobs can create data in the IBM Db2 log directory. In some cases, if more than one restore job is run, data will remain in the log directory from the previous job. As a result, the next attempt to restore a database to the original location fails unless the log directory is purged.

  For example, if the Db2 log directory is empty and a restore job runs with the options **Restore to original instance**, **Overwrite existing databases**, and **Recover until end of backup**, the restore job is successfully completed. If the job is followed by a second job with the options **Restore to original instance**, **Overwrite existing databases**, and **Recover until end of available logs**, this second restore attempt fails because the original restore job left data in the Db2 log directory.

Note: When you are restoring multi-partitioned databases to an alternative location, ensure that the target instance is configured with the same partition numbers as the original instance. All of those partitions must be on a single host. When you are restoring data to a new instance that is renamed, both instances instances required for the restore operation must be configured with the same number of partitions.

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

Procedure

1. In the navigation panel, expand **Manage Protection** > **Applications** > **Db2** and click **Create job** > **Restore**.

   The Rrestore wizard opens.

2. Optional: If you started the restore wizard from the **Jobs and Operations** page, click **Db2** as the source type and click **Next**.

   Tips:

   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

3. On the **Select source** page, click a Db2 instance to show the databases in that instance. Choose a database by clicking the plus icon for that database name. Click **Next** to continue.

4. In the **Source snapshot** page, choose the type of restore operation required.

   - **On-Demand: Snapshot**: creates a once-off restore operation from a database snapshot. The job is not set to recur.
   - **On-Demand: Point-in-Time**: creates a once-off restore operation from a point-in-time backup of the database. The job is not set to recur.
   - **Recurring**: creates a recurring job that runs on a schedule and repeats.

   Tip:

   For an **On-Demand: Snapshot** you can select no recovery or to recover until the end of the backup. For an **On-Demand: Point in Time** restore job you can select to recover until the end of the available logs, or recover until a specific point-in-time.

5. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape).

Option	Description
	• Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.

Option	Description
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

6. Choose a **restore method** appropriate for the destination chosen for the restore operation. Click **Next** to continue.
   - **Instant Access**: In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the volume from the vSnap repository. Use the data for custom recovery from the files in the mounted volume.
   - **Production**: In this mode, the Db2 application server first copies the files from the vSnap repository volume to the target host, which is either an alternative location or the original instance. That copied data is then used to start the database.
   - **Test**: In this mode, the agent creates a new database by using the data files directly from the vSnap repository.
   - Add a database name when you are restoring the database to a different location and you want to rename the database.

   Tip:

   Production is the only **restore method** that is available for restore operations to the original location. Any options not appropriate for the restore operation that you selected are not selectable.

   To restore data to the original instance, follow the instructions in Restoring to the original instance. To restore data to an alternative instance, follow the instructions in Restoring to an alternate instance.
7. Set the destination for the restore operation by choosing one of the following options. Click **Next** to continue.
   - **Restore to original instance**: this option restores data to the original server and original instance.
   - **Restore to alternate instance**: this option restores data to a different specified location, creating a copy of the data at that location.

   If you are restoring data to an alternative location, choose an instance in the **Instance** table before you click **Next**. The alternative instance must be on a different machine; unsuitable instances are not available for selection. For multi-partition databases, the target instance must have the same set of partitions on a single machine.
8. In the **Job Options** page, select the recovery, application, and advanced options for the restore operation you are defining.

   Tip:

   Recovery options are not available for instant access restore jobs.
   - **No Recovery**. This option skips any rollforward recovery after the restore operation. The database remains in a `Rollforward pending` state until you decide whether you want to run the rollforward operation manually.
   - **Recover until end of backup**. This option recovers the selected database to its state at the time the backup was created. The recovery process uses the log files that are included in the Db2 database backup.

- **Recover until end of available logs**. This option is available only if the logs are backed up in the Db2 backup job definition. IBM Spectrum Protect Plus uses the latest restore point. A temporary restore point for log backups is created automatically so that the Db2 database can be rolled forward to the end of the logs. This recovery option is not available if you selected a specific restore point from the list. This option is available only when you are running an on-demand point-in-time restore job which uses the latest backup.
- **Recover until specific point-in-time**. This option includes all the backup data up to a specific point-in-time. This option is available only if you enabled log backups in your Db2 backup job definition. Configure a point-in-time recovery by a specific date and time, for example, `Jan 1, 2019 12:18:00 AM`. IBM Spectrum Protect Plus finds the restore points directly before and after the selected point-in-time. During the recovery process, the older data backup volume and the newer log backup volume are mounted. If the point-in-time is after the last backup, a temporary restore point is created. This recovery option is not available if you selected a specific restore point from the list. This option is available only when you are running an on-demand point-in-time restore job that uses the newest backup.

Tip: To skip optional steps in the restore wizard, select **Skip optional steps** and click **Next**.

9. Optional: In the **Job Options** page, select the application options for the restore operation you are defining.

   Tip:

   Application options are not available for instant access restore jobs.

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

     Attention: Ensure that no other databases share the local database directory as the original database or that data is overwritten when this choice is selected.
   - **Maximum Parallel Streams per Database**. You can choose to run the restore operation of data in parallel streams. This option is useful if you are restoring a large database.
   - **Specify the size of the Db2 database memory set in KB**. Specify the memory, in KB, to be allocated for the database restore on the target machine. This value is used to modify the shared memory size of the Db2 database on the target server. To use the same shared memory size at both the source server and the target server, set the value to zero.
10. Optional: In the **Job Options** page, select the advanced options for the restore operation you are defining.
    - **Run cleanup immediately on job failure**. This option enables the automatic cleanup of backup data as part of a restore if recovery fails. This option is selected by default. Do not clear this option unless instructed by IBM® Support for troubleshooting purposes.
    - **Continue with restores of other selected databases even if one fails**. This option continues the restore operation if one database in the instance fails to be restored successfully. The process continues for all other databases that are being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.
    - **Mount point prefix**. For instant access restore operations, specify the prefix for the path where the mount point is to be directed.
11. Choose script options in the **Apply Scripts** page, and click **Next** to continue.
    - Select **Pre-Script** to select an uploaded script, and an application or script server where the pre-script runs. To select an application server where the script runs, clear the **Use Script Server** check box. Go to the **System Configuration** > **Script** page to configure scripts and script servers.
    - Select **Post-Script** to select an uploaded script and an application or script server where the post-script runs. To select an application server where the script runs, clear the **Use Script Server** check box. Go to the **System Configuration** > **Script** page to configure scripts and script servers.

- Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails. When this option is enabled and the prescript completes with a nonzero return code, the backup or restore job continues to run and the prescript task status returns `COMPLETED`. If a postscript completes with a nonzero return code, the postscript task status returns `COMPLETED`. When this option is not selected, the backup or restore job does not run, and the prescript or postscript task status returns with a `FAILED` status.

12. In the **Schedule** page, name the restore job and choose the frequency for the job to run. Schedule the start time, and click **Next** to continue.

    If the restore job you are specifying is an on-demand job, there is no option to enter a schedule. Specify a schedule only for recurrent restore jobs.
13. In the **Review** page, review your selections for the restore job. If all the details are correct for your restore job, click **Submit**, or click **Back** to make amendments.

Results

A few moments after you click **Submit**, the **onDemandRestore** record is added to the **Job Sessions** pane. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon  . All running jobs are viewable in the **Jobs and Operations Running Jobs** page.

To restore data to the original instance, follow the instructions in Restoring to the original instance. To restore data to an alternative instance, follow the instructions in Restoring to an alternate instance.

Restoring Db2 data to the original instance

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

Before you begin

If your Db2 environment includes partitioned databases, the data of all partitions is backed up during regular backup jobs. All instances are listed in the backup pane. Multi-partitioned instances are shown with partition numbers and host names.

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

- At least one Db2 backup job is set up and running successfully. For instructions about setting up a backup job, see “Backing up Db2 data” on page 504.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see Chapter 19, “Managing user access,” on page 643.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.
- Restore jobs can create data in the IBM Db2 log directory. In some cases, if more than one restore job is run, data will remain in the log directory from the previous job. As a result, the next attempt to restore a database to the original location fails unless the log directory is purged.

  For example, if the Db2 log directory is empty and a restore job runs with the options **Restore to original instance**, **Overwrite existing databases**, and **Recover until end of backup**, the restore job is successfully completed. If the job is followed by a second job with the options **Restore to original instance**, **Overwrite existing databases**, and **Recover until end of available logs**, this second restore attempt fails because the original restore job left data in the Db2 log directory.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Applications** > **Db2** and click **Create job** > **Restore**.

   The Rrestore wizard opens.

2. Optional: If you started the restore wizard from the **Jobs and Operations** page, click **Db2** as the source type and click **Next**.

   Tips:

   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

3. On the **Select source** page, click a Db2 instance to show the databases in that instance. Choose a database by clicking the plus icon for that database name. Click **Next** to continue.

4. In the **Source snapshot** page, choose the type of restore operation required.

   - **On-Demand: Snapshot**: creates a once-off restore operation from a database snapshot. The job is not set to recur.
   - **On-Demand: Point-in-Time**: creates a once-off restore operation from a point-in-time backup of the database. The job is not set to recur.
   - **Recurring**: creates a recurring job that runs on a schedule and repeats.

   Tip:

   For an **On-Demand: Snapshot** you can select no recovery or to recover until the end of the backup. For an **On-Demand: Point in Time** restore job you can select to recover until the end of the available logs, or recover until a specific point-in-time.

5. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape).

Option	Description
	• Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.

Option	Description
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

6. In the **Restore Method** page, choose **Production** for the restore operation.

   In **Production** mode, the Db2 application server first copies the files from the vSnap repository volume to the target host. That copied data is then used to start the database.

   Tip: Avoid entering a new database name when you are restoring a production operation to the original instance as it will not be implemented.
7. Set the destination for the restore operation to **Restore to original instance** to restore data to the original server. Click **Next** to continue.
8. Choose options as described in “Restoring Db2 data ” on page 511.
9. In the **Schedule** page, name the restore job and choose the frequency for the job to run. Schedule the start time, and click **Next** to continue.

   If the restore job you are specifying is an on-demand job, there is no option to enter a schedule. Specify a schedule only for recurrent restore jobs.
10. In the **Review** page, review your selections for the restore job. If all the details are correct for your restore job, click **Submit**, or click **Back** to make amendments.

Results

A few moments after you click **Submit**, the **onDemandRestore** record is added to the **Job Sessions** pane. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon  . All running jobs are viewable in the **Jobs and Operations Running Jobs** page.

Restoring Db2 databases to an alternative instance

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

Before you begin

Important: For all restore operations, Db2 must be at the same version level on the source and target hosts. In addition to that requirement, you must ensure that an instance with the same name as the instance that is being restored exists on each host. This requirement applies when the target instance has the same name, and when the names are different. In order for the restore operation to succeed, both instances must be provisioned, one with original name and the other with the new name.

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

- At least one Db2 backup job is set up and running successfully. For instructions about setting up a backup job, see “Backing up Db2 data” on page 504.

- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For more information about assigning roles, see Chapter 19, “Managing user access,” on page 643.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.
- Restore jobs can create data in the IBM Db2 log directory. In some cases, if more than one restore job is run, data will remain in the log directory from the previous job. As a result, the next attempt to restore a database to the original location fails unless the log directory is purged.

  For example, if the Db2 log directory is empty and a restore job runs with the options **Restore to original instance**, **Overwrite existing databases**, and **Recover until end of backup**, the restore job is successfully completed. If the job is followed by a second job with the options **Restore to original instance**, **Overwrite existing databases**, and **Recover until end of available logs**, this second restore attempt fails because the original restore job left data in the Db2 log directory.

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

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

Note: When you are restoring multi-partitioned databases to an alternative location, ensure that the target instance is configured with the same partition numbers as the original instance. All of those partitions must be on a single host. When you are restoring data to a new instance that is renamed, both instances instances required for the restore operation must be configured with the same number of partitions.

About this task

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

Procedure

1. In the navigation panel, expand **Manage Protection** > **Applications** > **Db2** and click **Create job** > **Restore**.

   The Rrestore wizard opens.
2. Optional: If you started the restore wizard from the **Jobs and Operations** page, click **Db2** as the source type and click **Next**.

   Tips:

   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
3. On the **Select source** page, click a Db2 instance to show the databases in that instance. Choose a database by clicking the plus icon for that database name. Click **Next** to continue.
4. In the **Source snapshot** page, choose the type of restore operation required.

- **On-Demand: Snapshot**: creates a once-off restore operation from a database snapshot. The job is not set to recur.
- **On-Demand: Point-in-Time**: creates a once-off restore operation from a point-in-time backup of the database. The job is not set to recur.
- **Recurring**: creates a recurring job that runs on a schedule and repeats.

Tip:

For an **On-Demand: Snapshot** you can select no recovery or to recover until the end of the backup. For an **On-Demand: Point in Time** restore job you can select to recover until the end of the available logs, or recover until a specific point-in-time.

5. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

6. Choose a **restore method** appropriate for the destination chosen for the restore operation. Click **Next** to continue.
   - **Production**: In this mode, the Db2 application server first copies the files from the vSnap repository volume to the target host, which is either an alternative location or the original instance. That copied data is then used to start the database.
   - **Test**: In this mode, the agent creates a new database by using the data files directly from the vSnap repository.

- **Instant Access**: In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the volume from the vSnap repository. Use the data for custom recovery from the files in the mounted volume.
- Add a database name when you are restoring the database to a different location and you want to rename the database.

7. Set the destination for the restore operation to **Restore to alternate instance** to restore data to a different location, which you can select from the list of eligible locations. Click **Next** to continue.

   When you are restoring to an alternative location, choose an instance in the **Instance** table before you click **Next**. Unsuitable target instances cannot be selected.
8. Choose options as described in "Restoring Db2 data " on page 511.
9. In the **Schedule** page, name the restore job and choose the frequency for the job to run. Schedule the start time, and click **Next** to continue.

   If the restore job you are specifying is an on-demand job, there is no option to enter a schedule. Specify a schedule only for recurrent restore jobs.
10. In the **Review** page, review your selections for the restore job. If all the details are correct for your restore job, click **Submit**, or click **Back** to make amendments.

Results

A few moments after you click **Submit**, the **onDemandRestore** record is added to the **Job Sessions** pane. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon . All running jobs are viewable in the **Jobs and Operations Running Jobs** page.

Exchange Server

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

Prerequisites for Exchange Server

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

For more information, see "Microsoft Exchange Server requirements" on page 89.

Virtualization support

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

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

Privileges

To help ensure that an Exchange agent can work in your IBM Spectrum Protect Plus environment, you must set up the appropriate privileges for the Exchange user account.

Role-based access control

You are required to register the Exchange Server with IBM Spectrum Protect Plus with an Exchange user who has local administrator privileges and the correct role-based access control (RBAC) permissions.

Also, for granular restore operations you are required to use an Exchange user who has local administrator privileges and the correct RBAC permissions.

To meet the minimum requirements for an Exchange user, complete the following steps:

1. Verify that the Exchange user is a member of a local Administrator group and has an active Exchange mailbox in the domain.

   By default, Windows adds the Exchange Organization Administrators group to other security groups, including the local Administrators group. For Exchange users who are not members of the Exchange Organization Management group, you must manually add the user account to the local Administrators group by taking one of the following actions:

   - On the computer of the domain member, click **Administrative tools** > **Computer Management** > **Local Users and Groups tool**.
   - On a domain controller computer that does not have a local Administrators group or Local Users and Groups tool, manually add the user account to the Administrators group in the domain: Click **Administrative tools** > **Active Directory Users and Computers tool**.
2. Set the role and scope.
   - Verify that the Exchange user has the correct RBAC permissions.

     You must assign the following management roles to each Exchange user who will complete mailbox restore operations:

     – Active Directory Permissions
     – ApplicationImpersonation
     – Databases
     – Disaster Recovery
     – Mailbox Import Export
     – Public Folders
     – View-Only Configuration
     – View-Only Recipients

     Place users who complete mailbox restore tasks into an Exchange Server role group that contains these roles.

     Exchange Server includes several built-in role groups. The Organization Management role group by default contains most, if not all, of the roles that are listed.

     Place users who must complete multiple mailbox restore tasks into the Organization Management role group (ensuring that the group contains all of the listed roles).

     Alternatively, you can place the user into another role group that you created or any other built-in role group that contains the roles that are listed. A user whose name is not in the Organization Management role group or subgroups might experience slower performance during restore operations.

     Important: You can manage Exchange role groups by using the Exchange Admin Center (EAC) or Exchange Powershell Cmdlets *only* if your user name is authorized by the security policy in your organization.
   - Management role scope

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

     – The Exchange Server that contains the required data
     – The recovery database that is created by IBM Spectrum Protect Plus
     – The database that contains the active mailbox
     – The database that contains the active mailbox of the user who completes the restore operation

Encrypting File System

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

Exchange certificates

Exchange digital certificates must be installed and configured for the mailbox browser to function during a granular restore operation. Ensure that the current Exchange certificates are installed and configured correctly in your environment.

Note: With Exchange 2016 and Exchange 2019, the Exchange Server is configured to use Transport Layer Security (TLS) by default. This TLS security encrypts communication between internal Exchange servers, and between Exchange services on the local server.

Adding an Exchange application server

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

About this task

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

Restriction: You can assign only one application server or file server per host. For example, if you register a host as a Microsoft Windows file system, you cannot register the same host as a Microsoft SQL Server or a Microsoft Exchange Server.

Procedure

To add an Exchange application server, complete the following steps:

1. In the navigation panel, expand **Manage Protection** > **Databases** > **Exchange**.
2. On the **Exchange** page, click **Manage Application Servers**, and then click **Add Application Server** to add the host system.
3. In the **Application Properties** form, enter the IP or host address.
4. Enter a user ID in the format of active directory domain and user account (domain\user), and the associated password.

   This user must have the correct Exchange roles and privileges. For more information about Exchange privileges, see “Privileges ” on page 523.
5. In the **Maximum concurrent databases** field, set the maximum number of databases per service level agreement (SLA) policy that can be backed up concurrently. The default is 10. Valid values are 1 - 99.

   This value might be higher or lower than the number of databases that are associated with an SLA policy. For example, if an SLA policy has 10 associated databases and this value is set to 2, a backup operation occurs for only 2 of the 10 databases at the same time. As each backup operation completes, a second backup operation starts until all databases are backed up. If an SLA policy has 5 associated databases and this value is set to 10, all 5 database backup operations occur at the same time.

   This option applies only to SLA policies that are associated with multiple databases. For SLA policies that are associated with only one database, this option provides no function.

   The maximum number of concurrent database backup operations is limited by your environment. Some things to consider are the vSnap server configuration, network bandwidth, and the physical disk configuration of your IBM Spectrum Protect Plus server.

For guidance about tuning your IBM Spectrum Protect Plus environment for best performance, see the IBM Spectrum Protect Plus Blueprints.

6. Click **Save**, and repeat the steps to add other Microsoft Exchange instances to IBM Spectrum Protect Plus.

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

What to do next

When you add your Exchange application server to IBM Spectrum Protect Plus, an inventory is automatically run on each instance. Databases must be detected to ensure that they can be backed up, and you can run a manual inventory at any time to detect updates. For instructions about running a manual inventory, see "Detecting Exchange databases by running an inventory" on page 526. For instructions about setting up Exchange database backup jobs, see "Defining a Service Level Agreement backup job" on page 527.

Detecting Exchange databases by running an inventory

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

Before you begin

Ensure that you added your Exchange instances to IBM Spectrum Protect Plus. For instructions about adding an Exchange instance, see "Adding an Exchange application server" on page 525.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Databases** > **Exchange**.
2. Click **Run Inventory**.

   When the inventory is running, the button label changes to **Inventory In Progress**. You can run an inventory on any available application server, but you can run only one inventory process at a time.
3. To monitor the inventory job, go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

   Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.
4. When the inventory job is complete, on the **Exchange Backup** pane, click an Exchange instance to open a view that shows the databases that are detected for that instance. If any databases are missing from the **Instances** list, check your Exchange application server and rerun the inventory.

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

Testing the Exchange connection

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

Procedure

1. In the navigation panel, expand **Manage Protection** > **Databases** > **Exchange**.
2. On the **Exchange** page, click **Manage Application Servers**.

   The Microsoft Exchange application servers that are available are shown.
3. Click **Actions** for the Microsoft Exchange application server that you want to test, and then click **Test**.

The test report shows you a list of the tests that ran and their status. Each test procedure includes a test of the physical host network configuration, a remote session test, and a test of Windows prerequisites such as user administrator privileges.

4. Click **OK** to close the test. Run the test again after you fix any issues.

Backing up Exchange databases

To protect Exchange databases, you can define a backup job that runs continuously to create incremental backups. You can also run on-demand backup jobs outside of the schedule.

Before you begin

Ensure that the application servers that contain the Exchange databases that you want to back up are registered with IBM Spectrum Protect Plus. For more information, see "Adding an Exchange application server" on page 525.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Databases** > **Exchange**.
2. On the **Exchange Backup** pane, click the Microsoft Exchange instance, and then select the database to back up.

   Each database is listed by instance or database name, the applied SLA policy, and the eligibility for log backup.
3. Click **Run**.

   The backup job begins, and you can view the details in **Jobs and Operations** > **Running Jobs**.

   Tip: The **Run** button is only enabled for a single database backup, and the database must have an SLA policy applied.

   To run an on-demand backup job for multiple databases that are associated with an SLA policy, click **Create job**, select **Ad hoc backup**, and follow the instructions in "Running an ad hoc backup job" on page 627.
4. To run backup jobs for multiple databases, select the databases in the Exchange backup pane, and click **Select an SLA Policy**.

   For more information about defining SLA policy backup jobs, and backup job options, see "Defining a Service Level Agreement backup job" on page 527.

Defining a Service Level Agreement backup job

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

About this task

IBM Spectrum Protect Plus supports single or multiple Exchange databases per Exchange backup job. Multiple database backup jobs run sequentially.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Databases** > **Exchange**.
2. Select an Exchange instance to back up all the data in that instance, or click an instance name, and then select individual databases that you want to back up.
3. Click **Select an SLA Policy** and choose an SLA Policy.

   Predefined choices are Gold, Silver, and Bronze, each with different frequencies and retention rates. Gold is the most frequent with the shortest retention rate. You can also create a custom SLA policy or edit an existing policy. For more information see "Creating an SLA policy for hypervisors, databases, and file systems" on page 324.

4. Click **Select Options** to define options for your backup, such as enabling log backups for future recovery options, and specifying the parallel streams to reduce the time that is taken to back up large databases. Save your changes.

   Tips for configuring options:

   Review the following tips to help you configure options for the backup job:

   - To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
   - If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
   - Options that are shown in yellow indicate that the option value has changed from the previously saved value.
   - To close the **Options** pane without saving changes, click **Select Options**.
5. Configure the SLA policy by clicking the icon in the **Policy Options** column of the **SLA Policy Status** table.

   For more information about SLA configuration options, see "Setting SLA configuration options for a backup job" on page 528.
6. To run the policy outside of the scheduled job, select the instance or database and then click **Actions** > **Start**.

   The status changes to **Running** for your chosen SLA. To pause the schedule, click **Actions** > **Pause Schedule**, and to cancel a job after it has started, click **Actions** > **Cancel**.

Setting SLA configuration options for a backup job

After you set up a service level agreement (SLA) for your backup job, you can choose to configure more options for that job. Extra SLA options include running scripts, excluding resources from the backup operation, and forcing a full base backup copy if required.

Procedure

1. In the **Policy Options** column of the **SLA Policy Status** table for the job that you are configuring, click the clipboard icon to specify additional configuration options.
2. To define a pre-script configuration, select **Pre-Script** and take one of the following actions:
   - To use a script server, select **Use Script Server** and choose an uploaded script from the **Script** or **Script Server** list.
   - To run a script on an application server, clear the **Use Script Server** check box, and choose an application server from the **Application Server** list.
3. To define a post-script configuration, select **Post-Script** and take one of the following actions:
   - To use a script server, select **Use Script Server** and choose an uploaded script from the **Script** or **Script Server** list.
   - To run a script on an application server, clear the **Use Script Server** check box, and choose an application server from the **Application Server** list.

   Scripts and script servers are configured on the **System Configuration** > **Script** page. For more information about working with scripts, see Configuring scripts.
4. Select **Continue job/task on script error** to continue running the job when the script that is associated with the job fails.

   If this option is selected, the backup or restore operation is attempted and the script task status is reported as COMPLETED when the script completes processing with a nonzero return code. If this option is not selected, the backup or restore is not attempted and the script task status is reported as FAILED.
5. Specify resources to exclude them from the backup job. Enter an exact resource name in the **Exclude Resources** field. If you are unsure of a name, use wildcard asterisks that are specified

before the pattern (**text*) or after the pattern (*text**). Multiple wildcards can be entered with standard alphanumeric characters and the following special characters: - _ and *. Separate entries with a semicolon.

6. If you want to create a full backup of a particular resource, enter the name of that resource in the **Force full backup of resources** field. Separate multiple resources with a semicolon.

   A full backup replaces the existing backup of that resource for one occurrence only. After that, the resource is backed up incrementally as before.

7. Click **Save**.

Backing up Exchange database logs

You can back up the database transaction logs for Exchange databases. Exchange log backups are scheduled by using Windows Task Scheduler. When log backups are available, you can run a rollforward data recovery during a restore operation to ensure that the data is recovered to the latest possible point in time.

About this task

When log backups are enabled, a Task Scheduler task is created on the Exchange server. The task runs a backup operation of your Exchange log files according to the SLA policy.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Databases** > **Exchange**.
2. Click the Exchange Server instance that you want to protect, and then select the databases whose logs you want to back up.

   Tip: The **Eligible for Log Backup** column shows the databases for which you can run log backups. If a database is registered as not eligible for log backup, a hover help explanation is provided.

3. Click **Select Options** and then select **Enable Log Backup**.

   If an on-demand job runs with the **Enable Log Backup** option enabled, log backup occurs. However, when the job runs again on a schedule, the option is disabled for that job run to prevent possible missing segments in the chain of backups.

4. For **Repeats**, enter the frequency of the log backups in **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.
5. Choose the **Start Time** and select the time for the log backups to begin, and then click **Save**.

Results

The database transaction logs are backed up to the vSnap server according to the selected frequency.

Restriction: The database logs are backed up on the preferred node only. Only one Exchange Server instance at a time can write log backups to the vSnap server.

Any log backup issues that occur are displayed in the Alert notifications in IBM Spectrum Protect Plus.

Backing up Exchange databases in a Database Availability Group

You can back up the mailbox databases in an Exchange Database Availability Group (DAG) and specify whether to use the active copy or a passive copy of the database for the backup. The Exchange servers in a DAG environment synchronize the data between active and passive copies for high availability.

About this task

By using the information from an inventory job, IBM Spectrum Protect Plus provides a DAG view that displays all of the databases in an Exchange DAG environment. Each database has an active copy on one server in the DAG, and one or more passive copies on the other servers. By default, scheduled backups

are taken from the server that the database is active on, but you can select a different server to back up a passive copy of the database.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Databases** > **Exchange**.
2. In the **Exchange Backup** pane, click the **View** menu and select **Database Availability Groups**.
3. Click the Exchange DAG that you want to view, and then select the databases to back up.
4. Click **Select Options**. In the **Backup preferred node** list, select the instance to run the backups on.
   With the **Backup preferred node** option, you can select a passive copy of the database for the backup.
5. Click **Select an SLA Policy** and then select an SLA policy from the list.
6. To create the job definition by using default options, click **Save**.
   The DAG databases are scheduled for backup jobs in accordance with the selected SLA policies and the preferred node choices.
7. To run the selected policy outside of the schedule, in the **SLA Policy Status** pane, click **Actions** > **Start**.

Incremental forever backup strategy

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

The incremental forever backup solution provides the following advantages:

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

The IBM Spectrum Protect Plus incremental forever process includes the following steps:

1. The first backup job creates a VSS snapshot of the Exchange application. As a result, the database files are in an application consistent state. The complete database files are copied to the vSnap location.
2. All subsequent backups create a VSS snapshot of the Exchange application. The database files are in an application consistent state. However, only the change blocks of the database files are copied to the vSnap location.
3. The backups are reconstructed at each point in time that a backup is performed, making it possible to recover the database from any single backup point.

Restoring Exchange databases

If data in an Exchange database is lost or corrupted, you can restore the data from a backup copy. Use the **Restore** wizard to set up a restore job schedule or an on-demand restore operation. You can define a job that restores data to the original instance or to an alternative instance, with different types of recovery options and configurations available.

Before you begin

Ensure that the following requirements are met:

- At least one Exchange backup job is defined and ran successfully. For instructions about defining a backup job, see “Defining a Service Level Agreement backup job” on page 527.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is defining the restore job. For more information about assigning roles, see Chapter 19, “Managing user access,” on page 643.

- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

Important: For granular restore operations, you must log on to the Exchange application server and use the Microsoft Management Console (MMC) GUI to complete mailbox batch restore and mailbox restore browser tasks.

Procedure

To restore data in an Exchange database, take one of the following actions:

- Restore a database to the original instance and location.
- Restore a database to the original instance with a different file location.
- Restore a database to an alternative instance.
- Restore mailbox data by using the granular restore function.
- Restore a database in a database availability group (DAG).

Restoring an Exchange database to the original instance

Restore an Exchange database to its original instance by using production mode or test mode. Choose between restoring the latest backup or an earlier Exchange database backup version.

Before you begin

Ensure that the following requirements are met:

- At least one Exchange backup job is defined and ran successfully.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is defining the restore job. For more information about assigning roles, see Chapter 19, “Managing user access,” on page 643.

About this task

When you restore a database to its original location in production mode, you cannot rename it. This restore option runs a full production restore operation, and existing data is overwritten at the target site.

Procedure

To define an Exchange restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **Exchange** > **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **Exchange**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
   a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
   b) Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.

c) Click **Next** to continue.

3. On the **Source snapshot** page, select the type of restore job that you want to create:

   On-demand: Snapshot
   Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

   On-demand: Point in Time
   Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

   Recurring
   Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Restore method** page, choose from the following options:

   - **Test**. In test mode, the agent creates a new recovery database by using the data files directly from the vSnap repository. This restore type might be used for testing purposes.

- **Production**. In production mode, the agent first restores the files from the vSnap volume back to primary storage and then creates the new database by using the restored files.

For Test restore only, in the **New Database Name** field, enter the new name for the restored database. The **New Database Name** field is also displayed when you choose Production restore, but this is for restoring to a new database location on the original instance. For detailed instructions on this task, see "Restoring an Exchange database to a new location on the original instance" on page 535.

6. On the **Set destination** page, select **Restore to original instance** and click **Next**.

Tip: In a Microsoft Exchange Server Database Availability Group (DAG) environment, you can restore to a recovery database on an active or passive member.

7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

Recovery Options
Choose from the following recovery options:

No Recovery
This option skips any rollforward recovery after the restore operation. The database remains in a `Rollforward pending` state until you decide whether you want to run the rollforward recovery manually.

Recover until end of backup
Restore the selected database to the state at the time the backup was created.

Recover until end of available logs
This option restores the database and applies all available logs (including logs newer than the backup that might exist on the application server) to recover the database up to the latest possible time. This option is available only if you selected **Enable Log Backup** in the backup job.

Recover until specific point in time
When log backups are enabled, this option restores the database and applies logs from the log backup volume to recover the database up to an intermediate, user-specified point in time. Choose the date and time by selecting from the **By Time** options.

Application Options
Set the application options:

Maximum Parallel Streams per Database
Set the maximum data stream from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might improve restore speed, but high-bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring an Exchange database to its original location by using its original database name.

Advanced Options
Set the advanced job definition options:

Run cleanup immediately on job failure
This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see Configuring scripts. Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:

- If you are running an on-demand job, click **Next**.
- If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.

10. On the **Review** page, review your restore job settings and click **Submit** to create the job.

    The restore job is created, and you can check on its status in **Jobs and Operations** > **Running Jobs**.

Restoring an Exchange database to a new location on the original instance

You can restore an Exchange database to its original instance, but to a new location on the application server. Choose between restoring the latest backup or an earlier Exchange database backup version.

About this task

When you restore a database to its original instance by using a production restore operation, you can restore the database to a new file location on the application server with a new name for the restored database. In production mode, the agent first restores the files from the vSnap volume back to primary storage and then creates a new database by using the restored files.

Procedure

To define an Exchange restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **Exchange** > **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **Exchange**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

2. On the **Select source** page, take the following actions:

   a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.

   b) Click the plus icon next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

   The selected sources are added to the restore list next to the database list. To remove an item from the list, click the minus icon next to the item.

   c) Click **Next** to continue.

3. On the **Source snapshot** page, select the type of restore job that you want to create:

   On-demand: Snapshot
   Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

   On-demand: Point in Time
   Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

   Recurring
   Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.

Option	Description
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.

Option	Description
	Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. In the **Restore Method** page, click the **Production** restore option.

   Tip: It is mandatory to select Production mode for this restore operation.

   a) In the **Name** field, expand the database name to see the path information for the existing database on the application server.

   b) In the **New Database Name** field, enter the new name for the restored database.

   c) In the **Destination Path** field, enter the new directory location for the database file on the server, including the .edb name, and the logs location.

   Warning: The destination directories that you enter in the **Destination Path** field must already exist on the application host. If not, then create the necessary directories on the server before you complete the restore operation.

   For example, for a database that is named Database_A, enter C:\<new_destination_path>\Database_A.edb, and for the location of the logs, enter C:\<new_logs_path>.

6. On the **Set destination** page, select **Restore to original instance** and click **Next**.

   Tip: In a Microsoft Exchange Server Database Availability Group (DAG) environment, you can restore to a recovery database on an active or passive member.

7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

   Recovery Options
   Choose from the following recovery options:

   No Recovery
   This option skips any rollforward recovery after the restore operation. The database remains in a `Rollforward pending` state until you decide whether you want to run the rollforward recovery manually.

   Recover until end of backup
   Restore the selected database to the state at the time the backup was created.

   Recover until end of available logs
   This option restores the database and applies all available logs (including logs newer than the backup that might exist on the application server) to recover the database up to the latest possible time. This option is available only if you selected **Enable Log Backup** in the backup job.

   Recover until specific point in time
   When log backups are enabled, this option restores the database and applies logs from the log backup volume to recover the database up to an intermediate, user-specified point in time. Choose the date and time by selecting from the **By Time** options.

   Application Options
   Set the application options:

   Maximum Parallel Streams per Database
   Set the maximum data stream from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might improve restore speed, but high-bandwidth consumption might affect overall system performance.

   This option is applicable only when you are restoring an Exchange database to its original location by using its original database name.

   Advanced Options
   Set the advanced job definition options:

   Run cleanup immediately on job failure
   This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.
8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see Configuring scripts. Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:
   - If you are running an on-demand job, click **Next**.
   - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.

    The restore job is created, and you can check on its status in **Jobs and Operations** > **Running Jobs**.

Restoring an Exchange database to an alternative instance

You can select a Microsoft Exchange database backup and restore it to an Exchange Server instance on an alternative host. You can restore the database in production mode or test mode to the alternative instance.

Before you begin

Ensure that the following requirements are met:

- Enough disk space and allocated dedicated volumes are available for the copying of files.
- The file system structure on the source server is the same as the file system structure on the target server. This file system structure includes table spaces, online logs, and the local database directory.

Procedure

1. In the navigation panel, click **Manage Protection** > **Databases** > **Exchange** > **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **Exchange**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
   a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
   b) Click the plus icon next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

      The selected sources are added to the restore list next to the database list. To remove an item from the list, click the minus icon next to the item.
   c) Click **Next** to continue.
3. On the **Source snapshot** page, select the type of restore job that you want to create:

   On-demand: Snapshot
   Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

   On-demand: Point in Time
   Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

   Recurring
   Creates a repeating point-in-time restore job that runs on a schedule.
4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order:

Option	Description
	Backup Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations:

Option	Description
	Primary The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Restore method** page, choose from the following options:
   - **Test**. In test mode, the agent creates a new recovery database by using the data files directly from the vSnap repository. This restore type might be used for testing purposes.
   - **Production**. In production mode, the agent first restores the files from the vSnap volume back to primary storage and then creates the new database by using the restored files.

   a) In the **New Database Name** field, enter a new database name.

   b) (Production restore only) Expand the database name to see the source and destination path information. In the **Destination Path** field, enter the directory location of the Exchange database file on the alternative host, including the .edb name, and the logs location.

   Warning: The destination directories that you enter in the **Destination Path** field must already exist on the alternative host. If not, then create the necessary directories on the alternative host before you complete the restore operation.

   For example, for a database that is named Database_A, enter C:\<new_destination_path>\Database_A.edb, and for the location of the logs , enter c:\<new_logs_path>.
6. On the **Set destination page**, choose **Restore to alternate instance**, select the target instance that you want to restore the database to and then click **Next**.
7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

   Recovery Options
   Choose from the following recovery options:

   No Recovery
   This option skips any rollforward recovery after the restore operation. The database remains in a `Rollforward pending` state until you decide whether you want to run the rollforward recovery manually.

   Recover until end of backup
   Restore the selected database to the state at the time the backup was created.

Recover until end of available logs
This option restores the database and applies all available logs (including logs newer than the backup that might exist on the application server) to recover the database up to the latest possible time. This option is available only if you selected **Enable Log Backup** in the backup job.

Recover until specific point in time
When log backups are enabled, this option restores the database and applies logs from the log backup volume to recover the database up to an intermediate, user-specified point in time. Choose the date and time by selecting from the **By Time** options.

Application Options
Set the application options:

Maximum Parallel Streams per Database
Set the maximum data stream from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might improve restore speed, but high-bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring an Exchange database to its original location by using its original database name.

Advanced Options
Set the advanced job definition options:

Run cleanup immediately on job failure
This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see Configuring scripts. Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:
   - If you are running an on-demand job, click **Next**.
   - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.
    The restore job is created, and you can check on its status in **Jobs and Operations** > **Running Jobs**.

Restoring individual mailbox items by using a granular restore operation

You can restore Exchange individual mailbox items by using a granular restore operation and the IBM Spectrum Protect Plus Microsoft Management Console (MMC) GUI.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations. If RBAC permissions were not assigned, you might encounter configuration errors in the IBM Spectrum Protect Plus MMC GUI for each missing role.

Tip:

If you encounter role-based configuration errors in the IBM Spectrum Protect Plus MMC GUI, you can set the required permissions manually to resolve the errors (see “Privileges ” on page 523), or you can run the IBM Spectrum Protect Plus configuration wizard to automatically configure permissions (see step “15” on page 546).

About this task

To start a granular restore operation, complete preparatory steps in the IBM Spectrum Protect Plus GUI, and then log in to the Exchange application server. Then, use the IBM Spectrum Protect Plus MMC GUI to restore user mailbox data from the recovery database that is created by the granular restore operation. A granular restore operation can be used to complete the following tasks:

- You can restore selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode `.pst` file.
- You can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
- You can restore an archive mailbox or a part of the mailbox, for example, a specific folder.
- You can restore archive mailbox messages to a mailbox that is on the Exchange Server, to an archive mailbox, or to an Exchange Server `.pst` file.

Procedure

1. In the navigation panel, click **Manage Protection** > **Databases** > **Exchange** > **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **Exchange**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

2. On the **Source select** page, complete the following steps:

   a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.

   b) Click the plus icon  next to the database that you want to use as the source of the restore operation.

   Tip: You must select only one database for a granular restore operation. If you select multiple databases, the granular restore option will not be available on the **Restore method** page.

   The selected source is added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.

   c) Click **Next** to continue.

3. On the **Source snapshot** page, select the type of restore job that you want to create:

   On-demand: Snapshot
   Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

   On-demand: Point in Time
   Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

   Recurring
   Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.

Option	Description
	Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Restore method** page, click **Granular Restore**.

   The recovery database name is displayed in the **New Database Name** field. The name consists of the existing database name with the suffix _RDB.

6. On the **Set destination** page, select **Restore to original instance** and click **Next**.

   Tip: In a Microsoft Exchange Server Database Availability Group (DAG) environment, you can restore to a recovery database on an active or passive member.

7. Optional: In the **Job Options** page, **Recover until end of backup** and **Run cleanup immediately on job failure** are selected by default. Click **Next** to continue.

   Restriction: Do not clear the **Run cleanup immediately on job failure** option unless instructed by IBM Support for troubleshooting purposes.

8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see Configuring scripts. Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:
   - If you are running an on-demand job, click **Next**.
   - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.

The restore job is created, and you can check on its status in **Jobs and Operations** > **Running Jobs**.

11. In the navigation panel, click **Jobs and Operations** > **Active Resources** > **Databases** to view the recovery database and mount point details.

    Tip: Click the icon to display an information message that describes the next steps for completing the granular restore task.

12. Connect to the Exchange application server instance by using Remote Desktop Connection (RDC) or Virtual Network Computing (VNC) if connecting remotely, or by logging on to the Exchange Server machine locally.

    The granular restore operation automatically installs and starts the IBM Spectrum Protect Plus MMC GUI on the application server. If the MMC GUI fails to start, start it manually by using the path that is provided in the **Active Resources** information message.

13. In the IBM Spectrum Protect Plus MMC GUI, click the **Protect and Recover Data** node, and select **Exchange Server**.
14. On the **Recover** tab for the Exchange Server instance, click **View** > **Mailbox Restore Browser** to view the mailbox from the recovery database.
15. Optional: Run the IBM Spectrum Protect Plus configuration wizard:
    a) In the navigation panel, click **Dashboard** > **Manage** > **Configuration** > **Wizards** > **IBM Spectrum Protect Plus Configuration**.
    b) In the **Actions** pane, click **Start**.

       The configuration wizard runs the requirements check.

    c) When the requirements checks have run, click the **Warnings** link next to **User Roles Check**.
    d) On the message dialog box, to add any missing roles, click **Yes**.
    e) On the configuration wizard, click **Next**, and then click **Finish**.
16. In the **Mailbox Restore Browser > Source** tree, click the mailbox that contains the items you want to restore, which enables you to browse the individual folders and messages.

    Choose from the following actions to select the folder or message to restore.

Table 103. Previewing and filtering mailbox items

Task	Action
Preview mailbox items	a. Select a mailbox item, such as **Inbox**, to display its contents in the preview pane. b. Click an individual item in the preview pane, such as an email message, to view the message text and details. c. If an item contains an attachment, click the attachment icon to preview its contents.

Table 103. Previewing and filtering mailbox items (continued)	
Task	**Action**
Filter mailbox items	Use the filter options to narrow the list of folders and messages to restore: a. Click **Show Filter Options** and **Add Row**. b. Click the down arrow in the **Column Name** field and select an item to filter. You can filter by folder name, subject text, and other options. **Restriction:** You can filter public mailbox folders only by the **Folder Name** column. When you select **All Content**, the mailbox items are filtered by attachment name, sender, subject, and message body. c. In the **Operator** field, select an operator: Contains. d. In the **Value** field, specify a filter value. e. To specify additional filtering criteria, click **Add Row**. f. Click **Apply Filter** to filter the messages and folders.

17. When you have selected the mailbox item to restore, in the **Actions** pane, click the restore task that you want to run. Choose from the following options:
    - **Restore Folder to Original Mailbox**
    - **Restore Messages to Original Mailbox**
    - **Save Mail Message Content**

    Tip: If you click **Save Mail Message Content**, a Windows Save File window is displayed. Specify the location and message name and click **Save**.

    When you choose the restore option, the **Restore Progress** window opens and shows the progress of the restore operation, and the mailbox item is restored.
18. To restore a mailbox item to another mailbox or `.pst` file, complete the following steps.

    Note: You can also restore a complete mailbox to another mailbox or `.pst` file.

    Choose from the actions in the following table:

Table 104. Restoring a mailbox item to another mailbox or `.pst` file	
Task	**Action**
Restore a mailbox item (or a mailbox) to a different mailbox	a. On the **Actions** pane, click **Open Exchange Mailbox**. b. Enter the alias of the mailbox to identify it as the restore destination. c. Drag the source mailbox item (or mailbox) to the destination mailbox on the results pane. **Restriction:** You cannot drag mail items or subfolders in the `Recoverable Items` folder to a destination mailbox.

Table 104. Restoring a mailbox item to another mailbox or `.pst` file (continued)	
Task	**Action**
Restore a mailbox item (or mailbox) to an Outlook personal folders (`.pst`) file	a. On the **Actions** pane, click **Open non-Unicode PST File**. b. When the **Open File** window opens, select an existing `.pst` file or create a `.pst` file. c. Drag the source mailbox item (or mailbox) to the destination `.pst` file on the results pane. **Restriction:** You can use the **Mailbox Restore Browser** view only with non-Unicode `.pst` files.
Restore a Public Folder	Select this action to restore a public folder to an existing online public folder mailbox. You can filter the mailbox and restore a specific public folder to an existing online public folder. In the **Folder to be restored** field, enter the name of the public folder that you want to restore. • To restore a subfolder in a parent folder, specify the full folder path in this format: *parent_folder_name*/*sub_folder_name*. • To restore all subfolders in a parent folder, use *parent_folder_name*/*. • If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\). You can also restore all or part of a public folder to a different public folder mailbox than the original mailbox. In the **Target public folder mailbox** field, specify the destination public folder mailbox that you want to restore to.

19. In the **Actions** pane, click **Close Exchange Mailbox** or **Close PST File** to close the destination mailbox or `.pst` file.

    Tip: You can enable the Microsoft Management Console to gather diagnostic information to assist in problem determination related to restore operations. The process gathers configuration files, trace files, and overall diagnostics of the MMC GUI. For more information, see the following technote: Enabling diagnostic information in the IBM Spectrum Protect Plus MMC GUI(http://www.ibm.com/support/docview.wss?uid=ibm10882270).
20. When the restore operation for the individual items is finished, return to IBM Spectrum Protect Plus and click **Jobs and Operations** > **Active Resources** > **Databases**.
21. Select the resource, and then click **Cancel File Restore** to end the granular restore process.

Restoring mailboxes by using a granular restore operation

You can restore Exchange mailboxes by using a granular restore operation and the IBM Spectrum Protect Plus Microsoft Management Console (MMC) GUI.

Before you begin

You must have role-based access control (RBAC) permissions to complete individual mailbox restore operations. If RBAC permissions were not assigned, you might encounter configuration errors in the IBM Spectrum Protect Plus MMC GUI for each missing role.

Tip:

If you encounter role-based configuration errors in the IBM Spectrum Protect Plus MMC GUI, you can set the required permissions manually to resolve the errors (see “Privileges ” on page 523), or you can run the IBM Spectrum Protect Plus configuration wizard to automatically configure permissions (see step “15” on page 552).

About this task

To start a granular restore operation, complete preparatory steps in the IBM Spectrum Protect Plus GUI, and then log in to the Exchange application server. Then use the IBM Spectrum Protect Plus MMC GUI to restore user mailbox data from the recovery database that is created by the granular restore operation. A granular restore operation can be used to complete the following tasks:

- You can restore an entire mailbox or selected mailbox items to the original mailbox, another online mailbox on the same server, or to a Unicode `.pst` file.
- You can restore a public folder mailbox database, a public folder mailbox, or only a part of the mailbox, for example, a specific public folder.
- You can restore an archive mailbox or a part of the mailbox, for example, a specific folder.
- You can restore archive mailbox messages to a mailbox that is on the Exchange Server, to an archive mailbox, or to an Exchange Server `.pst` file.

Procedure

1. In the navigation panel, click **Manage Protection** > **Databases** > **Exchange** > **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **Exchange**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Source select** page, complete the following steps:
   a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.

   b) Click the plus icon  next to the database that you want to use as the source of the restore operation.

      Tip: You must select only one database for a granular restore operation. If you select multiple databases, the granular restore option will not be available on the **Restore method** page.

      The selected source is added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.

   c) Click **Next** to continue.
3. On the **Source snapshot** page, select the type of restore job that you want to create:

   On-demand: Snapshot
   Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

   On-demand: Point in Time
   Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

   Recurring
   Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane.

Option	Description
	Repository server The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Restore method** page, click **Granular Restore**.

   The recovery database name is displayed in the **New Database Name** field. The name consists of the existing database name with the suffix _RDB.

6. On the **Set destination** page, select **Restore to original instance** and click **Next**.

   Tip: In a Microsoft Exchange Server Database Availability Group (DAG) environment, you can restore to a recovery database on an active or passive member.

7. Optional: In the **Job Options** page, **Recover until end of backup** and **Run cleanup immediately on job failure** are selected by default. Click **Next** to continue.

   Restriction: Do not clear the **Run cleanup immediately on job failure** option unless instructed by IBM Support for troubleshooting purposes.

8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see Configuring scripts. Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:

- If you are running an on-demand job, click **Next**.
- If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.

10. On the **Review** page, review your restore job settings and click **Submit** to create the job.

    The restore job is created, and you can check on its status in **Jobs and Operations** > **Running Jobs**.

11. In the navigation panel, click **Jobs and Operations** > **Active Resources** > **Databases** to view the recovery database and mount point details.

    Tip: Click the icon to display an information message that describes the next steps for completing the granular restore task.

12. Connect to the Exchange application server instance by using Remote Desktop Connection (RDC) or Virtual Network Computing (VNC) if connecting remotely, or by logging on to the Exchange Server machine locally.

    The granular restore operation automatically installs and starts the IBM Spectrum Protect Plus MMC GUI on the application server. If the MMC GUI fails to start, start it manually by using the path that is provided in the **Active Resources** information message.

13. In the IBM Spectrum Protect Plus MMC GUI, click the **Protect and Recover Data** node, and select **Exchange Server**.
14. On the **Recover** tab for the Exchange Server instance, select **View** > **Mailbox Restore**.

    A list of user mailboxes from all databases that are included in the backup is displayed.

15. Optional: Run the IBM Spectrum Protect Plus configuration wizard:
    a) In the navigation panel, click **Dashboard** > **Manage** > **Configuration** > **Wizards** > **IBM Spectrum Protect Plus Configuration**.
    b) In the **Actions** pane, click **Start**.

       The configuration wizard runs the requirements check.

    c) When the requirements checks have run, click the **Warnings** link next to **User Roles Check**.
    d) On the message dialog box, to add any missing roles, click **Yes**.
    e) On the configuration wizard, click **Next**, and then click **Finish**.
16. Select one or more mailboxes from the recovery database to restore. Mailboxes are listed by Mailbox Name, Alias, Server, Database, and Mailbox Type.

    You can restore only user mailboxes that are located in the recovery database.

    Tip: Mailboxes from other databases are shown in this view for informational purposes only. If the mailbox that you want to restore is not in the recovery database, use this view to determine which Exchange database the user mailbox was assigned to. You can then run the granular restore task again for that database.

17. To complete the restore operation, in the **Actions** pane, click one of the following restore options.

Table 105. Restore options

Option	Action
Restore Mail to Original Location	Restore mail items to their location at the time of the backup operation.
Restore Mail to Alternate Location	Restore the mail items to a different mailbox. • On the **Alternate Mailbox Options** window, enter the **Mailbox alias** name. **Tip:** If deleted mail items or tasks are flagged in the `Recoverable Items` folder of a mailbox, the items are restored with the flag attribute to the **Flagged Items and Tasks** view in the target mailbox.

Table 105. Restore options (continued)

Option	Action
Restore Mail to non-Unicode PST file **Restriction:** • This option is available only for Exchange Server 2013. • Each folder can contain a maximum of 16,383 mail items.	Restore mail items to a non-Unicode personal folders (`.pst`) file. When you restore mail items to a `.pst` file with one selected mailbox, you are prompted for a file name. When you restore mail items to a `.pst` file with more than one selected mailbox, you are prompted for a directory location. Each mailbox is restored to a separate `.pst` file that reflects the name of the mailbox at the specified directory. If the `.pst` file exists, the file is used. Otherwise, the file is created.
Restore Mail to Unicode PST file	Restore mail items to a Unicode `.pst` file. When you restore mail items to a `.pst` file with one selected mailbox, you are prompted for a file name. When you restore mail items to a `.pst` file with more than one selected mailbox, you are prompted for a directory location. **Tip:** You can enter a standard path name (for example, `c:\PST\mailbox.pst`) or a UNC path (for example, `\\server\c$\PST\mailbox.pst`). When you enter a standard path, the path is converted to a UNC path. If the UNC is a non-default UNC path, enter the UNC path directly. Each mailbox is restored to a separate `.pst` file that reflects the name of the mailbox at the specified directory. If the `.pst` file exists, the file is used. Otherwise, the file is created.
Restore Public Folder Mailbox	Restore a public folder mailbox to an online public folder mailbox. In the **Folder to be restored** field, enter the name of the public folder that you want to restore: • To restore a subfolder in a parent folder, specify the full folder path in this format: *parent_folder_name/sub_folder_name*. • To restore all subfolders in a parent folder, use *parent_folder_name/**. • If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\). You can also restore all or part of a public folder mailbox to a different public folder mailbox than the original mailbox. In the **Target public folder mailbox** field, specify the destination public folder mailbox.

Table 105. Restore options (continued)	
Option	**Action**
Restore Mail to Archive Mailbox	This action applies to a primary mailbox or an archive mailbox. Select this action to restore all or part of either type of mailbox to the original archive mailbox or to an alternative archive mailbox. You can filter the archive mailbox and restore a specific mailbox folder. In the **Folder to be restored** field, enter the name of the folder in the archive mailbox that you want to restore. • To restore a subfolder in a parent folder, specify the full folder path in this format: *parent_folder_name*/*sub_folder_name*. • To restore all subfolders in a parent folder, use *parent_folder_name*/*. • If the full folder path includes spaces, enclose the folder path in double quotation marks, and do not append a backslash character (\). In the **Target archive mailbox** field, specify the archive mailbox destination.
Exclude recoverable mail items while restoring the mailbox	Apply this action if you are restoring an online, public folder, or archive mailbox to an original mailbox, alternative mailbox, or to a Unicode `.pst` file. Specify a value of **Yes** to exclude the mail items in the `Recoverable Items` folder in mailbox restore operations. **No** is the default value.

Tip: You can enable the Microsoft Management Console to gather diagnostic information to assist in problem determination related to restore operations. The process gathers configuration files, trace files, and overall diagnostics of the MMC GUI. For more information, see the following technote: Enabling diagnostic information in the IBM Spectrum Protect Plus MMC GUI(http://www.ibm.com/support/docview.wss?uid=ibm10882270).

18. When the mailbox restore operation is finished, return to IBM Spectrum Protect Plus and click **Jobs and Operation** > **Active Resources** > **Databases**.
19. Select the resource, and then click **Cancel File Restore** to end the granular restore process.

Restoring Database Availability Group backups

With IBM Spectrum Protect Plus, you can restore a Microsoft Exchange Server Database Availability Group (DAG) backup to the original instance or to an alternative instance.

About this task

If you select the production restore method for the restore job, you must restore a replicated database copy to an active database copy. If you selected a passive database copy as the preferred target of backup operations, IBM Spectrum Protect Plus attempts to restore the database to this passive copy by default. The restore operation fails. In this situation, you can choose to restore the database to an alternative instance, and then select the active database copy.

If you select another restore method, the target location is a stand-alone recovery database that can be used on an active or passive DAG member.

Procedure

To define an Exchange restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **Exchange** > **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **Exchange**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

2. In the **Source select** page, complete the following steps:

   a) Click the **View** menu and select **Database Availability Groups**.

   b) In the **Availability Groups** list, click an Exchange instance to see the list of restore points for that instance and select the backup versions that you want to restore. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.

   c) Click the add to restore list icon next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

   The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the icon next to the item.

   d) Click **Next** to continue.

3. On the **Source snapshot** page, select the type of restore job that you want to create:

   On-demand: Snapshot
   Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

   On-demand: Point in Time
   Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

   Recurring
   Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server.

Option	Description
	Replication Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations:

Option	Description
	Primary The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. In the **Restore method** page, choose from the following options:
   - **Test**. Choose this option to restore the data from the vSnap repository directly. This restore type might be used for testing purposes.
   - **Production**. Choose this option to restore the full database with a full-copy data restore operation. This restore operation is for permanent use of the restored database.

   Click **Next** to continue.
6. In the **Set destination** page, specify where you want to restore the database and click **Next**.

   Restore to original instance
   Select this option to restore the database to the original server.

   Restore to alternate instance
   Select this option to restore the database to a local destination that is different from the original server, then select the alternative location from the list of available servers.

   Attention: When you choose the destination, you must select an active node as the destination; otherwise, the restore operation fails.
7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

   Recovery Options
   Choose from the following recovery options:

   No Recovery
   This option skips any rollforward recovery after the restore operation. The database remains in a `Rollforward pending` state until you decide whether you want to run the rollforward recovery manually.

   Recover until end of backup
   Restore the selected database to the state at the time the backup was created.

   Recover until end of available logs
   This option restores the database and applies all available logs (including logs newer than the backup that might exist on the application server) to recover the database up to the latest

possible time. This option is available only if you selected **Enable Log Backup** in the backup job.

Recover until specific point in time
When log backups are enabled, this option restores the database and applies logs from the log backup volume to recover the database up to an intermediate, user-specified point in time. Choose the date and time by selecting from the **By Time** options.

Application Options
Set the application options:

Maximum Parallel Streams per Database
Set the maximum data stream from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might improve restore speed, but high-bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring an Exchange database to its original location by using its original database name.

Advanced Options
Set the advanced job definition options:

Run cleanup immediately on job failure
This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see Configuring scripts. Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:
   - If you are running an on-demand job, click **Next**.
   - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.
    The restore job is created, and you can check on its status in **Jobs and Operations** > **Running Jobs**.

Accessing Exchange database files with instant access mode

You can access the Exchange database files by using the instant access restore type and mount the database files from the vSnap volume to an application server.

About this task

In instant access mode, no further action is taken after IBM Spectrum Protect Plus mounts the share. Use the data for custom recovery of data from the files in the vSnap volume.

Procedure

1. In the navigation panel, click **Manage Protection** > **Databases** > **Exchange** > **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **Exchange**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:

a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.

b) Click the plus icon next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list, click the minus icon next to the item.

c) Click **Next** to continue.

3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot
Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time
Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring
Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu.

Option	Description
	When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the

Option	Description
	operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Set destination** page, specify where you want to mount the database files and click **Next**.

Option	Description
Restore to original location	Select this option to mount the database files to the original server.
Restore to alternate location	Select this option to mount the database files to a local destination that is different from the original server, and then select the alternative location from the list of available servers.

6. On the **Restore Method** page, choose **Instant Access**, and then click **Next**.
7. Optional: On the **Job options** page, configure additional options if necessary and click **Next** to continue.
8. Optional: On the **Apply scripts** page, select the **Pre-Script** or **Post-Script** to apply, or choose **Continue job/task on script error**. For more information about working with scripts, see Configuring scripts. Click **Next** to continue.
9. Take one of the following actions on the **Schedule** page:
   - If you are running an on-demand job, click **Next**.
   - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.

    The restore job is created, and you can check on its status in **Jobs and Operations** > **Running Jobs**.
11. You can now access the Exchange database files on the application server mount point, and carry out any Exchange related or custom actions you want to do.

    Note: The Exchange database files on the mount point are read/write. However, updating them does not modify the original backup.
12. When you are finished with the instant access restore operation, click **Jobs and Operations** > **Active Resources** > **Databases**.
13. Select the resource, and then click **End Instant Disk Restore** to remove the mounted database and end the restore process.

MongoDB

After you successfully add MongoDB instances to IBM Spectrum Protect Plus, you can start to protect the data in your MongoDB databases. Create service level agreement (SLA) policies to back up and maintain MongoDB data.

Ensure that your MongoDB environment meets the system requirements. For more information, see “MongoDB requirements” on page 96.

Prerequisites for MongoDB

All system requirements and prerequisites for the IBM Spectrum Protect Plus MongoDB application server must be met before you start protecting MongoDB data with IBM Spectrum Protect Plus.

For MongoDB system requirements, see MongoDB system requirements.

To meet the prerequisites for MongoDB, complete the following checks and actions.

1. Ensure you have met the space prerequisites, as described in Space requirements for MongoDB protection.
2. Set the file size limit for the MongoDB instance user with the command **ulimit** `-f` to unlimited. Alternatively, set the value to sufficiently high to allow the copying of the largest database files in your backup and restore jobs. If you change the **ulimit** setting, restart the MongoDB instance to finalize the configuration.
3. If you are running MongoDB in a Linux environment, ensure that the installed sudo version is at a supported level.

   For more information about the version level, see “MongoDB requirements” on page 96. For information about setting sudo privileges, see “Setting sudo privileges” on page 564.
4. If your MongoDB databases are protected by authentication, you must set up role-based access control. For more information, see “Roles for MongoDB” on page 562.
5. Each MongoDB instance to be protected must be registered on IBM Spectrum Protect Plus. After the instances are registered, IBM Spectrum Protect Plus runs an inventory to detect MongoDB resources. Ensure that all instances that you want to protect are detected and listed correctly.
6. Ensure that the SSH service is running on port 22 on the server, and that firewalls are configured to allow IBM Spectrum Protect Plus to connect to the server with SSH. The SFTP subsystem for SSH must be enabled.
7. Ensure that you do not configure nested mount points.

Restrictions

The following restrictions apply to the MongoDB application server:

- MongoDB sharded cluster configurations are detected when you run an inventory, but these resources are not eligible for backup or restore operations.
- Unicode characters in MongoDB file path names cannot be handled by IBM Spectrum Protect Plus. All names must be in ASCII.

Virtualization

Protect your MongoDB environment with IBM Spectrum Protect Plus. The MongoDB application server that is installed on a VMware or Kernel-based Virtual Machine (KVM) virtual machine is protected when MongoDB is running on a supported operating system.

Roles for MongoDB

You must define role-based access control (RBAC) roles for the MongoDB agent users if authentication is enabled on the MongoDB database. When the roles are set up, users can protect and monitor MongoDB resources with IBM Spectrum Protect Plus in accordance with the users' defined roles.

Role-based access control for MongoDB

For each MongoDB user, specify access roles by using a command similar to the following example:

```
use admin
db.grantRolesToUser("<username>",
[ { role: "hostManager", db: "admin" },
{ role: "clusterManager", db: "admin" } ] )
```

The following roles are available:

hostManager
: This role provides access to the **fsyncLock** command. This access is required for application-consistent backups of MongoDB databases where journaling is not enabled. This role also provides access to the shutdown command, which is used during a restore operation to shut down the MongoDB server instance that the restore is directed to.

clusterMonitor
This role provides access to commands for monitoring and reading the state of the MongoDB database. The following commands are available to users with this role:

- **getCmdLineOpts**
- **serverVersion**
- **replSetGetConfig**
- **replSetGetStatus**
- **isMaster**
- **listShards**

clusterManager
This role is only required only for running test restore operations of replica sets. Users who run the **replSetReconfig** command can create the restored instance of a single node replica set. This role enables read and write access during test restore operations of replica sets. Without this access, the node in the replica set would remain in the REMOVED state without read and write access. In addition, this role provides access to commands for reading the state of the MongoDB database. The following commands are available for this role:

- **replSetReconfig**
- **getCmdLineOpts**
- **serverVersion**
- **replSetGetConfig**
- **replSetGetStatus**
- **isMaster**
- **listShards**

Space prerequisites for MongoDB protection

Before you start backing up MongoDB data, ensure that you have enough free space on the target and source hosts, and in the vSnap repository. Extra space is required to store temporary Logical Volume Manager (LVM) backups of logical volumes where the MongoDB data is located. These temporary backups, that are known as LVM snapshots, are created automatically by the MongoDB agent.

LVM snapshots

LVM snapshots are point-in-time copies of LVM logical volumes. After the file copy operation finishes, earlier LVM snapshots are removed by the IBM Spectrum Protect Plus MongoDB agent in a cleanup operation.

For each LVM snapshot logical volume, you must allocate at least 10 percent free space in the volume group. If there is enough free space in the volume group, the IBM Spectrum Protect Plus MongoDB agent reserves up to 25 percent of the source logical volume size for the snapshot logical volume.

Linux LVM2

When you run a MongoDB backup operation, MongoDB requests a snapshot. This snapshot is created on a Logical Volume Management (LVM) system for each logical volume with data or logs for the selected database. On Linux systems, logical volumes are managed by LVM2.

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

On Linux, the LVM2 volume manager stores the snapshot of a logical volume within the same volume group. There must be enough space available to store the logical volume. The logical volume grows in size as the data changes on the source volume for the lifetime of the snapshot.

Setting sudo privileges

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

About this task

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

Procedure

1. Create an agent user by issuing the following command:

   `useradd -m` *`agent`*

   where *agent* specifies the name of the IBM Spectrum Protect Plus agent user.
2. Set a password for the new user by issuing the following command:

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

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

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

Adding a MongoDB application server

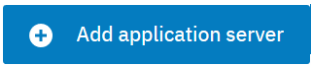
To start protecting MongoDB resources, you must add the server that hosts your MongoDB instances, and set credentials for the instances. Repeat the procedure to add all the servers that host MongoDB resources.

About this task

To add a MongoDB application server to IBM Spectrum Protect Plus, you must have the host address of the machine.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Applications** > **MongoDB**.
2. In the **MongoDB** window, click **Manage Application Servers**, and click **Add Application Server** to add the host machine.

3. In the **Application Properties** form, enter the host address.
4. Choose to register the host with a user or an SSH key.

   If you select **User**, you can choose to enter a new user and password, or an existing user. If you select **SSH Key**, select the SSH key from the menu.

   Restriction: Any user that is specified must have sudo privileges set up.

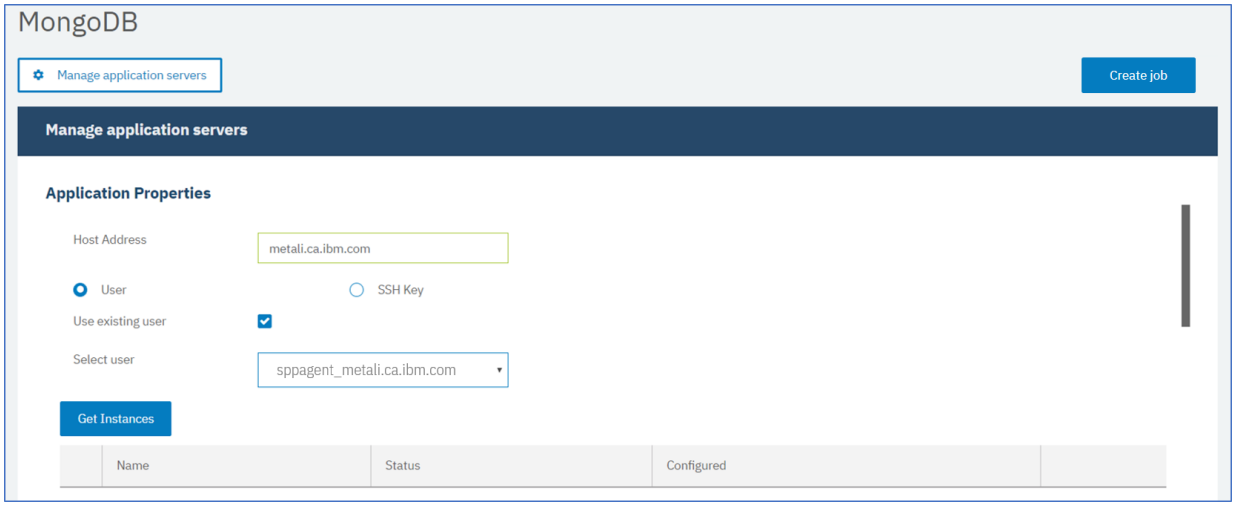


Figure 53. Adding a MongoDB agent

5. Click **Get Instances** to detect and list the MongoDB instances that are available on the host server that you are adding.

   Each MongoDB instance is listed with its connection host address, status, and an indication of whether it is configured.

   Attention: If you register more than one application server for one replica set, the instance name that is displayed might change after each inventory, backup, or restore operation. The host name of the most recently added application server that belongs to the replica set is used as part of the instance name. An inventory operation is run as part of backup and restore operations.

6. If you are using access control, configure an instance by setting credentials. Click **Set Credential**, and set the user ID, and password. Alternatively, you can select to use an existing user profile.

   For more information about access control, see Chapter 19, “Managing user access,” on page 643.

   When you set credentials, you assign MongoDB user roles for the backup and restore operations with access to role-protected MongoDB servers by using Salted Challenge Response Authentication Mechanism (SCRAM), or Challenge and response authentication. The MongoDB user that is assigned for the role-protected MongoDB server requires one of the following access levels to protect resources:

   - *Host Manager*: manages the database as the administrator. This role is required for taking and managing snapshots.
   - *Cluster Administrator*: retrieves configuration information and runs test mode restore operations of MongoDB replica sets. This role is required to reconfigure test mode restore operations of MongoDB replica sets for data queries.
   - *Cluster Monitor*: monitors the protection of MongoDB resources, and retrieves configuration information.

7. Optional: Set the option **Maximum concurrent databases** by entering a number in the field.
8. Save the form, and repeat the steps to add other MongoDB application servers to IBM Spectrum Protect Plus.

What to do next

After you add MongoDB application servers to IBM Spectrum Protect Plus, an inventory is automatically run on each application server to detect the relevant databases in those instances.

To verify that the databases are added, review the job log. Go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as a wildcard in the name.

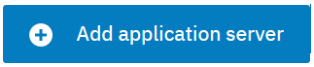
Databases must be detected to ensure that they can be protected. For instructions about running a manual inventory, see Detecting MongoDB resources.

Registering a MongoDB Ops Manager Application Database for protection

To protect your MongoDB Ops Manager Application Database, you must first register the Ops Manager host address with IBM Spectrum Protect Plus.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Applications** > **MongoDB**.
2. In the **MongoDB** window, click **Manage Application Servers**, and click **Add Application Server**.

3. In the Application Properties form, enter the host address for the Ops Manager Application Database. Get instances and set credentials by following the steps outlined in “Adding a MongoDB application server” on page 564.

   The Ops Manager Application Database is listed in the Instances table as shown in the following example:

   ```
   metali8.limerick.ie.ibm.com Connection: '333.0.5.1:88888' Ops Manager Application Database
   ```

What to do next

The MongoDB Ops Manager Application Database is available for backing up. You can define backup and restore jobs to protect your data. To regularly back up your data, define a backup job that includes a service level agreement (SLA) policy. For more information, see “Backing up MongoDB data” on page 568 and “Defining a regular service level agreement job” on page 569.

Detecting MongoDB resources

After you add your MongoDB application servers to IBM Spectrum Protect Plus, an inventory is run automatically to detect all MongoDB instances and databases. You can run a manual inventory on any application server to detect, list, and store all MongoDB databases for the selected host.

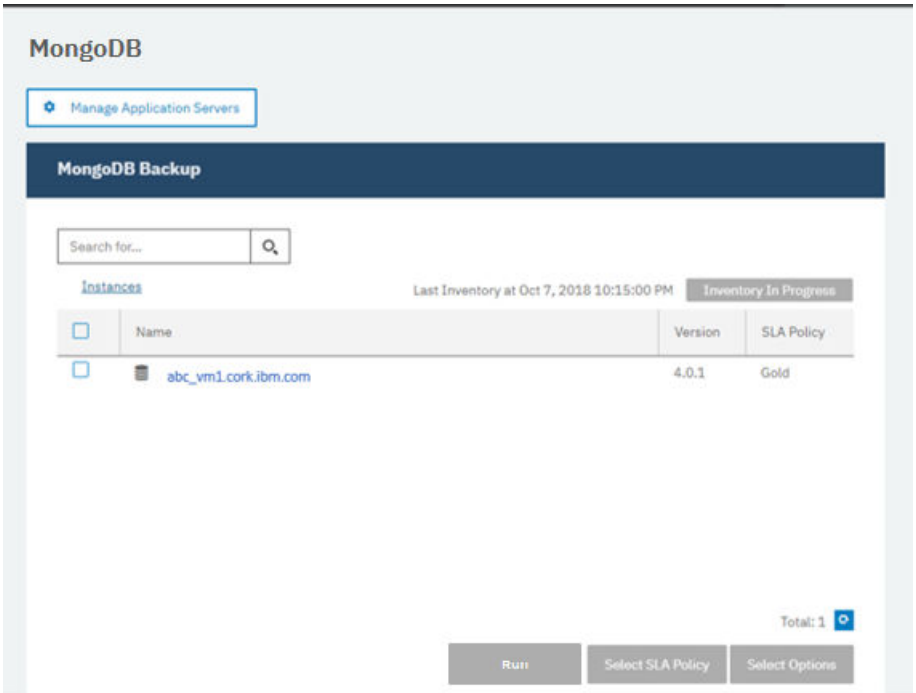
Before you begin

Ensure that you added your MongoDB application servers to IBM Spectrum Protect Plus. For instructions, see Adding a MongoDB application server.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Applications** > **MongoDB**.

   Tip: To add more MongoDB instances to the **Instances** pane, follow the instructions in Adding a MongoDB application server.
2. Click **Run Inventory**.



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

To monitor the inventory job, go to **Jobs and Operations**. Click the **Running Jobs** tab, and look for the latest Application Server Inventory log entry.

Completed jobs are shown on the **Job History** tab. You can use the **Sort By** list to sort jobs based on start time, type, status, job name, or duration. Use the **Search by name** field to search for jobs by name. You can use asterisks as wildcard characters in the name.

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

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

   Attention: If you register more than one application server for one replica set, the instance name that is displayed might change after each inventory, backup, or restore operation. The host name of the most recently inventoried application server that belongs to the replica set is used as part of the instance name. An inventory operation is run as part of backup and restore operations.

What to do next

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

Testing the MongoDB connection

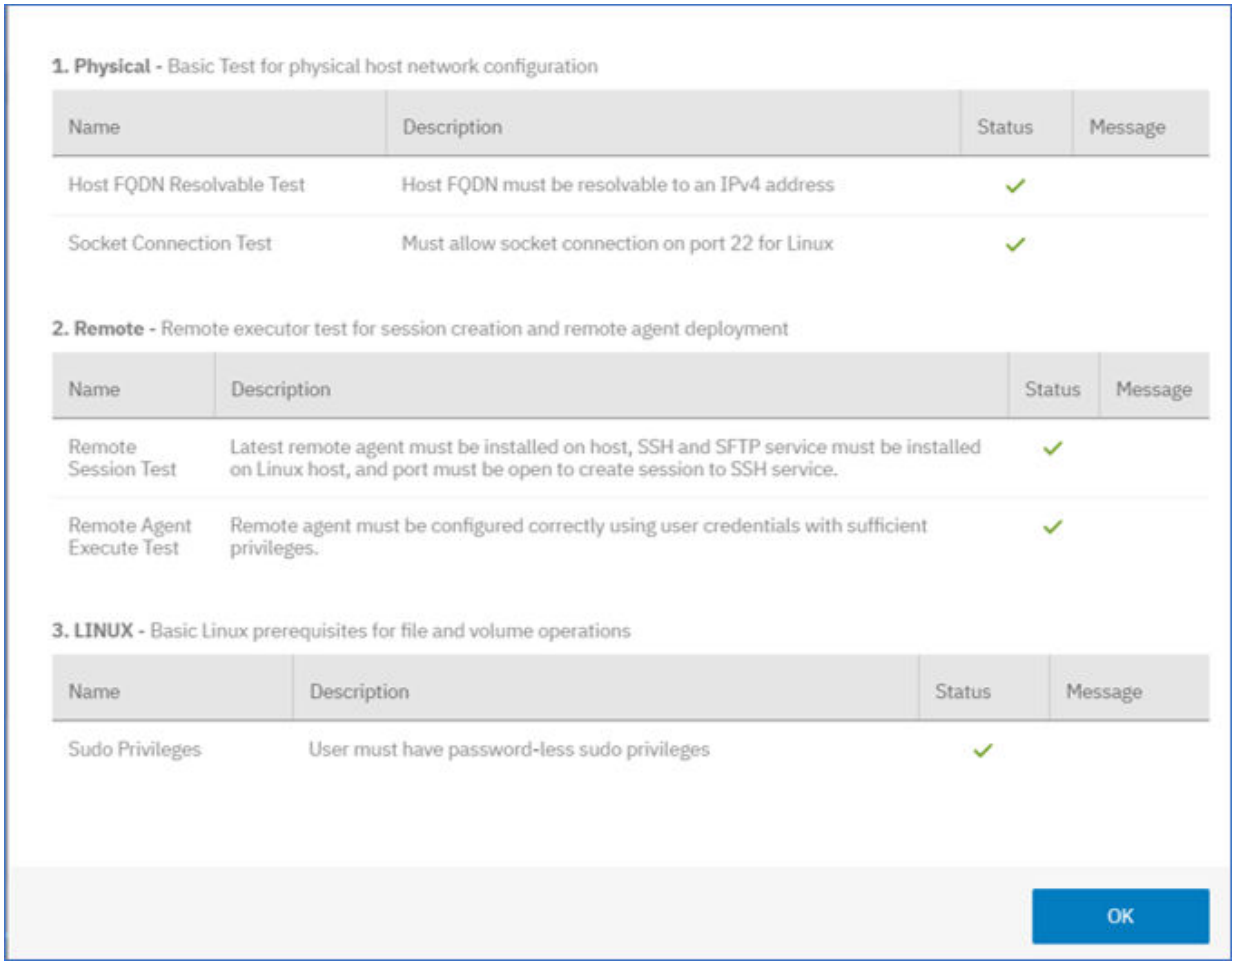
After you add a MongoDB application server, you can test the connection. The test verifies communication between IBM Spectrum Protect Plus and the MongoDB server. It also checks that the correct sudo permissions area available for the user who is running the test.

Procedure

1. In the navigation panel, click **Manage Protection** > **Applications** > **MongoDB**.

2. In the **MongoDB** window, click **Manage Application Servers**, and select the host address that you want to test.

   A list of the MongoDB application servers that are available is shown.
3. Click **Actions** and choose **Test** to start the verification tests for physical and remote system connections and settings.



   The test report displays a list that includes tests for the physical host network configuration, and tests for the remote server installation on the host.
4. Click **OK** to close the test report. If issues are reported, fix the issues and rerun the test to verify the fixes.

Backing up MongoDB data

You can define backup jobs to protect your MongoDB data. To regularly back up your data, define a backup job that includes a service level agreement (SLA) policy.

Before you begin

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

Review the following prerequisites before you create a backup job definition:

- Add the application servers that you want to back up. For the procedure, see Adding a MongoDB application server.
- Configure an SLA Policy. For the procedure, see Defining a Service Level Agreement backup job.

- Before an IBM Spectrum Protect Plus user can set up backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources, and backup and restore operations, in the **Accounts** pane. For more information, see Chapter 19, “Managing user access,” on page 643 and “Roles for MongoDB” on page 562.

Restriction: Do not run inventory jobs at the same time that backup jobs are scheduled.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Applications** > **MongoDB**.
2. Select the check box for the instance that you want to back up.

   Under each MongoDB instance, data to be backed up is listed as **ALL**. Each instance in the Instances pane is listed by instance name, version, and the applied SLA policy.
3. Click **Select Options** to specify the number of parallel streams for the backup operation, and then click **Save**. By selecting an appropriate number of parallel streams, you can minimize the time that is required for the backup job.

   The saved options are used for all backup jobs for this instance as selected.
4. To run the backup job with these options, click the instance name, select the **ALL** database representation, and click **Run**.

   The backup job begins, and you can view the details in **Jobs and Operations** > **Running Jobs**.

   Tip: The **Run** button is only enabled if an SLA policy is applied to the **ALL** representation of the databases.

   To run an on-demand backup job for multiple databases that are associated with an SLA policy, click **Create job**, select **Ad hoc backup**, and follow the instructions in “Running an ad hoc backup job” on page 627.
5. Select the instance again, and click **Select an SLA Policy** to choose an SLA policy.
6. Save the SLA selection.

   To define a new SLA or to edit an existing policy with custom retention and frequency rates, select **Manage Protection** > **Policy Overview**. In the **SLA Policies** pane, click **Add SLA Policy**, and define policy preferences.

What to do next

After the SLA policy is saved, you can run the policy at any time by clicking **Actions** for that policy name, and selecting **Start**. The status in the log changes to show that the backup job is in the `Running` state.

To cancel a job that is running, click **Actions** for that policy name and select **Cancel**. A message asks whether you want to keep the data that is already backed up. Choose **Yes** to keep the backed up data, or **No** to discard the backup.

Defining a regular service level agreement job

After your MongoDB instances are listed, select and apply an SLA policy to start protecting your data.

Procedure

1. In the navigation panel, expand **Manage Protection** > **Applications** > **MongoDB**.
2. Select the MongoDB instance to back up all the data in that instance.

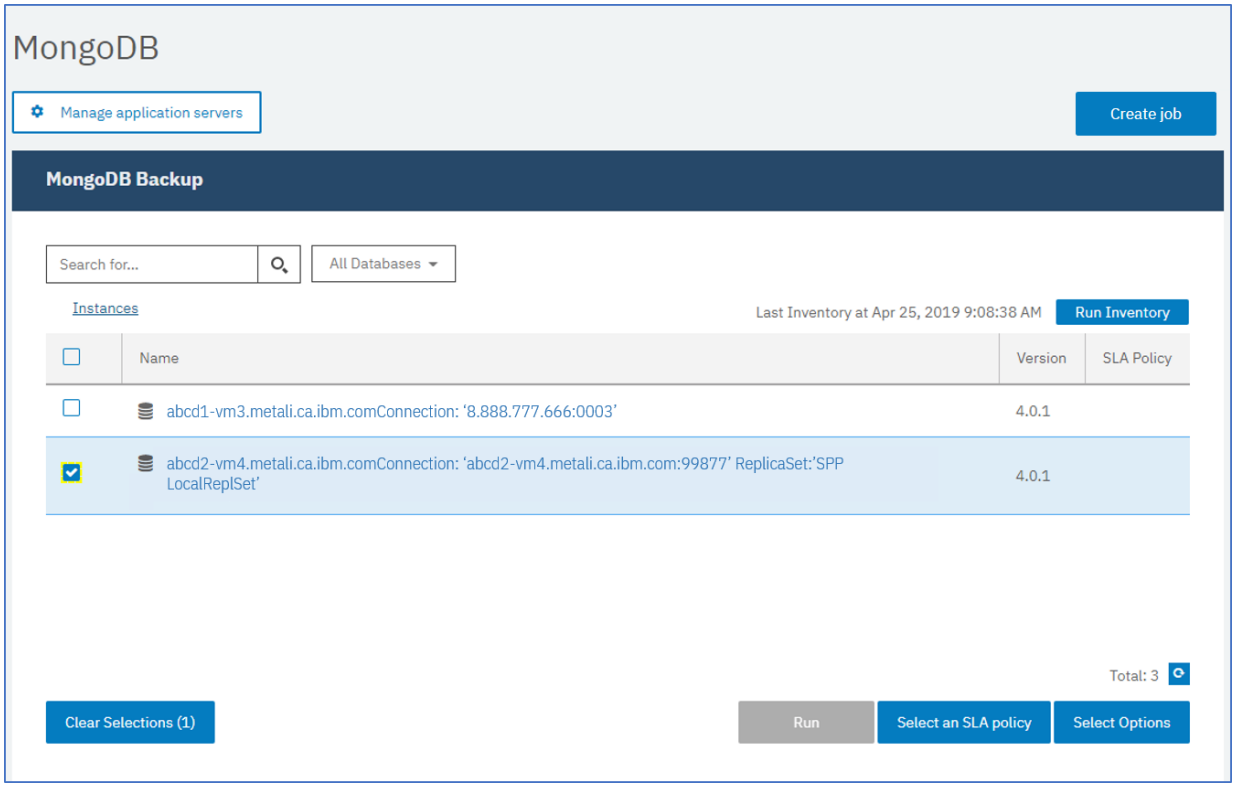


Figure 54. MongoDB Backup pane showing instances

3. Click **Select an SLA policy** and choose an SLA policy. Save your choice.

   Predefined choices are Gold, Silver, and Bronze, each with different frequencies and retention rates. You can also create a custom SLA policy by navigating to **Policy Overview** > **Add SLA Policy**.

4. Optional: To enable multiple backup streams to reduce the time that is taken to back up large databases, click **Select Options** and enter a number of parallel streams. Save your changes.

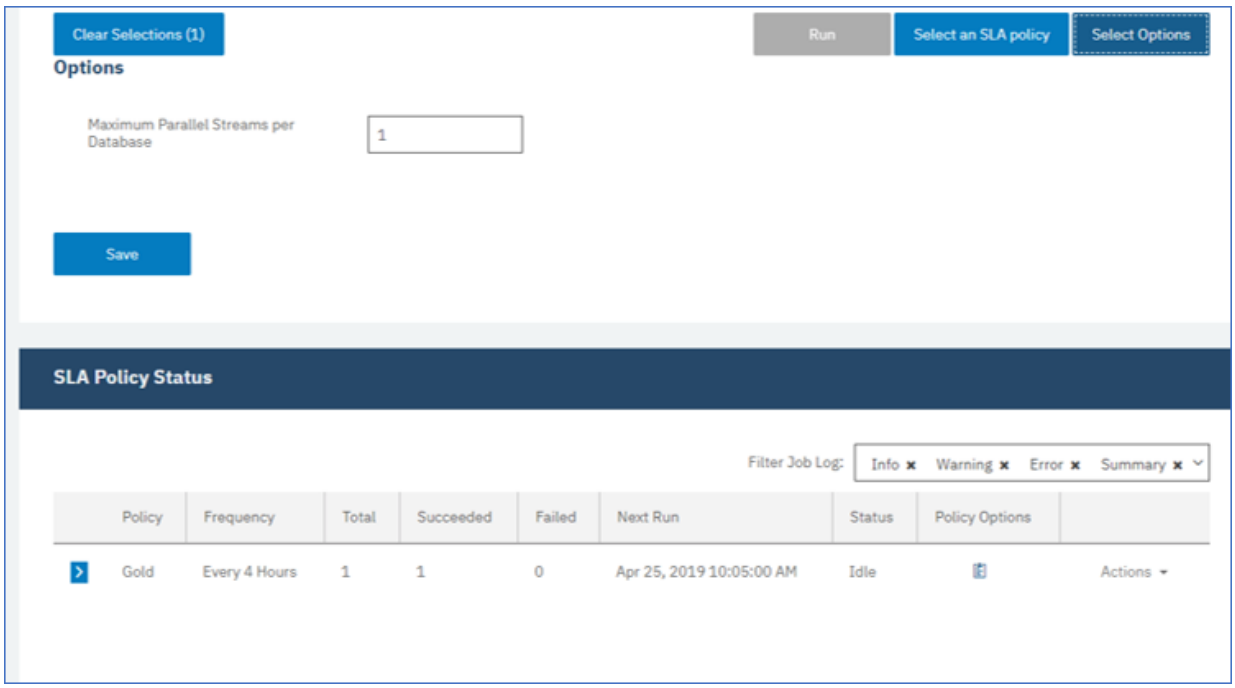


Figure 55. Backup options and SLA Policy Status

Tips for configuring the option:

Review the following tips to help you configure the option for the backup job:

- To set the option for child resources to the same value as the parent, click **Set all options to inherit**.
- If multiple resources were selected for the backup job, the option value is indeterminate. If you change a value for the option, that value is used for all selected resources after you click **Save**.
- If an option value is changed from the previously saved value, the option is shown in yellow.
- To close the **Options** pane without saving changes, click **Select Options**.

5. Configure the SLA policy by clicking the icon in the **Policy Options** column of the **SLA Policy Status** table.

   For more information about SLA configuration options, see "Setting SLA configuration options for your backup" on page 571.
6. To run the policy outside of the scheduled job, select the instance. Click the **Actions** button and select **Start**. The status changes to **Running** for your chosen SLA and you can follow the progress of the job in the log shown.

What to do next

After the SLA policy is saved, you can run the policy at any time by clicking **Actions** for that policy name, and selecting **Start**. The status in the log changes to show that the backup job is in the `Running` state.

To cancel a job that is running, click **Actions** for that policy name and select **Cancel**. A message asks whether you want to keep the data that is already backed up. Choose **Yes** to keep the backed up data, or **No** to discard the backup.

Setting SLA configuration options for your backup

After you set up a service level agreement (SLA) policy for your backup job, you can choose to configure extra options for that job. Additional SLA options include running scripts, and forcing a full base backup.

Procedure

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

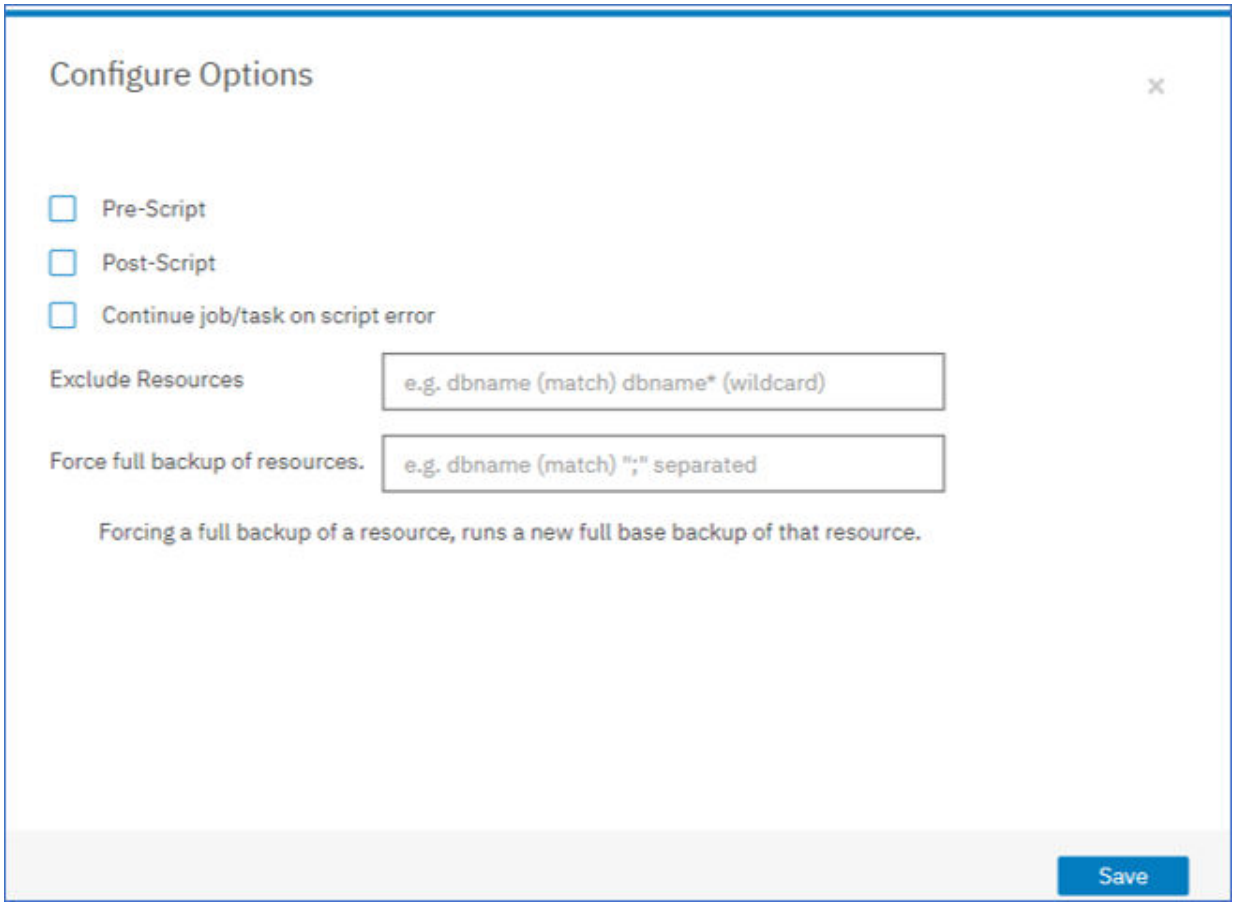


Figure 56. Specifying additional SLA configuration options

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

   Scripts and script servers are configured on the **System Configuration** > **Script** page. For more information about working with scripts, see **Configuring scripts**.
4. To continue running the job when the script that is associated with the job fails, select **Continue job/task on script error**.

   If this option is selected, the backup or restore operation is reattempted after an initial fail, and the script task status is reported as COMPLETED when the script completes processing with a nonzero return code. If this option is not selected, the backup or restore is not reattempted and the script task status is reported as FAILED.
5. Skip **Exclude Resources** for MongoDB SLA options, as you cannot specify resources to exclude. Instances are backed up rather than individual databases.
6. To create a full, new backup of a MongoDB instance, select **Force full backup of resources**.

   A full new backup of that resource is created to replace the existing backup of that resource for one occurrence only. After that the resource is backed up incrementally as before.

Restoring MongoDB data

To restore data, define a job that restores data to the latest backup or select an earlier backup copy. Choose to restore data to the original instance or to an alternative instance on a different machine, creating a cloned copy. Define and save the restore job to run as an ad hoc operation, or to run regularly as a scheduled job.

Before you begin

Before you create a restore job for MongoDB, ensure that the following requirements are met:

- At least one MongoDB backup job is set up and running successfully. For instructions about setting up a backup job, see "Backing up MongoDB data" on page 568.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For instructions about assigning roles, see Chapter 19, "Managing user access," on page 643, and "Roles for MongoDB" on page 562.
- Enough disk space is allocated at the target server for the restore operation.
- Dedicated volumes are allocated for file copying.
- The same directory structure and layout are available on both the target and source servers.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

For restore operations to alternative instances, MongoDB must be at the same version level on the target and host machines.

For more information about space requirements, see Space prerequisites for MongoDB protection. For more information about prerequisites and setup, see Prerequisites for MongoDB.

Procedure

To define a MongoDB restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Applications** > **MongoDB** > **Create job**, and then select **Restore** to open the Restore wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **MongoDB**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

2. On the **Select source** page, take the following actions:
   a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
   b) Click the add to restore list icon next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

      The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the remove from restore list icon next to the item.
   c) Click **Next** to continue.
3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot
Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time
Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring
Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Restore method** page, choose the type of restore operation, and click **Next** to continue.
   - **Test**: In this mode, the agent creates a database by using the data files directly from the vSnap repository. This option is available only when you are restoring data to an alternative instance. Members of replica sets will not be reconfigured after the MongoDB server is started. The server is started as a single-node replica set.
   - **Production**: In this mode, the MongoDB application server first copies the files from the vSnap repository to the target host. The copied data is then used to start the database. MongoDB

instances that are members of a replica set are not started during a production restore operation. This action prevents data from being overwritten when connecting to the replica set.

- **Instant Access**: In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the share. Use the data for custom recovery from the files in the vSnap repository.

For test mode or production mode, you can optionally enter a new name for the restored database.

For production mode, you can also specify a new folder for the restored database by expanding the database and entering a new folder name.

6. On the **Set destination** page, select **Restore to original instance** to restore to the original server, or **Restore to alternate instance** to restore to a different location that you can select from the locations listed.

   For more information about restoring data to the original instance, see Restoring to the original instance. For more information about restoring your data to an alternative instance, see Restoring to an alternate instance.

7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

   In the **Recovery Options** section, the **Recover until end of backup** for MongoDB is selected by default. This option recovers the selected data to the state it was in at the time the backup was created. The recovery operation makes use of the log files that are included in the MongoDB backup.

   Application Options

   Set the application options:

   Overwrite existing database

   Enable this option to allow the restore job to overwrite the selected database. If this option is not selected, the restore job fails when data with the same name is found during the restore process.

   Attention: Ensure that no other data shares the same local database directory as the original data or the data will be overwritten.

   Maximum Parallel Streams per Database

   Set the maximum number of parallel data streams from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might speed up restore operations, but high bandwidth consumption might affect overall system performance.

   This option is applicable only when you are restoring a MongoDB database to its original location by using its original database name.

   Advanced Options

   Set the advanced job definition options:

   Run cleanup immediately on job failure

   This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

   Allow session overwrite

   Select this option to replace existing databases with the same name during a restore operation. During an instant disk restore operation, the existing database is shut down and overwritten, and then the recovered database is restarted. If this option is not selected and a database with the same name is encountered, the restore operation fails with an error.

   Continue with restores of other selected databases even if one fails

   If one database in the instance is not successfully restored, the restore operation continues for all other data that is being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.

Mount Point Prefix
For **Instant Access** restore operations, specify a mount point prefix for the path where the mount is to be directed.

8. Optional: On the **Apply scripts** page, specify scripts that can be run before or after a job runs. Batch and PowerShell scripts are supported on Windows operating systems while shell scripts are supported on Linux operating systems.

   Pre-Script
   Select this check box to choose an uploaded script and an application or script server where the pre-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration** > **Script**.

   Post-Script
   Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration** > **Script** page.

   Continue job/task on script error
   Select this option to continue running the job when the script that is associated with the job fails. When this option is enabled, in the event that a script completes processing with a nonzero return code, the backup or restore job continues to run and the pre-script task status is reported as `COMPLETED`. If a post-script completes processing with a nonzero return code, the post-script task status is reported as `COMPLETED`. When this option is not selected, the backup or restore job does not run, and the pre-script or post-script task is reported as `FAILED`.

   Click **Next** to continue.
9. On the **Schedule** page, click **Next** to start on-demand jobs after you complete the Restore wizard. For recurring jobs, enter a name for the job schedule, and specify how often and when to start the restore job.
10. On the **Review** page, review your restore job settings.

    Attention: Review the selected options before you proceed to **Submit** because data will be overwritten when the **Overwrite existing data** application option is selected. You can cancel a restore job when it is in progress, but if the **Overwrite existing data** option is selected, data is overwritten even if you cancel the job.

11. To proceed with the job, click **Submit**. To cancel the job, navigate to **Jobs and Operations** and click the **Schedule** tab. Find the restore job you want to cancel. Click **Actions**, and select **Cancel**.

Results

A few moments after you select **Restore**, the **onDemandRestore** job is added to the **Jobs and Operations** > **Running Jobs** pane. Click the record to show the step-by-step details of the operation. You can also download the zipped log file by clicking **Download.zip**. For any other jobs, click the **Running Jobs** or **Job History** tabs and click the job to display its details.

The IP address and port for the restored server can be found in the log file for the restore operation. Navigate to **Jobs and Operations** > **Running Jobs** to find the logs for your restore operation.

For information about restoring data to the original instance, see Restoring to the original instance. For information about restoring your data to an alternative instance, see Restoring to an alternate instance.

Restoring MongoDB data to the original instance

You can restore a MongoDB instance to the original host and choose between restoring to the latest backup or an earlier MongoDB database backup version. When you restore data to its original instance, you cannot rename it. This restore option runs a full production restoration of data, and existing data is overwritten at the target site if the **Overwrite existing databases** application option is selected.

Before you begin

Before you create a restore job for MongoDB, ensure that the following requirements are met:

- At least one MongoDB backup job is set up and running successfully. For instructions about setting up a backup job, see "Backing up MongoDB data" on page 568.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For instructions about assigning roles, see Chapter 19, "Managing user access," on page 643, and "Roles for MongoDB" on page 562.
- Enough disk space is allocated at the target server for the restore operation.
- Dedicated volumes are allocated for file copying.
- The same directory structure and layout are available on both the target and source servers.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

For more information about space requirements, see Space prerequisites for MongoDB protection. For more information about prerequisites and setup, see Prerequisites for MongoDB.

Procedure

1. In the navigation panel, click **Manage Protection** > **Applications** > **MongoDB** > **Create job**, and then select **Restore** to open the Restore wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **MongoDB**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.
2. On the **Select source** page, take the following actions:
   a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
   b) Click the add to restore list icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

      The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the remove from restore list icon  next to the item.
   c) Click **Next** to continue.
3. On the **Source snapshot** page, select the type of restore job that you want to create:

   On-demand: Snapshot
   Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

   On-demand: Point in Time
   Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

   Recurring
   Creates a repeating point-in-time restore job that runs on a schedule.
4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.

Option	Description
	Cloud service archive The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Restore method** page, choose the type of restore operation, and click **Next** to continue.
   - **Production**

     To recover an entire instance to the original instance, the preferred method is to choose this option with the overwrite application option. MongoDB instances that are members of a replica set are not started during a production restore operation. This action prevents data from being overwritten when connecting to the replica set.
   - **Test**

     Choose this option to restore data to the same server but using a different port.
   - **Instant Access**

     Choose this option to mount the backup to the application server without restoring the data or overwriting the data.

     Click **Next** to continue.

   For test mode or production mode, you can optionally enter a new name for the restored database.

   For production mode, you can also specify a new folder for the restored database by expanding the database and entering a new folder name.
6. On the **Set destination** page, choose **Restore to original instance** and click **Next**.

7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

   In the **Recovery Options** section, the **Recover until end of backup** for MongoDB is selected by default. This option recovers the selected data to the state it was in at the time the backup was created. The recovery operation makes use of the log files that are included in the MongoDB backup.

   Application Options

   Set the application options:

   Overwrite existing database

   Enable this option to allow the restore job to overwrite the selected database. If this option is not selected, the restore job fails when data with the same name is found during the restore process.

   Attention: Ensure that no other data shares the same local database directory as the original data or the data will be overwritten.

   Maximum Parallel Streams per Database

   Set the maximum number of parallel data streams from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might speed up restore operations, but high bandwidth consumption might affect overall system performance.

   This option is applicable only when you are restoring a MongoDB database to its original location by using its original database name.

   Advanced Options

   Set the advanced job definition options:

   Run cleanup immediately on job failure

   This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

   Allow session overwrite

   Select this option to replace existing databases with the same name during a restore operation. During an instant disk restore operation, the existing database is shut down and overwritten, and then the recovered database is restarted. If this option is not selected and a database with the same name is encountered, the restore operation fails with an error.

   Continue with restores of other selected databases even if one fails

   If one database in the instance is not successfully restored, the restore operation continues for all other data that is being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.

   Mount Point Prefix

   For **Instant Access** restore operations, specify a mount point prefix for the path where the mount is to be directed.

8. Optional: On the **Apply scripts** page, specify scripts that can be run before or after a job runs. Batch and PowerShell scripts are supported on Windows operating systems while shell scripts are supported on Linux operating systems.

   Pre-Script

   Select this check box to choose an uploaded script and an application or script server where the pre-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration** > **Script**.

   Post-Script

   Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration** > **Script** page.

Continue job/task on script error
Select this option to continue running the job when the script that is associated with the job fails. When this option is enabled, in the event that a script completes processing with a nonzero return code, the backup or restore job continues to run and the pre-script task status is reported as `COMPLETED`. If a post-script completes processing with a nonzero return code, the post-script task status is reported as `COMPLETED`. When this option is not selected, the backup or restore job does not run, and the pre-script or post-script task is reported as `FAILED`.

Click **Next** to continue.

9. On the **Schedule** page, click **Next** to start on-demand jobs after you complete the Restore wizard. For recurring jobs, enter a name for the job schedule, and specify how often and when to start the restore job.
10. On the **Review** page, review your restore job settings.



Attention: Review the selected options before you proceed to **Submit** because data will be overwritten when the **Overwrite existing data** application option is selected. You can cancel a restore job when it is in progress, but if the **Overwrite existing data** option is selected, data is overwritten even if you cancel the job.

11. To proceed with the job, click **Submit**. To cancel the job, navigate to **Jobs and Operations** and click the **Schedule** tab. Find the restore job you want to cancel. Click **Actions**, and select **Cancel**.

Restoring MongoDB data to an alternative instance

You can select a MongoDB database backup and restore it to an alternative host. You can also choose to restore a database to a different vSnap repository, or you can rename the database. This process creates an exact copy of the instance on a different host.

Before you begin

Before you create a restore job for MongoDB, ensure that the following requirements are met:

- At least one MongoDB backup job is set up and running successfully. For instructions about setting up a backup job, see “Backing up MongoDB data” on page 568.
- IBM Spectrum Protect Plus roles and resource groups are assigned to the user who is setting up the restore job. For instructions about assigning roles, see Chapter 19, “Managing user access,” on page 643, and “Roles for MongoDB” on page 562.
- Enough disk space is allocated at the target server for the restore operation.
- Dedicated volumes are allocated for file copying.
- The same directory structure and layout are available on both the target and source servers.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

For restore operations to alternative instances, MongoDB must be at the same version level on the target and host machines.

For more information about space requirements, see Space prerequisites for MongoDB protection. For more information about prerequisites and setup, see Prerequisites for MongoDB.

Procedure

1. In the navigation panel, click **Manage Protection** > **Applications** > **MongoDB** > **Create job**, and then select **Restore** to open the Restore wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **MongoDB**.

- For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
- The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

2. On the **Select source** page, take the following actions:
   a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
   b) Click the add to restore list icon next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

      The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the remove from restore list icon next to the item.
   c) Click **Next** to continue.
3. On the **Source snapshot** page, select the type of restore job that you want to create:

   On-demand: Snapshot
   Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

   On-demand: Point in Time
   Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

   Recurring
   Creates a repeating point-in-time restore job that runs on a schedule.
4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

   The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

   Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.

Option	Description
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.

Option	Description
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Restore method** page, choose the type of restore operation, and click **Next** to continue.
   - **Test**: In this mode, the agent creates a database by using the data files directly from the vSnap repository. This option is available only when you are restoring data to an alternative instance. Members of replica sets will not be reconfigured after the MongoDB server is started. The server is started as a single-node replica set.
   - **Production**: In this mode, the MongoDB application server first copies the files from the vSnap repository to the target host. The copied data is then used to start the database. MongoDB instances that are members of a replica set are not started during a production restore operation. This action prevents data from being overwritten when connecting to the replica set.
   - **Instant Access**: In this mode, no further action is taken after IBM Spectrum Protect Plus mounts the share. Use the data for custom recovery from the files in the vSnap repository.

   For test mode or production mode, you can optionally enter a new name for the restored database.

   For production mode, you can also specify a new folder for the restored database by expanding the database and entering a new folder name.
6. In the **Set destination** page, choose **Restore to alternate instance** and select the target instance that you want to restore the data to.

   The original instance is not selectable because you cannot overwrite the original data when you select **Restore to alternate instance**. You also cannot select instances on different versions levels or instances on the same host as the original instance.

   Click **Next** to continue.
7. Optional: On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

   In the **Recovery Options** section, the **Recover until end of backup** for MongoDB is selected by default. This option recovers the selected data to the state it was in at the time the backup was created. The recovery operation makes use of the log files that are included in the MongoDB backup.

   Application Options

   Set the application options:

   Overwrite existing database

   Enable this option to allow the restore job to overwrite the selected database. If this option is not selected, the restore job fails when data with the same name is found during the restore process.

   Attention: Ensure that no other data shares the same local database directory as the original data or the data will be overwritten.

   Maximum Parallel Streams per Database

   Set the maximum number of parallel data streams from the backup storage per database. This setting applies to each database in the job definition. Multiple databases can still be restored in parallel if the value of the option is set to 1. Multiple parallel streams might speed up restore operations, but high bandwidth consumption might affect overall system performance.

This option is applicable only when you are restoring a MongoDB database to its original location by using its original database name.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow session overwrite

Select this option to replace existing databases with the same name during a restore operation. During an instant disk restore operation, the existing database is shut down and overwritten, and then the recovered database is restarted. If this option is not selected and a database with the same name is encountered, the restore operation fails with an error.

Continue with restores of other selected databases even if one fails

If one database in the instance is not successfully restored, the restore operation continues for all other data that is being restored. When this option is not selected, the restore job stops when the recovery of a resource fails.

Mount Point Prefix

For **Instant Access** restore operations, specify a mount point prefix for the path where the mount is to be directed.

8. Optional: On the **Apply scripts** page, specify scripts that can be run before or after a job runs. Batch and PowerShell scripts are supported on Windows operating systems while shell scripts are supported on Linux operating systems.

   Pre-Script

   Select this check box to choose an uploaded script and an application or script server where the pre-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration** > **Script**.

   Post-Script

   Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server, clear the **Use Script Server** check box. To configure scripts and script servers, click **System Configuration** > **Script** page.

   Continue job/task on script error

   Select this option to continue running the job when the script that is associated with the job fails. When this option is enabled, in the event that a script completes processing with a nonzero return code, the backup or restore job continues to run and the pre-script task status is reported as `COMPLETED`. If a post-script completes processing with a nonzero return code, the post-script task status is reported as `COMPLETED`. When this option is not selected, the backup or restore job does not run, and the pre-script or post-script task is reported as `FAILED`.

   Click **Next** to continue.

9. On the **Schedule** page, click **Next** to start on-demand jobs after you complete the Restore wizard. For recurring jobs, enter a name for the job schedule, and specify how often and when to start the restore job.
10. On the **Review** page, review your restore job settings.

    Attention: Review the selected options before you proceed to **Submit** because data will be overwritten when the **Overwrite existing data** application option is selected. You can cancel a restore job when it is in progress, but if the **Overwrite existing data** option is selected, data is overwritten even if you cancel the job.

11. To proceed with the job, click **Submit**. To cancel the job, navigate to **Jobs and Operations** and click the **Schedule** tab. Find the restore job you want to cancel. Click **Actions**, and select **Cancel**.

Using a granular restore operation for MongoDB

You can restore specific MongoDB databases or collections by using a granular restore operation. For a granular restore operation, first run a test restore job and then run the appropriate MongoDB commands.

Before you begin

If authentication is enabled, you must provide credentials for users so that they can correct permissions on the instance in the test restore operation.

About this task

The granular restore operation for MongoDB is based on a test mode restore job. When you run the test restore job on IBM Spectrum Protect Plus, and the **mongodump** and **mongorestore** commands on the MongoDB server, you can access individual databases or collections from the recovery source.

Use this procedure to complete either of the following tasks:

- Restore any number of databases by using the **mongodump** and **mongorestore** commands for the databases that you require.
- Restore any number of collections by using the **mongodump** and **mongorestore** commands for the collections that you require.

Procedure

1. In the navigation panel, click **Manage Protection** > **Applications** > **MongoDB** > **Create job**, and then select **Restore** to open the **Restore** wizard.
2. On the **Select source** page, take the following actions:
   a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
   b) Click the add to restore list icon next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

      The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the remove from restore list icon next to the item.
   c) Click **Next** to continue.
3. On the **Restore method** page, select **Test**, and click **Next** to continue with the test restore process.
4. On the **Set destination** page, choose **Restore to alternate instance**, and select the target instance that you want to restore the data to.

   You cannot select the original instance is not selectable as you cannot overwrite the original data when you select **Restore to alternate instance**. Instances on different versions levels cannot be selected. Other instances on the same host as the original instance, cannot be selected either.

   Click **Next** to continue.
5. Proceed through the restore wizard pages and select the required options.
6. On the **Review** page, review your restore job settings.

   Attention: Review the selected options before you proceed to **Submit** because data will be overwritten when the **Overwrite existing data** application option is selected. You can cancel a restore job when it is in progress, but if the **Overwrite existing data** option is selected, data is overwritten even if you cancel the job.
7. Log on to the MongoDB server to which the test restore job is directed.
8. Run the MongoDB system command `ps -ef | grep mongod` to find the temporary recovery MongoDB instance location.
9. Run the MongoDB `mongodump` command to create a dump file of any specific database or collection.

Use the appropriate command. The first command is for a database and the second command is for a collection:

```
mongodump --host <hostname> --port <port> --db <dbname> <dumpfolder>
```

Or,

```
mongodump --host <hostname> --port <port> --collection <collectionname> <dumpfolder>
```

10. Run the **mongorestore** command to restore the dump file into any MongoDB instance. Choose either the original MongoDB instance that the backup was created for, or any alternative instance.

    Use the appropriate command. The first command is for a database and the second command is for a collection:

    ```
    mongorestore --host <hostname> --port <port> --db <dbname> <dumpfolder>\<dbname>
    ```

    Or,

    ```
    mongorestore --host <hostname> --port <port> --collection <collectionname>
    <dumpfolder>\<dbname>
    ```

11. When the database or collection restore operation finishes, go to **Jobs and Operations** > **Active Resources** > **Databases**.
12. Click **Cancel Restore** to end the granular restore procedure.

Backing up and restoring Oracle data

To protect Oracle content, first register the Oracle instance so that IBM Spectrum Protect Plus recognizes it. Then create jobs for backup and restore operations.

Ensure that your Oracle environment meets the system requirements in "Oracle Server database backup and restore requirements" on page 106.

Adding an Oracle application server

When an Oracle application server is added, an inventory of the instances and databases that are associated with the application server is captured and added to IBM Spectrum Protect Plus. This process enables you to complete backup and restore jobs, as well as run reports.

Procedure

To register an Oracle application server, complete the following steps.

1. In the navigation panel, click **Manage Protection** > **Databases** > **Oracle**.
2. Click **Manage Application Servers**.
3. Click **Add Application Server** to add the host machine.
4. In the **Application Properties** pane, enter the host address.

   The host address is a resolvable IP address, or a resolvable path and machine name.
5. Select **User** or **SSH key**.

Option	Description
User	Click this option to specify an existing user or enter a user ID and password. The user must have **sudo** privileges set up. Populate the fields as follows: **Use existing user** Select this check box to use a previously entered user name and password for the application server. Select a user name from the **Select user** list.

Option	Description
	UserID Enter your user name for the application server. If the virtual machine is attached to a domain, the user identity follows the default *domain**name* format. If the user is a local administrator, use the *local_administrator* format. **Password** Enter your password for the application server.
SSH Key	Click this option to use an SSH key. Select a key from the **Select a SSH key** list.

6. To protect multithreaded databases in Oracle 12c and later versions, provide credentials for the databases:
   a) Click **Get databases** to detect and list the Oracle databases on the host server that you are adding.
      Each Oracle database is listed with its name, status, and an indication of whether credentials were previously specified for the database.
   b) For each multithreaded database that you want to protect, click **Set Credential** and specify the user ID and password. Alternatively, you can select an existing user from the **Select user** list.
      You must specify the credentials of an Oracle database user who has SYSDBA privileges.
7. In **Maximum concurrent databases**, set the maximum number of databases to back up concurrently on the server.
   Server performance is impacted when many databases are backed up concurrently, as each database utilizes multiple threads and consumes bandwidth when copying data. Use this option to control the impact on server resources and minimize the impact on production operations.
8. Click **Save**. IBM Spectrum Protect Plus confirms a network connection, adds the application server to the IBM Spectrum Protect Plus database, and then catalogs the instance.
   If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

What to do next

After you add the Oracle application server, complete the following action:

Action	How to
Assign user permissions to the application server.	See “Creating a role” on page 650.

Related concepts

“Managing user access” on page 643
By using role-based access control, you can set the resources and permissions available to IBM Spectrum Protect Plus user accounts.

Related tasks

“Backing up Oracle data” on page 590
Use a backup job to back up Oracle environments with snapshots.

“Restoring Oracle data” on page 593
Use a restore job to restore an Oracle environment from snapshots. IBM Spectrum Protect Plus creates a vSnap clone from the version that is selected during the job definition creation and creates a Network Files System (NFS) share. The IBM Spectrum Protect Plus agent then mounts the share on the Oracle server where the restore job is to be run. For Oracle Real Application Clusters (RAC), the restore job is run on all nodes in the cluster.

Detecting Oracle resources

Oracle resources are automatically detected after the application server is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the application server was added.

Before you begin

IBM Spectrum Protect Plus uses the Oracle instance inventory files that are named `inventory.xml` to locate databases. These files are contained in a `ContentsXML` subdirectory for each instance that is pointed to by the `/etc/oraInst.loc` file. Ensure that the inventory files and the `oraInst.loc` file are updated with the latest changes to your environment. If the inventory files are outdated or corrupted, or if the `oraInst.loc` file does not point to the inventory locations, the associated databases are not be recognized in IBM Spectrum Protect Plus. For more information about managing the Oracle instance inventory, see https://docs.oracle.com/cd/E11882_01/em.112/e12255/oui2_manage_oracle_homes.htm#OUICG143.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **Oracle**.
2. In the list of Oracle instances, select an instance or click the link for the instance to navigate to the resource that you want. For example, if you want to run an inventory job for an individual database in the instance, click the instance link and then select a virtual machine.
3. Click **Run Inventory**.

Testing connection to an Oracle application server

You can test the connection to an Oracle host. The test function verifies communication with the host and tests DNS settings between the IBM Spectrum Protect Plus virtual appliance and the host.

Procedure

To test the connection, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **Oracle**.
2. Click **Manage Application Servers**.
3. In the list of hosts, click **Test** in the **Actions** menu for the host.

Backing up Oracle data

Use a backup job to back up Oracle environments with snapshots.

Before you begin

Review the following information:

- To ensure that file system permissions are retained correctly when IBM Spectrum Protect Plus moves Oracle data between servers, ensure that the user and group IDs of the Oracle users (for example, oracle, oinstall, dba) are consistent across all the servers. Refer to Oracle documentation for recommended uid and gid values.
- If an Oracle Inventory job runs at the same time or short period after an Oracle backup job, copy errors might occur because of temporary mounts that are created during the backup job. As a best practice, schedule Oracle Inventory jobs so that they do not overlap with Oracle backup jobs.
- Avoid configuring log backup for a single Oracle database by using multiple backup jobs. If a single Oracle database is added to multiple job definitions with log backup enabled, a log backup from one job could truncate a log before it is backed up by the next job. This might cause point-in-time restore jobs to fail.

- Avoid scheduling log backups at the same time as an SLA backup job for the same Oracle database. If a log backup occurs at the same time as the backup task of an SLA backup, the SLA backup job may fail. Additionally, ad-hoc backups should not be started if they will run at the same time as scheduled log backups.
- Point-in-time recovery is not supported when one or more data files are added to the database in the period between the chosen point-in-time and the time that the preceding backup job ran.

Take the following actions:

- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see Chapter 19, "Managing user access," on page 643.
- Register the providers that you want to back up. For more information, see "Adding an Oracle application server" on page 588.
- Configure SLA policies. For more information, see "Create backup policies" on page 259.

About this task

During the initial base backup, IBM Spectrum Protect Plus creates a vSnap volume and an NFS share. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus agent mounts the share on the Oracle server where the backup is to be completed.

In the case of Oracle Real Application Clusters (RAC), the backup is completed from any one node in the cluster. When the backup job is completed, the IBM Spectrum Protect Plus agent unmounts the share from the Oracle server and creates a vSnap snapshot of the backup volume.

IBM Spectrum Protect Plus can protect multithreaded databases in Oracle 12c and later versions. For instructions about enabling IBM Spectrum Protect Plus to protect multithreaded databases, see "Adding an Oracle application server" on page 588.

Procedure

To define an Oracle backup job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **Oracle**.
2. Select Oracle homes, databases, and ASM diskgroups to back up. Use the search function to search for available instances.
3. Click **Select an SLA Policy** to add one or more SLA policies that meet your backup data criteria to the job definition.
4. To create the job definition by using default options, click **Save**.

   The job runs as defined by the SLA policies that you selected. To run the job manually, click **Jobs and Operations** > **Schedule**. Select the job and click **Actions** > **Start**.

   Tip: When the job for the selected SLA policy runs, all resources that are associated with that SLA policy are included in the backup operation. To back up only selected resources, you can run an on-demand job. An on-demand job runs the backup operation immediately.

   - To run an on-demand backup job for a single resource, select the resource and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
   - To run an on-demand backup job for one or more resources, click **Create job**, select **Ad hoc backup**, and follow the instructions in "Running an ad hoc backup job" on page 627.
5. To edit options before you create the job definition, click **Select Options**. Set the job definition options.

   Tips for configuring options:

   Review the following tips to help you configure options for the backup job:

   - To set the options for child resources to the same values as the parent, click **Set all options to inherit**.

- If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
- Options that are shown in yellow indicate that the option value has changed from the previously saved value.
- To close the **Options** pane without saving changes, click **Select Options**.

Enable Log Backup

Enable Log Backup must be selected to allow for Oracle point-in-time restore.

Select **Enable Log Backup** to permit IBM Spectrum Protect Plus to automatically create a log backup volume and mount it to the application server. IBM Spectrum Protect Plus then automatically discovers the location of the existing primary archived log and uses cron to configure a scheduled job. The scheduled job completes a transaction log backup from the primary location to that log backup volume at the frequency specified through the **Frequency** setting.

If an on-demand job runs with the **Enable Log Backup** option enabled, log backup occurs. However, when the job runs again on a schedule, the option is disabled for that job run to prevent possible missing segments in the chain of backups.

The **Frequency** can be set to a value independent of the database backup frequency specified in the SLA Policy settings. For example, the SLA Policy may be configured to back up the database once per day while the log backup frequency could be set to once per 30 minutes.

For Oracle RAC, IBM Spectrum Protect Plus mounts the volume and configures the cron job on each of the cluster nodes. When the schedule is triggered, the jobs internally coordinate to ensure that any one active node completes the log backup and the other nodes take no action.

IBM Spectrum Protect Plus automatically manages the retention of logs in its own log backup volume based on the retention settings in the SLA policy.

Choose an option for **Truncate source logs after successful backup**. Select **Never**, the default, if you do not want to truncate source logs after a successful backup job. If this option is selected, archived logs on the primary log destination are not deleted, and Database Administrators must continue to manage those logs using their existing log retention policies. Select **Older than a specified number of days**, **Older than a specified number of hours**, or **Immediately after log backup** if you do want source logs to be truncated after successful backup. Selecting one of these options will automatically delete older archived logs from the database's primary archived log location after a number of days, hours or at the end of every successful database backup.

When the option **Truncate source logs after successful backup** is set to **Older than a specified number of days** or **Older than a specified number of hours**, set the retention of primary logs by entering a day or hour value in the **Primary log retention in days** or **Primary log retention in hours** field. This setting controls the quantity of archived logs that are retained in the primary archived log locations. For example, if **Older than a specified number of days** is selected and **Primary log retention in days** is set to **3**, IBM Spectrum Protect Plus deletes all archived logs older than three days from the primary archived log location at the end of every successful database backup.

Note: Users that upgrade from a previous version of IBM Spectrum Protect Plus that did not select the **Truncate source logs after successful backup** will have that option set to **Never** after upgrading. Similarly, users that opted to select the **Truncate source logs after successful backup** option and entered a number of days for **Primary log retention in days** will have the option set for **Older than a specified number of days** and the associated number of days set in **Primary log retention in days** field after upgrading.

Maximum Parallel Streams per Database

Set the maximum data stream per database to the backup storage. This setting applies to each database in the job definition. Multiple databases can be backed up in parallel if the value of the option is set to **1**. Multiple parallel streams might improve backup speed, but high bandwidth consumption might affect overall system performance.

6. When you are satisfied that the job-specific information is correct, click **Save**.

7. To configure additional options, click the **Policy Options** clipboard icon  icon that is associated with the job in the **SLA Policy Status** section. Set the following additional policy options:

   Pre-scripts and Post-scripts

   Run a pre-script or a post-script. Pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job level. Windows-based machines support Batch and PowerShell scripts while Linux-based machines support shell scripts.

   In the **Pre-script** or **Post-script** section, select an uploaded script and an application or script server where the script will run. To select an application server where the script will run, clear the **Use Script Server** check box. Scripts and script servers are configured through the **System Configuration** > **Script** page.

   To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

   When this option is enabled, if a pre-script or post-script completes processing with a non-zero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes with a non-zero return code, the post-script task status is reported as COMPLETED.

   When this option is disabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.

   Exclude Resources

   Exclude specific resources from the backup job through single or multiple exclusion patterns. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

   Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters as well as the following special characters: - _ and *.

   Separate multiple filters with a semicolon.

   Force Full Backup of Resources

   Force base backup operations for specific virtual machines or databases in the backup job definition. Separate multiple resources with a semicolon.

What to do next

After you create the backup job definition, complete the following action:

Action	How to
Create an Oracle Restore job definition.	See “Restoring Oracle data” on page 593.

Related concepts

“Configuring scripts for backup and restore operations” on page 628
Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Restoring Oracle data

Use a restore job to restore an Oracle environment from snapshots. IBM Spectrum Protect Plus creates a vSnap clone from the version that is selected during the job definition creation and creates a Network Files System (NFS) share. The IBM Spectrum Protect Plus agent then mounts the share on the Oracle server where the restore job is to be run. For Oracle Real Application Clusters (RAC), the restore job is run on all nodes in the cluster.

Before you begin

Complete the following prerequisites:

- Create and run an Oracle backup job. For instructions, see “Backing up Oracle data” on page 590.
- Before an IBM Spectrum Protect Plus user can restore data, the appropriate roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations by using the **Accounts** pane. For instructions, see Chapter 19, “Managing user access,” on page 643.

Review the following restrictions:

- Point-in-time recovery is not supported if one or more data files were added to the database in the period between the chosen point in time and the time that the preceding backup job ran.
- If an Oracle database is mounted but not opened during a backup job, IBM Spectrum Protect Plus cannot determine the database **tempfile** settings that are related to **autoextensibility** and maximum size. When a database is restored from this restore point, IBM Spectrum Protect Plus cannot re-create the **tempfiles** with the original settings because they are unknown. Instead, **tempfiles** are created with default settings, `AUTOEXTEND ON` and `MAXSIZE 32767M`. After the restore job is completed, you can manually update the settings.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.

About this task

The following restore modes are supported:

Instant access mode
In instant access mode, no further action is taken after mounting the share. Users can complete any custom recovery by using the files in the vSnap volume.

Test mode
In test mode, the agent creates a new database by using the data files directly from the vSnap volume.

Production mode
In production mode, the agent first restores the files from the vSnap volume back to primary storage and then creates the new database by using the restored files.

Procedure

To define an Oracle restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **Oracle > Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **Oracle**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

2. On the **Select source** page, take the following actions:
   a) Click a source in the list to show the databases that are available for restore operations. You can also use the search function to search for available instances and toggle the displayed instances through the **View** filter.
   b) Click the plus icon  next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

The selected sources are added to the restore list next to the database list. To remove an item from the list, click the minus icon  next to the item.

c) Click **Next** to continue.

3. On the **Source snapshot** page, select the type of restore job that you want to create:

On-demand: Snapshot
Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

On-demand: Point in Time
Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

Recurring
Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server** The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Restore method** page, set the restore job to run in test, production, or instant access mode by default.

   For test or production mode, you can optionally enter a new name for the restored database.

For production mode, you can also specify a new folder for the restored database by expanding the database and entering a new folder name.

Click **Next** to continue.

After the job is created, it can be run in test, production, or instant access mode in the **Job Sessions** pane.

6. On the **Set destination** page, specify where you want to restore the database and click **Next**.

   Restore to original location
   Select this option to restore the database to the original server.

   Restore to alternate location
   Select this option to restore the database to a local destination that is different from the original server, and then select the alternative location from the list of available servers.

7. On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

   Recovery Options
   Set the following point-in-time recovery options:

   Recover until end of backup
   Restore the selected database to the state at the time that the backup was created.

   Recover until specific point in time
   When log backup is enabled by using an Oracle Backup job definition, point-in-time restore options will be available when you create an Oracle Restore job definition. Select one of the following options, and then click **Save**:

   - **By Time**. Select this option to configure a point-in-time recovery from a specific date and time.
   - **By SCN**. Select this option to configure a point-in-time recovery by System Change Number (SCN).

   IBM Spectrum Protect Plus finds the restore points that directly proceed and follow the selected point in time. During the recovery, the older data backup volume and the newer log backup volume are mounted. If the point in time occurred after the last backup, a temporary restore point is created.

   Application Options
   Set the application options:

   Overwrite existing database
   Enable this option to allow the restore job to overwrite the selected database. By default, this option is not selected.

   Maximum Parallel Streams per Database
   Set the maximum number of parallel data stream from the backup storage per database. This setting applies to each database in the job definition. If the value of the option is set to 1, multiple databases can still be restored in parallel. Multiple parallel streams might improve restore speed, but high bandwidth consumption might affect overall system performance.

   Init Params
   This option controls the initialization parameters that are used to start the recovered database in Oracle test and production workflows.

   Source. This option is the default. IBM Spectrum Protect Plus uses the same initialization parameters as the source database, but with the following changes:

   - Parameters that contain paths such as **`control_files`**, **`db_recovery_file_dest`**, or **`log_archive_dest_*`** are updated to reflect the new paths based on the renamed mount points of the recovered volumes.
   - Parameters such as **`audit_file_dest`** and **`diagnostic_dest`** are updated to point to the appropriate location under the Oracle base directory on the destination server if the path differs from the source server.

- If a new name is specified for the database, the **db_name** and **db_unique_name** parameters are updated to reflect the new name.
- Cluster-related parameters such as **instance_number**, **thread**, and **cluster_database** are set automatically by IBM Spectrum Protect Plus, depending on the appropriate values for the destination.

Target. Customize the initialization parameters by specifying a template file that contains the initialization parameters that are used by IBM Spectrum Protect Plus.

The specified path must point to a plain text file that exists on the destination server and is readable by the IBM Spectrum Protect Plus user. The file must be in Oracle `pfile` format, consisting of lines in the following format:

```
name = value
```

Comments that begin with the # character are ignored.

IBM Spectrum Protect Plus reads the template `pfile` and copies the entries to the new `pfile` that is used to start the recovered database. However, the following parameters in the template are ignored. Instead, IBM Spectrum Protect Plus sets their values to reflect appropriate values from the source database or to reflect new paths based on the renamed mount points of the recovered volumes.

- **control_files**
- **db_block_size**
- **db_create_file_dest**
- **db_recovery_file_dest**
- **log_archive_dest**
- **spfile**
- **undo_tablespace**

Additionally, cluster-related parameters like **instance_number**, **thread**, and **cluster_database** are set automatically by IBM Spectrum Protect Plus, depending on the appropriate values for the destination.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow session overwrite

Select this option to replace an existing database with a database of the same name during recovery. When an Instant Disk Restore is performed for a database and another database with the same name is already running on the destination host or cluster, IBM Spectrum Protect Plus shuts down the existing database before starting up the recovered database. If this option is not selected, the restore job fails when IBM Spectrum Protect Plus detects a running database with the same name.

Continue with restores of other databases even if one fails

Toggle the recovery of a resource in a series if the previous resource recovery fails. If this option is not enabled, the restore job stops if the recovery of a resource fails.

Protocol Priority (Instant access only)

If more than one storage protocol is available, select the protocol to take priority in the job. The available protocols are **iSCSI** and **Fibre Channel**.

Mount Point Prefix

For instant access restore operations, specify the prefix for the path where the mount point is to be directed.

8. Optional: On the **Apply scripts** page, specify scripts that can be run before or after an operation runs at the job level. Batch and PowerShell scripts are supported on Windows operating systems, and shell scripts are supported on Linux operating systems.

   Pre-Script
   Select this check box to choose an uploaded script and an application or script server where the pre-script will run. To select an application server where the pre-script will run, clear the **Use Script Server** check box. Scripts and script servers are configured on the **System Configuration** > **Script** page.

   Post-Script
   Select this check box to choose an uploaded script and an application or script server where the post-script will run. To select an application server where the post-script will run, clear the **Use Script Server** check box. Scripts and script servers are configured on the **System Configuration** > **Script** page.

   Continue job/task on script error
   Select this check box to continue running the job if the script that is associated with the job fails.

   When you select this check box, if a pre-script or post-script completes processing with a nonzero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes processing with a nonzero return code, the post-script task status is reported as COMPLETED.

   If you clear this check box, the backup or restore is not attempted, and the pre-script or post-script task status is reported as FAILED.
9. Take one of the following actions on the **Schedule** page:
   - If you are running an on-demand job, click **Next**.
   - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.

Results

An on-demand job begins after you click **Submit**, and the **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view the progress of the restore operation, expand the job. You can also download the log file by clicking the download icon  .

A recurring job will begin at the scheduled start time when you start the schedule in the **Jobs and Operations** > **Schedule** page.

All running jobs are viewable in the **Jobs and Operations** > **Running Jobs** page.

What to do next

Oracle databases are always restored in non-multithreaded mode. If the databases that you restored were originally in multithreaded mode, after the restore operation is completed, you must manually configure credentials and switch the databases to the multithreaded mode.

Related concepts

“Configuring scripts for backup and restore operations” on page 628
Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

“Adding an Oracle application server” on page 588

When an Oracle application server is added, an inventory of the instances and databases that are associated with the application server is captured and added to IBM Spectrum Protect Plus. This process enables you to complete backup and restore jobs, as well as run reports.

Backing up and restoring SQL Server data

To protect content on a SQL Server server, first register the SQL Server instance so that IBM Spectrum Protect Plus recognizes it. Then create jobs for backup and restore operations.

System requirements

Ensure that your SQL Server environment meets the system requirements in "Microsoft SQL Server database backup and restore requirements" on page 113.

Registration and authentication

Register each SQL Server server in IBM Spectrum Protect Plus by name or IP address. When registering a SQL Server Cluster (AlwaysOn) node, register each node by name or IP address. Note that the IP addresses must be public-facing and listening on port 5985. The fully qualified domain name and virtual machine node DNS name must be resolvable and route-able from the IBM Spectrum Protect Plus appliance.

The user identity must have sufficient rights to install and start the IBM Spectrum Protect Plus Tools Service on the node, including the **Log on as a service** right. For more information about this right, see Add the Log on as a service Right to an Account.

The default security policy uses the Windows NTLM protocol, and the user identity format follows the default *domain**name* format.

When you are using Windows group policy objects (GPO), the group policy object setting, **Network security: LAN Manager** authentication level must be set correctly. Set it with one of the following options:

- `Not Defined`
- `Send NTLMv2 response only`
- `Send NTLMv2 response only. Refuse LM`
- `Send NTLMv2 response only. Refuse LM & NTLM`

Privileges

On the SQL Server server, the system login credential must have public and sysadmin permissions enabled, plus permission to access cluster resources in a SQL Server AlwaysOn environment. If one user account is used for all SQL Server functions, a Windows login must be enabled for the SQL Server server, with public and sysadmin permissions enabled.

Every Microsoft SQL Server host can use a specific user account to access the resources of that particular SQL server instance.

To complete log backup operations, the SQL Server user registered with IBM Spectrum Protect Plus must have the sysadmin permission enabled to manage SQL Server agent jobs.

The Windows Task Scheduler is used to schedule log backups. Depending on a the environment, users may receive the following error: `A specified logon session does not exist. It may already have been terminated.` This is because of a Network access Group Policy setting that needs to be disabled. For more information on how to disable this GPO, please see the following Microsoft Support article: Task Scheduler Error "A specified logon session does not exist"

Adding an SQL Server application server

When an SQL Server application server is added, an inventory of the instances and databases that are associated with the application server is captured and added to IBM Spectrum Protect Plus. This process enables you to complete backup and restore jobs, as well as run reports.

Procedure

Restriction: You can assign only one application server or file server per host. For example, if you register a host as a Microsoft Windows file system, you cannot register the same host as a Microsoft SQL Server or a Microsoft Exchange Server.

To add an SQL Server host, complete the following steps.

1. In the navigation panel, click **Manage Protection** > **Databases** > **SQL**.
2. Click **Manage Application Servers**.
3. Click **Add Application Server**.
4. Populate the fields in the **Application Properties** pane:

   Host Address

   Enter the resolvable IP address or a resolvable path and machine name.

   Use existing user

   Enable to select a previously entered user name and password for the provider.

   UserID

   Enter your user name for the provider. The user identity follows the default *domain\name* format if the virtual machine is attached to a domain. The format *local _administrator* is used if the user is a local administrator.

   Password

   Enter your password for the provider.

   Maximum concurrent databases

   Set the maximum number of databases to back up concurrently on the server. Server performance is impacted when backing up a large number of databases concurrently, as each database utilizes multiple threads and consumes bandwidth when copying data. Use this option to control the impact on server resources and minimize the impact on production operations.
5. Click **Save**. IBM Spectrum Protect Plus confirms a network connection, adds the application server to the IBM Spectrum Protect Plus database, and then catalogs the instance.

   If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

What to do next

After you add the SQL Server application server, complete the following action:

Action	How to
Assign user permissions to the application server.	See “Creating a role” on page 650.

Related concepts

“Managing user access” on page 643
By using role-based access control, you can set the resources and permissions available to IBM Spectrum Protect Plus user accounts.

Related tasks

“Backing up SQL Server data” on page 602

Use a backup job to back up SQL Server environments with snapshots.

"Restoring SQL Server data" on page 607
Use a restore job to restore a Microsoft SQL Server environment from snapshots. After you run IBM Spectrum Protect Plus Instant Disk Restore jobs, your SQL Server clones can be used immediately. IBM Spectrum Protect Plus catalogs and tracks all cloned instances.

Detecting SQL Server resources

SQL Server resources are automatically detected after the application server is added to IBM Spectrum Protect Plus. However, you can run an inventory job to detect any changes that occurred since the application server was added.

Procedure

To run an inventory job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **SQL**.
2. In the list of SQL Server instances, select an instance or click the link for the instance to navigate to the resource that you want. For example, if you want to run an inventory job for an individual database in the instance, click the instance link and then select a virtual machine.
3. Click **Run Inventory**.

Testing the connection to a SQL Server application server

You can test the connection to a SQL Server host. The test function verifies communication with the host and tests DNS settings between the IBM Spectrum Protect Plus virtual appliance and the host.

Procedure

To test the connection, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **SQL**.
2. Click **Manage Application Servers**.
3. In the list of hosts, click **Test** in the **Actions** menu for the host.

Backing up SQL Server data

Use a backup job to back up SQL Server environments with snapshots.

Before you begin

During the initial base backup, IBM Spectrum Protect Plus creates a vSnap LUN volume and creates an NTFS share on that iSCSI LUN. During incremental backups, the previously created volume is reused. The IBM Spectrum Protect Plus agent maps the LUN to the SQL Server server and mounts the NTFS volume to where the backup is completed. If log backups are enabled, IBM Spectrum Protect Plus creates a separate vSnap volume and creates a CIFS on that volume. Log backup transaction files are copied to this share according to the schedule created for log backup.

When the backup job is completed, the IBM Spectrum Protect Plus agent unmounts the share from the SQL Server server and creates a vSnap snapshot of the backup volume.

Review the following information:

- Before an IBM Spectrum Protect Plus user can implement backup and restore operations, roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations through the **Accounts** pane. For more information, see Chapter 19, "Managing user access," on page 643.
- Microsoft iSCSI Initiator must be enabled and running on the Windows server. An iSCSI route must be enabled between the SQL system and vSnap server. For more information, see Microsoft iSCSI Initiator Step-by-Step Guide.

- IBM Spectrum Protect Plus does not support log backup of Simple recovery models.
- Failover of an SQL cluster instance during backup is not supported.
- If you plan to back up a large number of databases, you might have to increase the number of maximum worker threads on each associated SQL Server instance to ensure that backup jobs are completed successfully. The default value for maximum worker threads is 0. The server automatically determines the maximum worker threads value based on the number of processors available to the server. SQL Server uses the threads from this pool for network connections, database checkpoints, and queries. Additionally, a backup of each database requires one additional thread from this pool. If you have a large number of databases in a backup job, the default max worker threads might not be enough to back up all of the databases and the job will fail. For more information about increasing the maximum worker threads option, see Configure the max worker threads Server Configuration Option.
- IBM Spectrum Protect Plus supports database backups and transaction log backups. The product name is populated in the `msdb.dbo.backupset` for records created by backups initiated from IBM Spectrum Protect Plus.
- SQL databases protected by transparent data encryption (TDE) require that a master key and certificate to be created prior to encryption. The master key and certificate are managed by the user outside of IBM Spectrum Protect Plus and are not protected as part of a backup job. Prior to restoring a TDE protected database, the master key and certificate must be restored to the restore destination so that the data protected by IBM Spectrum Protect Plus can be unencrypted. The password used during the initial encryption process must be used to unencrypt the data. For more information about moving SQL databases protected with TDE, see https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/move-a-tde-protected-database-to-another-sql-server?view=sql-server-ver15.
- For more information about log backups for SQL, see "Log backups" on page 605.

Note: Due to limitations with the Volume Shadow Copy Services (VSS) framework, leading spaces, trailing spaces, and unprintable characters should not be used in database names. For more information, see https://support.microsoft.com/en-sg/help/2014054/backing-up-a-sql-server-database-using-a-vss-backup-application-may-fa

Take the following actions:

- Register the SQL Servers that you want to back up. For more information, see "Adding an SQL Server application server" on page 601.
- Configure SLA policies. For more information, see "Create backup policies" on page 259.
- Before you set up and run SQL backup jobs, configure the Shadow Copy storage settings for the volumes where your SQL databases are located. This setting is configured one time for each volume. If new databases are added to the job, the setting must be configured for any new volumes that contain SQL databases. In Windows Explorer, right-click the source volume and select the **Shadow Copies** tab. Set the **Maximum size** to **No limit** or a reasonable size based on the source volume size and I/O activities, and then click **OK**. The shadow copy storage area must be on the same volume or another available volume during backup job.

Procedure

To define an SQL backup job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **SQL**.
2. Select an SQL Server instance to back up.

   Use the search function to search for available instances and toggle the displayed instances through the **View** filter. The available options are **Standalone/Failover Cluster** and **Always On**.
3. Click **Select an SLA Policy** to add one or more SLA policies that meet your backup data criteria to the job definition.
4. To create the job definition by using default options, click **Save**.

   The job runs as defined by the SLA policies that you selected. To run the job manually, click **Jobs and Operations** > **Schedule**. Select the job and click **Actions** > **Start**.

Tip: When the job for the selected SLA policy runs, all resources that are associated with that SLA policy are included in the backup operation. To back up only selected resources, you can run an on-demand job. An on-demand job runs the backup operation immediately.

- To run an on-demand backup job for a single resource, select the resource and click **Run**. If the resource is not associated with an SLA policy, the **Run** button is not available.
- To run an on-demand backup job for one or more resources, click **Create job**, select **Ad hoc backup**, and follow the instructions in "Running an ad hoc backup job" on page 627.

5. Click **Select Options** to specify more options before you save the backup job.

Tips for configuring options:

Review the following tips to help you configure options for the backup job:

- To set the options for child resources to the same values as the parent, click **Set all options to inherit**.
- If multiple resources were selected for the backup job, the options are indeterminate. If you change the value for an option, that value is used for all selected resources after you click **Save**.
- Options that are shown in yellow indicate that the option value has changed from the previously saved value.
- To close the **Options** pane without saving changes, click **Select Options**.

Enable Log Backup

Select this option to enable the backing up of transaction logs. These logs are used for recovery options such as point-in-time restore operations. If log backups are enabled for your backup jobs, transactions are continuously logged during the backup time. Notification is sent if any discontinuity is detected in log file backups.

Restriction: It is possible that the same databases that are on a VM might be backed up as part of a VM backup job and a SQL Server backup job. If you want to enable log backup for a SQL Server backup job, ensure that **Truncate SQL logs** option is not selected for a VM backup job that backs up the same databases. The log truncation deletes all inactive logs from the log file. The deleted log sequence causes discontinuity in the log backup.

To enable log backup schedule creation for multiple databases on the same SQL Server instance, ensure that all databases are added to the same SLA policy. A staging area for the process of log backing up is not required.

If an on-demand job runs with the **Enable Log Backup** option enabled, log backup occurs. However, when the job runs again on a schedule, the option is disabled for that job run to prevent possible missing segments in the chain of backups.

Select one of the following options:

Back up database files one at a time using parallel streams Select this option to use parallel streams to back up your databases sequentially.

Back up database files in parallel using parallel streams Select this option to use parallel streams to backup your databases in parallel.

Finally, set the **Maximum Parallel Streams per Database** by selecting the maximum number of data streams to be used per database during the backing up process. This setting applies to each database in the job definition. Multiple databases can be backed up in parallel if the value of the option is set to **1**. Specifying Multiple parallel streams can improve backup speed in some cases.

6. Click **Save** to save the options for your backup jobs.

   The job runs as defined by your SLA policy, or can be run manually from the **Job and Operations** window.

7. To configure additional options, click the **Policy Options** clipboard icon  icon that is associated with the job in the **SLA Policy Status** section. Set the following additional policy options:

Pre-scripts and post-scripts

Run a pre-script or a post-script. Pre-scripts and post-scripts are scripts that can be run before or after a job runs. Batch and PowerShell scripts are supported.

In the **Pre-script** or **Post-script** section, select an uploaded script and an application or script server where the script is due to run. To select an application server where the script runs, clear the **Use Script Server** check box. Scripts and script servers are configured on the **System Configuration** > **Script** page.

To continue running the job if the script associated with the job fails, select **Continue job/task on script error**.

When this option is enabled, if a pre-script or post-script finishes processing with a nonzero return code, the backup or restore operation is attempted and the pre-script task status is reported as `COMPLETED`. If a post-script completes with a nonzero return code, the post-script task status is reported as `COMPLETED`.

When this option is not enabled, the backup or restore is not attempted, and the pre-script or post-script task status is reported as `FAILED`.

Exclude Resources

Exclude specific resources from the backup job through single or multiple exclusion patterns. Resources can be excluded through an exact match or with wildcard asterisks specified before the pattern (*test) or after the pattern (test*).

Multiple asterisk wildcards are also supported in a single pattern. Patterns support standard alphanumeric characters in addition to the following special characters: - _ and *.

Separate multiple filters with a semicolon.

Force Full Backup of Resources

Force base backups operations for specific virtual machines or databases in the backup job definition. Separate multiple resources with a semicolon.

8. To save any additional options that you configured, click **Save**.

What to do next

After you create the backup job definition, complete the following action:

Action	How to
Create an SQL Restore job definition.	See “Restoring SQL Server data” on page 607.

Related concepts

“Configuring scripts for backup and restore operations” on page 628
Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

“Starting jobs on demand” on page 621
You can run any job on demand, even if the job is set to run on a schedule.

Log backups

Archived log files for databases contain committed transaction data. This transaction data can be used to run a rollforward recovery process as part of a restore operation. Using archive log backups enhances the recovery point objective for your data. Ensure that log backups are enabled in your backup jobs to allow rollforward recovery when you restore Microsoft SQL Server data.

Restriction: It is possible that the same databases that are on a VM might be backed up as part of a VM backup job and a SQL Server backup job. If enable log backup for a SQL Server backup job, ensure that **Truncate SQL logs** option is not selected for a VM backup job that backs up the same databases. The log truncation deletes all inactive logs from the log file. The deleted log sequence causes discontinuity in the log backup.

When you enable log backups for the first time, you must run a backup job for the SLA policy to activate log archiving to IBM Spectrum Protect Plus on the database. This backup creates a separate volume on the vSnap repository, and the volume is mounted persistently on the SQL application server. The volume remains mounted on the SQL application server unless the **Enable Log Backup** option is cleared and a new backup job is run. To enable log backups, follow the instructions in "Backing up SQL Server data" on page 602.

Review the following criteria before you set up log backup operations:

- To run log backups, the SQL Server agent user must be a local Windows administrator. This user must have sysadmin permission to manage SQL Server agent jobs. The agent uses that administrator account to enable and access log backup jobs. For each SQL Server instance, the SQL Server agent user also must be the user of the SQL Server service and the SQL Server agent service account. This rule is true for every SQL Server instance to be protected.
- IBM Spectrum Protect Plus does not support log backup operations for Simple recovery models.
- Avoid configuring log backups for a single SQL database by using multiple backup jobs. Logs are truncated during log backup operations. If a single SQL database is added to multiple job definitions with log backup enabled, a log backup from one job will truncate a log before the next job backs it up. This overlap might cause point-in-time restore jobs to fail.
- Before the logs are copied to the vSnap repository, IBM Spectrum Protect Plus uses the backup folder that is configured for the SQL Server instance as the staging area to collect logs. The volume where this folder is located must have sufficient space to contain the transaction logs between backup jobs. The staging area can be modified by changing the backup folder configuration in SQL Server Management Studio (SSMS).
- IBM Spectrum Protect Plus supports database backups and transaction log backups. The product name is populated in the `msdb.dbo.backupset` for records that are created by backups that are initiated from IBM Spectrum Protect Plus.
- IBM Spectrum Protect Plus automatically truncates post log backups of databases that it backs up. If database logs are not backed up with IBM Spectrum Protect Plus, logs are not truncated and must be managed separately.
- When an SQL backup job is completed with log backups enabled, all transaction logs up to the completion of that job are purged from the SQL Server. Log purging occurs only if the SQL backup job is completed successfully. If log backups are not backed up during a rerun of the job, log purging does not occur.
- A log backup operation for a secondary SQL Server Always On database can fail with the following error:

  ```
  Log backup for database 'DatabaseName' on a secondary replica failed because a
  synchronization point could not be established on the primary database.
  ```

  If this error occurs, change the backup preference of the availability group to `Primary`. Logs are then backed up from the primary replica. After a successful log backup of the primary replica is successfully completed, the backup preference can be changed.
- If a source database is overwritten, all previous transaction logs up to that point are placed in a *condense* directory after the original database is restored. When the next run of the SQL Server backup job is completed, the contents of the condense folder are removed.

Restoring SQL Server data

Use a restore job to restore a Microsoft SQL Server environment from snapshots. After you run IBM Spectrum Protect Plus Instant Disk Restore jobs, your SQL Server clones can be used immediately. IBM Spectrum Protect Plus catalogs and tracks all cloned instances.

Before you begin

Complete the following prerequisites:

- Create and run an SQL backup job. For instructions, see "Backing up SQL Server data" on page 602.
- Before an IBM Spectrum Protect Plus user can restore data, the appropriate roles and resource groups must be assigned to the user. Grant users access to resources and backup and restore operations by using the **Accounts** pane. For instructions, see Chapter 19, "Managing user access," on page 643.
- If you are planning to run a point-in-time recovery, ensure that both the restore target SQL instance service and the IBM Spectrum Protect Plus SQL Server service use the same user account.

Review the following restrictions and considerations:

- If you are planning to run a production restore operation to an SQL Server failover cluster, the root volume of the alternative file path must be eligible to host database and log files. The volume should belong to the destination SQL Server cluster server resource group, and be a dependency of the SQL Server cluster server.
- You cannot restore data to an NTFS or FAT compressed volume because of SQL Server database restrictions. For more information, see Description of support for SQL Server databases on compressed volumes.
- If you are planning to restore data to an alternative location, the SQL Server destination must be running the same version of SQL Server or a later version. For more information, see Compatibility Support.
- When you are restoring data to a primary instance in an SQL Always On Availability Group environment, the database is added to the target Always On database group. After the primary restore operation, the secondary database is seeded by the SQL server in environments where automatic seeding is supported (Microsoft SQL Server 2016 and later). The database is then enabled on the destination availability group. The synchronization time depends on the amount of data that is being transferred and the connection between the primary and secondary replicas.

  If automatic seeding is not supported or is not enabled, a secondary restore from the restore point with the shortest Log Sequence Number (LSN) gap of the primary instance must be completed. Log backups with the latest point-in-time restore point that is created by IBM Spectrum Protect Plus must be restored if the log backup was enabled on the primary instance. The secondary database restore operation is completed in the RESTORING state and you must issue the **T-SQL** command to add the database to the target group. For more information, see https://docs.microsoft.com/en-us/sql/t-sql/language-reference?view=sql-server-2017.
- When restoring from a IBM Spectrum Protect archive, files will be migrated to a staging pool from the tape prior to the job beginning. Depending on the size of the restore, this process could take several hours.
- If you are restoring SQL Server system databases such as the master, MSDB, or model, you must restore using the Instant Access method.

  1. Complete a SQL restore with the **Instant Access** option set for the system databases.
  2. Next, stop the SQL Server instance. For more information about starting and stopping SQL Server services, see https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/start-stop-pause-resume-restart-sql-server-services?view=sql-server-ver15.
  3. Rename all of the data files for the system databases.
  4. Copy all data files from the Instant Access location to the location of the system databases. This is typically located at the path `C:\ProgramData\SPP\mnt\id`. Verify that all data files are copied.
  5. Start the SQL Server instance. Verify that the master database is successfully restored.

About this task

Instant Disk Restore uses the iSCSI protocol to immediately mount LUNs without transferring data. Databases for which snapshots were taken are cataloged and instantly recoverable with no physical transfer of data.

The following restore modes are supported:

Instant access mode
In instant access mode, no further action is taken after mounting the share. Users can complete any custom recovery by using the files in the vSnap volume. An instant access restore of an Always On database is restored to the local destination instance.

Test mode
In test mode, the agent creates a new database by using the data files directly from the vSnap volume.

Production mode
In production mode, the agent first restores the files from the vSnap volume back to primary storage and then creates the new database by using the restored files.

Procedure

To define an SQL restore job, complete the following steps:

1. In the navigation panel, click **Manage Protection** > **Databases** > **SQL**. Click on **Create job**, and then select **Restore** to open the **Restore** wizard.

   Tips:

   - You can also open the wizard by clicking **Jobs and Operations** > **Create job** > **Restore** > **SQL**.
   - For a running summary of your selections in the wizard, click **Preview Restore** in the navigation panel in the wizard.
   - The wizard is opened in the default setup mode. To run the wizard in advanced setup mode, select **Advanced Setup**. With advanced setup mode, you can set more options for your restore job.

2. On the **Select source** page, take the following actions:
   a) Click a source in the list to show the databases that are available for restore operations. You can toggle the displayed sources to show either SQL Server instances in a stand-alone or cluster environment or Always On availability groups by using the **View** filter.

      You can also use the search function to search for databases in the instances or availability groups.

   b) Click the plus icon next to the database that you want to use as the source of the restore operation. You can select more than one database from the list.

      The selected sources are added to the restore list next to the database list. To remove an item from the list source, click the minus icon next to the item.

   c) Click **Next** to continue.

3. On the **Source snapshot** page, select the type of restore job that you want to create:

   On-demand: Snapshot
   Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.

   On-demand: Point in Time
   Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.

   Recurring
   Creates a repeating point-in-time restore job that runs on a schedule.

4. Complete the fields on the **Source snapshot** page and click **Next** to continue.

The fields that are shown depend on the number of items that were selected on the **Select source** page and on the restore type. Some fields are also not shown until you select a related field.

Fields that are shown for an on-demand snapshot, single resource restore

Option	Description
Date range	Specify a range of dates to show the available snapshots within that range.
Backup storage type	All backups in the selected date range are listed in rows that show the time that the backup operation occurred and the service level agreement (SLA) policy for the backup. Select the row that contains the backup time and SLA policy that you want, and then take one of the following actions: • Click the backup storage type that you want to restore from. The storage types that are shown depend on the types that are available in your environment and are shown in the following order: **Backup** Restores data that is backed up to a vSnap server. **Replication** Restores data that is replicated to a vSnap server. **Object Storage** Restores data that is copied to a cloud service or to a repository server. **Archive** Restores data that is copied to a cloud service archive or to a repository server archive (tape). • Click anywhere on the row. The first backup type that is shown sequentially from the left of the row is selected by default. For example, if the storage types **Backup**, **Replication**, and **Archive** are shown, **Backup** is selected by default.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Fields that are shown for an on-demand snapshot, multiple resources restore; point-in-time restore; or recurring restore

Option	Description
Restore Location Type	Select a type of location from which to restore data: **Site** The site to which snapshots were backed up. The site is defined in the **System Configuration** > **Site** pane. **Cloud service** The cloud service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane.

Option	Description
	Repository server The repository server to which snapshots were copied. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane. **Cloud service archive** The cloud archive service to which snapshots were copied. The cloud service is defined in the **System Configuration** > **Backup Storage** > **Object Storage** pane. **Repository server archive** The repository server to which snapshots were copied to tape. The repository server is defined in the **System Configuration** > **Backup Storage** > **Repository Server** pane.
Select a location	If you are restoring data from a site, select one of the following restore locations: **Primary** The primary site from which to restore snapshots. **Secondary** The secondary site from which to restore snapshots. If you are restoring data from a cloud or repository server, select a server from the **Select a location** menu.
Date selector	For on-demand restore operations, specify a range of dates to show the available snapshots within that range.
Restore Point	For on-demand restore operations, select a snapshot from the list of available snapshots in the selected date range.
Use alternate vSnap server for the restore job	If you are restoring data from a cloud service or a repository server, select this box to specify an alternative vSnap server, and then select a server from the **Select alternate vSnap** menu. When you restore data from a restore point that was copied to a cloud service or repository server, a vSnap server is used as a gateway to complete the operation. By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the backup and copy operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

5. On the **Restore method** page, set the restore job to run in test, production, or instant access mode by default.

   For test or production mode, you can optionally enter a new name for the restored database.

   For production mode, you can also specify a new folder for the restored database by expanding the database and entering a new folder name.

   Optionally, for Production and Test restores, in the **New Database Name** field, enter the new name for the restored database. The **New Database Name** field is also displayed when you choose Production restore, but this is for restoring to a new database location on the original instance. When renaming a SQL database, the naming rules for identifiers apply. For more information, see https://docs.microsoft.com/en-us/sql/relational-databases/databases/database-identifiers. When restoring with a new name, the **Global Preferences** option **Rename SQL data and log files when database is restore in production mode with new name** must be enabled. For more information, see "Configuring global preferences" on page 304.

   Click **Next** to continue.

After the job is created, you can run it in test, production, or instant access mode in the **Job Sessions** pane.

6. On the **Set destination** page, specify where you want to restore the database and click **Next**.

   Restore to original instance
   Select this option to restore the database to the original instance.

   Restore to primary instance
   For restore operations in an SQL Always On environment, select this option to restore the database to the primary instance of the Always On Availability Group. The database is added back to the group.

   Restore to alternate instance
   Select this option to restore the database to a local destination that is different from the original instance, and then select the alternative location from the list of available servers.

   For restore operations in an SQL Always On environment in test mode, the source availability database is restored to the selected target instance.

   For restore operations in an SQL Always On environment in production mode, the restored database is added to the target availability group if the destination instance is a primary replica. If the destination instance is a secondary replica of the target availability group, the database is restored to the secondary replica and left in restoring state.

   If the automatic seeding option is enabled for the destination availability group, the secondary database file paths are synchronized with the primary database. If the primary database log is not truncated, the secondary database can be added to the availability group by SQL.

7. On the **Job options** page, configure additional options for the restore job and click **Next** to continue.

   Recovery Options
   Set the following point-in-time recovery options:

   No Recovery
   Set the selected database to a RESTORING state. If you are managing transaction log backups without using IBM Spectrum Protect Plus, you can manually restore log files, and add the database to an availability group, assuming that the LSN of the secondary and primary database copies meets the criteria.

   Restriction: The **No Recovery** option does not support production mode restore operations to SQL Always On groups.

   Recover until end of backup
   Restore the selected database to the state at the time that the backup was created.

   Recover until specific point in time
   When log backup is enabled by using an SQL backup job definition, point-in-time restore options will be available when you create an SQL restore job definition. Select one of the following options:

   - **By Time**. Select this option to configure a point-in-time recovery from a specific date and time.
   - **By Transaction ID**. Select this option to configure a point-in-time recovery by transaction ID.

   Standby mode
   When the Standby mode option is selected, this leaves the SQL database in a read-only state. Uncommitted transactions are undone and saved into an undo file which may subsequently be used for bringing the database online. Transactions stored in the standby file can be applied when the database is ready to be recovered.

   Note: The location of a database restored using Standby mode may be reported to be in the original database location when viewing the database in SQL Management Studio. The location will actually be the directory that is specified by the user for a Production mode restore and the `C:\ProgramData\mnt\uuid_subdirectory` for a Test mode restore.

In a stand-alone restore operation, IBM Spectrum Protect Plus finds the restore points that directly proceed and follow the selected point in time. During the recovery, the older data backup volume and the newer log backup volume are mounted. If the point in time is after the last backup operation, a temporary restore point is created.

When you run restore operations in an SQL Always On environment in test mode, the restored database will join the instance where the availability group resides.

When you run restore operations in an SQL Always On environment in production mode, the restored primary database is joined to the availability group. If the automatic seeding option is enabled for the destination availability group, the secondary database file paths are synchronized with the primary database. If the primary database log is not truncated, the secondary database can be added to the availability group by SQL.

Application Options

Set the application options:

Overwrite existing database

Enable the restore job to overwrite the selected database. By default, this option is not enabled.

Tip: Before you run restore operations in an SQL Always On environment by using the production mode with the **Overwrite existing database** option, ensure that the database is not present on the replicas of the target availability group. To do so, you must manually clean up the original databases (to be overwritten) from all replicas of the target availability group.

Maximum Parallel Streams per Database

Set the maximum number of parallel data streams from the backup storage per database. This setting applies to each database in the job definition. If the value of the option is set to 1, multiple databases can still be restored in parallel. Multiple parallel streams might improve restore speed, but high bandwidth consumption might affect overall system performance.

This option is applicable only when you restore an SQL Server database to its original location using its original database name.

Advanced Options

Set the advanced job definition options:

Run cleanup immediately on job failure

This option enables the automatic cleanup of backup data as part of a restore job if the job fails. This option is selected by default. Do not clear this option unless instructed by IBM Software Support for troubleshooting purposes.

Allow session overwrite

Select this option to replace an existing database with a database of the same name during recovery. When an Instant Disk Restore is performed for a database and another database with the same name is already running on the destination host or cluster, IBM Spectrum Protect Plus shuts down the existing database before starting up the recovered database. If this option is not selected, the restore job fails when IBM Spectrum Protect Plus detects a running database with the same name.

Continue with restores of other databases even if one fails

Toggle the recovery of a resource in a series if the previous resource recovery fails. If this option is not enabled, the restore job stops if the recovery of a resource fails.

Protocol Priority (Instant Access only)

If more than one storage protocol is available, select the protocol to take priority in the job. The available protocols are **iSCSI** and **Fibre Channel**.

Mount Point Prefix

For instant access restore operations, specify the prefix for the path where the mount point is to be directed.

8. Optional: On the **Apply scripts** page, specify scripts that can be run before or after an operation runs at the job level. Batch and PowerShell scripts are supported.

Pre-Script
Select this check box to choose an uploaded script and an application or script server where the pre-script will run. To select an application server where the pre-script will run, clear the **Use Script Server** check box. Scripts and script servers are configured on the **System Configuration** > **Script** page.

Post-Script
Select this option to choose an uploaded script and an application or script server where the post-script will run. To select an application server where the post-script will run, clear the **Use Script Server** check box. Scripts and script servers are configured on the **System Configuration** > **Script** page.

Continue job/task on script error
Select this check box to continue running the job if the script that is associated with the job fails.

When you select this check box, if a pre-script or post-script completes processing with a nonzero return code, the backup or restore operation is attempted and the pre-script task status is reported as COMPLETED. If a post-script completes processing with a nonzero return code, the post-script task status is reported as COMPLETED.

If you clear this check box, the backup or restore operation is not attempted, and the pre-script or post-script task status is reported as FAILED.

9. Take one of the following actions on the **Schedule** page:
   - If you are running an on-demand job, click **Next**.
   - If you are setting up a recurring job, enter a name for the job schedule, and specify how often and when to start the restore job. Click **Next**.
10. On the **Review** page, review your restore job settings and click **Submit** to create the job.

Results

An on-demand job begins after you click **Submit**, and the **onDemandRestore** record is added to the **Job Sessions** pane shortly. To view progress of the restore operation, expand the job. You can also download the log file by clicking the download icon  .

A recurring job will begin at the scheduled start time when you start the schedule in the **Jobs and Operations** > **Schedule** page.

All running jobs are viewable in the **Jobs and Operations** > **Running Jobs** page.

Related concepts

"Configuring scripts for backup and restore operations" on page 628
Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Related tasks

"Adding an SQL Server application server" on page 601
When an SQL Server application server is added, an inventory of the instances and databases that are associated with the application server is captured and added to IBM Spectrum Protect Plus. This process enables you to complete backup and restore jobs, as well as run reports.

"Backing up SQL Server data" on page 602
Use a backup job to back up SQL Server environments with snapshots.

Chapter 16. Protecting IBM Spectrum Protect Plus

Protect the IBM Spectrum Protect Plus application by backing up the catalog. The catalog consists of data such as application configuration settings, SLA policies, registered resources, restore points, backup storage settings, and job information.

Backing up the IBM Spectrum Protect Plus catalog

Define a backup job to protect the IBM Spectrum Protect Plus catalog.

Before you begin

Ensure that an appropriate SLA policy is available. To optimize backup jobs, create SLA policies specifically for backing up IBM Spectrum Protect Plus. To reduce system load, ensure that other jobs are not scheduled to run during the IBM Spectrum Protect Plus backup job. To create an SLA policy, follow the instructions in "Creating an SLA policy for hypervisors, databases, and file systems" on page 324.

An IBM Spectrum Protect Plus catalog can be restored to the same location, or an alternate IBM Spectrum Protect Plus location in disaster recovery scenarios.

Procedure

To back up the IBM Spectrum Protect Plus catalog:

1. In the navigation panel, click **Manage Protection** > **IBM Spectrum Protect Plus** > **Backup**.
2. Select an SLA policy to associate with the IBM Spectrum Protect Plus catalog backup operation.
3. Click **Save** to create the job definition.

Results

The job runs as defined by the SLA policies that you selected, or you can manually run the job by clicking **Jobs and Operations** > **Schedule**. Then, select the job in the **Schedule** tab and click **Actions** > **Start**. For instructions, see "Start a backup job" on page 268.

Restoring the IBM Spectrum Protect Plus catalog

Define a job to restore the IBM Spectrum Protect Plus catalog to the original or an alternate location.

About this task

Attention: An IBM Spectrum Protect Plus catalog restore operation overwrites all data in the IBM Spectrum Protect Plus virtual appliance or alternate virtual appliance location. All IBM Spectrum Protect Plus operations stop while the data is being restored. The user interface is not accessible, and all jobs that are running are canceled. Any snapshots that are created between the backup and restore operations are not saved.

If restoring a cloud backup, the cloud resource or repository server must be registered on the alternate IBM Spectrum Protect Plus location.

When a catalog restore job is started, a job session identifier (ID) is assigned. During the initial phase, the job will be available to be monitored in the IBM Spectrum Protect Plus UI on the job management screen until the recovery step initiates the internal database restore. Once the job enters this state, IBM Spectrum Protect Plus is no longer available. During this phase, log information is written to the location: `/data/log/catalogprotection/managedb-catalogrestore-`*`time`*`.log`, where *time* is epoch time. Data contained in this log is relates to the restore of the mongo configuration and recovery catalog. After the process is complete, the `virgo` service will start and the data is written to the virgo log. The IBM Spectrum Protect Plus user interface is again accessible, but the complete loading of job

© Copyright IBM Corp. 2017, 2021

logs is delayed which results in them not being immediately visible in the IBM Spectrum Protect Plus user interface. An alert is generated after the job log recovery is completed.

Procedure

To restore the IBM Spectrum Protect Plus catalog:

1. In the navigation panel, click **Manage Protection** > **IBM Spectrum Protect Plus** > **Restore**.
2. Select a vSnap server, cloud resource, or repository server.

   Data can be restored to the same location, or an alternate location in disaster recovery scenarios.

   Available snapshots for the server are displayed.
3. Click **Restore** for the catalog snapshot that you want to restore.
4. Select one of the following restore modes:

   Restore the catalog and suspend all scheduled jobs
   The catalog is restored and all scheduled jobs are left in a suspended state. No scheduled jobs are started, which allows for the validation and testing of catalog entries and the creation of new jobs. Typically, this option is used in DevOps use cases.

   Restore the catalog
   The catalog is restored and all scheduled jobs continue to run as captured in the catalog backup. Typically, this option is used in disaster recovery.
5. Click **Restore**.
6. To run the restore job, in the dialog box, click **Yes**.

Managing IBM Spectrum Protect Plus restore points

You can use the **Restore Point Retention** pane to search for restore points in the IBM Spectrum Protect Plus catalog by backup job name, view their creation and expiration dates, and override the assigned retention.

Related concepts

"Job types" on page 619

Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Expiring job sessions

You can expire a job session to override the snapshot retention settings that were assigned during backup creation.

About this task

Expiring a job session will not remove a snapshot and related recovery point if the snapshot is locked by a replication or copy relationship. Run the replication or copy-enabled job to change the lock to a later snapshot. The snapshot and recovery point will be removed during the next run of the maintenance job.

Procedure

To set a job session to expire:

1. In the navigation panel, click **Manage Protection** > **IBM Spectrum Protect Plus** > **Restore Point Retention**.
2. On the Backup Sessions tab, search for the job session or restore point. Alternatively, on the Virtual Machines / Databases tab, select either Applications or Hypervisors to search for the desired catalog entry by entering the name. Names can be searched by entering partial text, using the asterisk (*) as a wildcard character, or using the question mark (?) for pattern matching.

For more information about using the search function, see Appendix A, “Search guidelines,” on page 681.

3. If you are searching from the Backup Sessions tab, use filters to fine-tune your search across job types and date range when the associated backup job started.
4. Click the search icon .
5. Select the job sessions that you want to expire.
6. From the **Actions** list, select one of the following options:
   - **Expire** is used to expire a single job session.
   - **Expire All Job Sessions** is used to expire all unexpired job sessions for the selected job.

   Note: When IBM Cloud Object Storage with a WORM policy is the destination provider type, the **Expire** and **Expire All Job Sessions** options are not listed in the **Actions** menu. In this case, IBM Spectrum Protect Plus does not control the retention and it is instead controlled by the provider.
7. To confirm the expiration, in the dialog box, click **Yes**.

Deleting resource metadata from the IBM Spectrum Protect Plus catalog

When you run an inventory job, resources are added to the IBM Spectrum Protect Plus catalog. To release space in the catalog, you can expire the metadata from the restore points that are associated with the resources.

About this task

Expiring a resource from the catalog does not remove associated snapshots from a vSnap server or secondary backup storage.

Procedure

To expire a resource from the catalog:

1. In the navigation panel, click **Manage Protection** > **IBM Spectrum Protect Plus** > **Restore Point Retention**.
2. Click the **Virtual Machines/Databases** tab.
3. Use the filter to search by resource type, and then enter a search string to search for a resource by name.

   For more information about using the search function, see Appendix A, “Search guidelines,” on page 681.
4. Click the search icon .
5. Click the delete icon that is associated with a resource.
6. To confirm the expiration, in the dialog box, click **Yes**.

Results

The catalog metadata that is associated with the resource is removed from the catalog.

Related concepts

“Job types” on page 619

Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Chapter 17. Managing jobs and operations

You can manage and monitor jobs in the **Jobs and Operations** window. You can also configure scripts to run before or after jobs.

Job types

Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Backup and restore jobs are user defined. After you create these jobs, you can modify the jobs at any time. Maintenance, inventory, and report jobs are predefined and not modifiable. However, you can modify the schedules of maintenance, inventory, and report jobs.

You can run all jobs on demand, even if they are set to run on a schedule. You can also hold and release jobs that are set to run on a schedule.

The following job types are available:

Backup
A backup job defines the resources that you want to back up and the service level agreement (SLA) policy or policies that you want to apply to those resources. Each SLA policy defines when the job runs. You can run the job by using the schedule that is defined by the SLA policy or you can run the job on demand.

You can also run backup jobs for a single resource or multiple selected resources that are associated with an SLA policy rather than backing up all resources that are associated with the policy.

The job name is auto generated and is constructed of the resource type followed by the SLA policy that is used for the job. For example, a backup job for SQL Server resources that are associated with the SLA policy Gold is `sql_Gold`.

Restore
A restore job defines the restore point that you want to restore data from. For example, if you are restoring hypervisor data, the restore point might be a virtual machine. If you are restoring application data, the restore point might be a database.

Restore jobs are ran on a schedule or on demand.

For scheduled jobs, the job name is defined by the user who creates the job.

For on-demand jobs, the job name `onDemandRestore` is auto generated when the job is run. On-demand restore jobs are removed from job history when the job is 3 days old and it is not in a pending or running state, and when it is not in **Active Resources**.

Maintenance
The maintenance job runs once a day to remove resources and associated objects that are created by IBM Spectrum Protect Plus when a job that is in a pending state is deleted.

The cleanup procedure reclaims space on storage devices, cleans up the IBM Spectrum Protect Plus catalog, and removes related snapshots. The maintenance job also removes cataloged data that is associated with deleted jobs.

The job name is `Maintenance`

Inventory
An inventory job is run automatically when you add a resource to IBM Spectrum Protect Plus. However, you can run an inventory job at any time to detect any changes that occurred since the resource was added.

The inventory job names are `Default Application Server Inventory`, `Default Hypervisor Inventory`, and `Default Storage Server Inventory`.

Report
A report job runs a scheduled report. The job name is the report name preceded by `Report_`.

© Copyright IBM Corp. 2017, 2021

Report names are similar to the following example:

```
Report_VM Backup History
```

Related concepts

"Protecting virtualized systems" on page 335
You must register the virtualized systems that you want to protect in IBM Spectrum Protect Plus and then create jobs to back up and restore the resources that are associated with the systems.

"Protecting databases" on page 497
You must register the database applications that you want to protect in IBM Spectrum Protect Plus and then create jobs to back up and restore the databases and resources that are associated with the applications.

Related tasks

"Creating an SLA policy for hypervisors, databases, and file systems" on page 324
You can create custom service level agreement (SLA) policies to define backup frequency, retention, replication, and copy policies that are specific for your environment.

"Running an ad hoc backup job" on page 627
With an ad hoc backup job, you can select one or more resources that are associated with an SLA policy and run an on-demand backup operation for those resources.

Creating jobs and job schedules

The method for creating jobs and job schedules depends on the job type.

You can create jobs and schedules for backup and restore jobs. The following table describes the available backup and restore jobs and provides links to the steps that are required to create the jobs and job schedules or run the jobs on demand.

Maintenance jobs are created by default. Inventory and report jobs are created automatically when an inventory operation runs or when a report is scheduled.

Job type	Description	How to create the job
Backup	You can create a job definition and assign one or more service level agreement (SLA) policies to that definition. The job definition defines the resources to back up and the SLA policy defines the schedule, targets, and other options for the backup operation.	See the topics that contain instructions for backing up data by resource type in the following sections: • Chapter 11, "Protecting virtualized systems," on page 335 • Chapter 12, "Protecting file systems," on page 383 • Chapter 13, "Protecting containers," on page 401 • Chapter 14, "Protecting data on cloud systems," on page 491 • Chapter 15, "Protecting databases," on page 497 For example, the backup topic for VMware is "Backing up VMware data" on page 340.

Job type	Description	How to create the job
Ad hoc backup	When a job is run for the selected SLA policy, all resources that are associated with that SLA policy are included in the backup operation. If you want to back up only selected resources by using a selected SLA policy, you can run an ad hoc job, which runs the backup operation immediately.	See "Running an ad hoc backup job" on page 627.
Restore	After you have run a backup job at least once, you can run a restore job to restore the data. You can create a restore job that runs on a schedule or that runs on demand.	See the topics that contain instructions for restoring data by resource type in the following sections: • Chapter 11, "Protecting virtualized systems," on page 335 • Chapter 12, "Protecting file systems," on page 383 • Chapter 13, "Protecting containers," on page 401 • Chapter 14, "Protecting data on cloud systems," on page 491 • Chapter 15, "Protecting databases," on page 497 For example, the restore topic for VMware is "Restoring VMware data" on page 350.

Related concepts

"Job types" on page 619
Jobs are used to run backup, restore, maintenance, inventory, and report operations in IBM Spectrum Protect Plus.

Related tasks

"Creating an SLA policy for hypervisors, databases, and file systems" on page 324
You can create custom service level agreement (SLA) policies to define backup frequency, retention, replication, and copy policies that are specific for your environment.

Starting jobs on demand

You can run any job on demand, even if the job is set to run on a schedule.

Procedure

Complete the following steps to start a job:

1. In the navigation panel, click **Jobs and Operations**, and click the **Schedule** tab.
2. Choose the job that you want to run, click the actions menu icon ⋮, and then click **Start**.
   The job is started and added to the **Running Jobs** tab.

What to do next

To view the job log for the job, select the job on the **Running Jobs** tab and click **Job Log**. To download the log for the job, click **Download.zip**.

To view all jobs that are running or ran concurrently with the job, click **Concurrent Jobs**.

Viewing jobs

View information about the jobs that are running in your environment, the job history, the active resources that are associated with restore jobs, and scheduled jobs.

About this task

Jobs are grouped on the following job pages:

Running Jobs
This page shows jobs that are running.

Job History
This page shows jobs that ran successfully, failed, or completed processing with warnings.

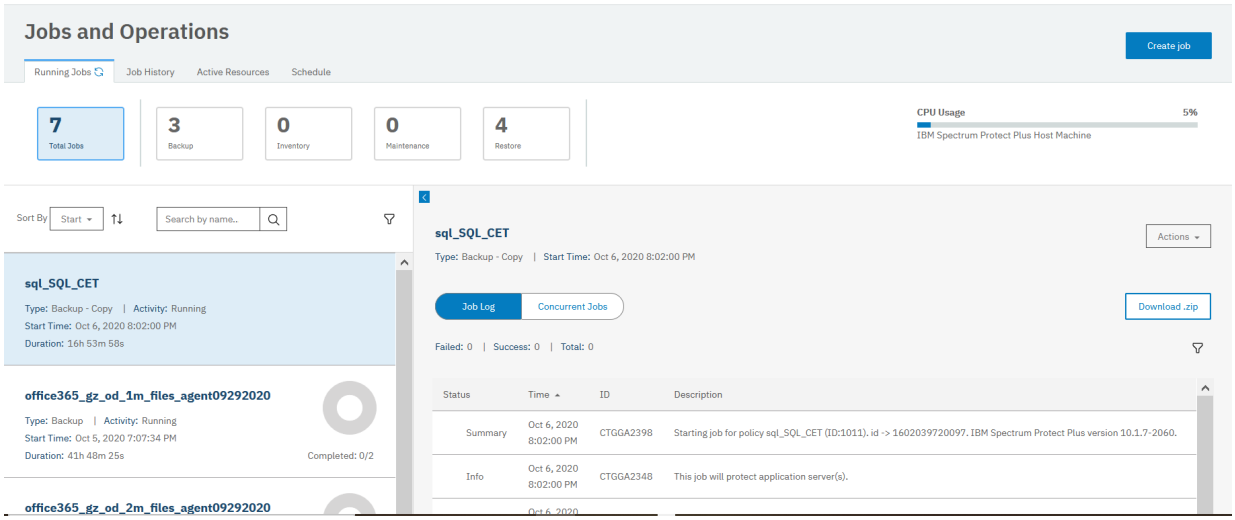
Active Resources
This page shows active resources that are associated with a restore job. An example of an active resource is a file system or database that is mounted as part of a restore operation.

Schedule
This page shows scheduled jobs.

Procedure

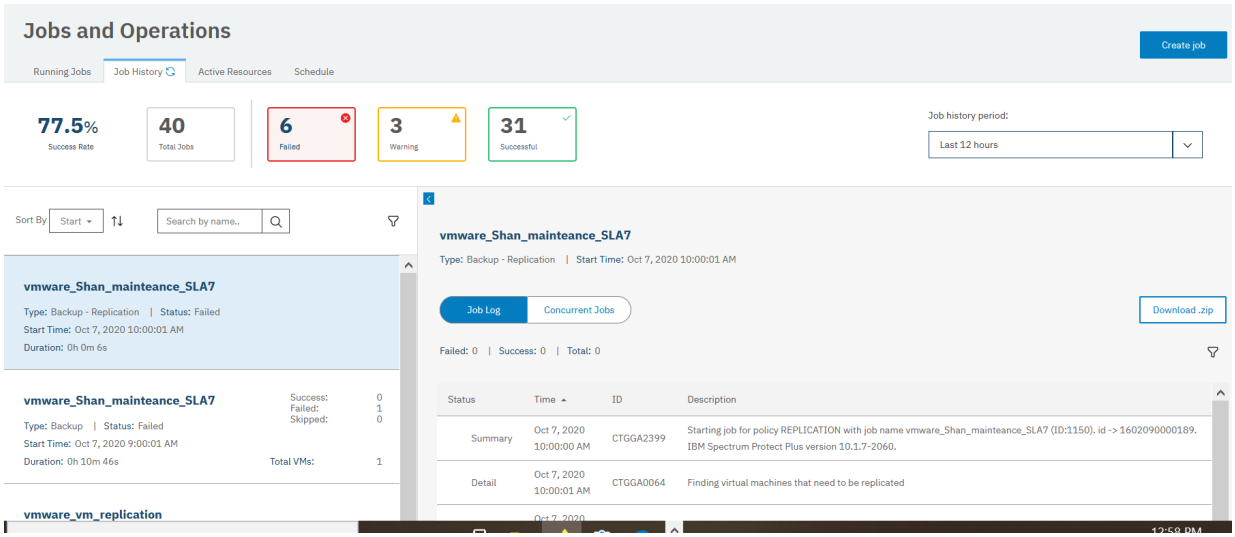
To view jobs, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**.
2. On the **Running Jobs** page, view the status of the jobs that are currently running, as shown in the following example.



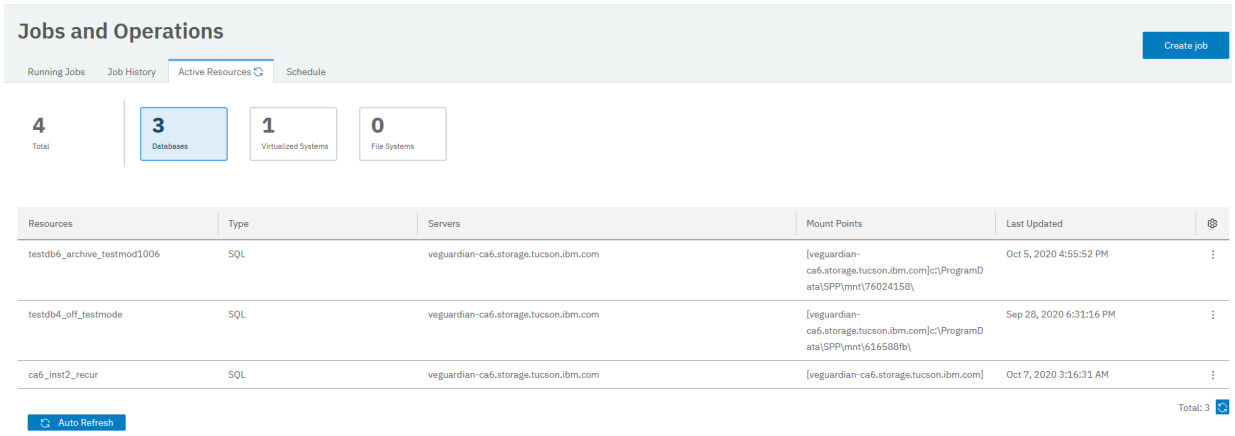
3. To view completed jobs, click **Job History**.

   The ribbon across this screen shows the status of historical jobs. Use the filter to define the duration of the job history to display.



4. To view the active resources in your environment, click **Active Resources**, and then click **Databases**, **Virtualized Systems**, or **File Systems** to view the active resources by resource type.

   To customize the columns that are displayed on each resource type, click the settings icon to select the columns.



5. To view scheduled jobs, click **Schedule**.
   You can complete the following actions for scheduled jobs:
   - Start or pause a job by selecting the job and clicking **Start** or **Pause Schedule**.
   - Edit some recurring and maintenance job schedules by selecting the job and clicking the schedule icon.
   - Edit restore jobs by selecting a job and clicking click the edit icon.
   - Customize the columns that are displayed for the job table by clicking the settings icon to select the columns.
   - Filter the jobs by job type by using the filter icon to select the job types that you want. For example, if you want to see only backup and restore jobs, select the **Backup** and **Restore** checkboxes and clear the others.
6. Optional: To download a job log and other files that reflect the information that is shown on the **Jobs and Operations** window, click **Download.zip**.

Viewing backup job progress at the resource level

View the status of individual resources in a backup job. Viewing the job at the resource level enables you to determine the backup performance of each resource. This feature provides information to help you to optimize backup performance and resolve possible issues.

About this task

This feature is available only for backup jobs. The progress of individual resources is not shown for other job types.

Procedure

To view the progress of individual resources in a backup job, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**.
2. Click **Running Jobs** for jobs that are in progress or **Job History** for jobs that are complete.
3. Select the job that contains the resources that you want to view, and then click **Progress**.

   Information about each resource is shown in a table. This information includes the progress of the backup operation for each resource in the **Overall Progress** column.

   If applicable for the resource type, the task that is running for the backup operation is also shown in the **Task Progress** column. This column is not included for some resource types, such as hypervisors, whose backup operations do not include individual tasks.

   The following example shows the progress information for a Kubernetes backup job. In this example, the overall backup progress for the resource is 60% as shown in the **Overall Progress** column. The current backup task that is running, data transfer, is shown in the **Task Progress** column. The table was expanded by clicking the twistie .

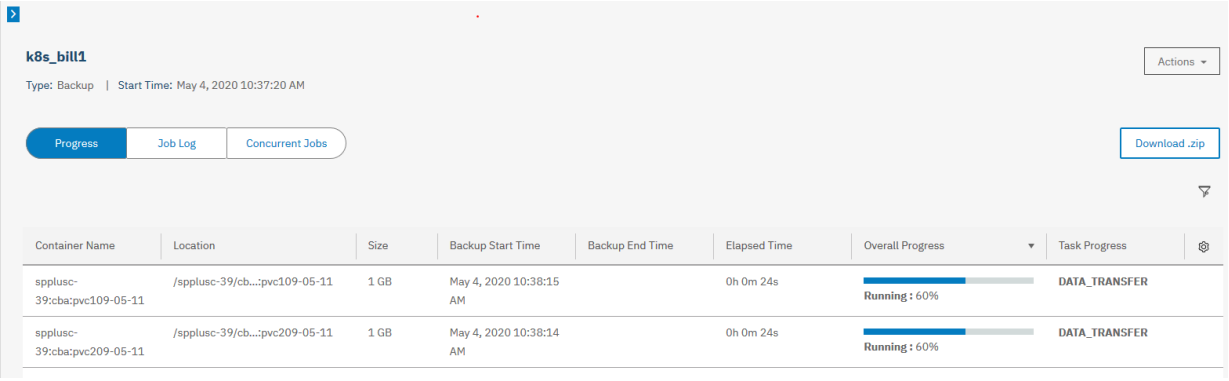


Figure 57. Viewing job information at the resource level

4. Optional: You can customize the columns that are shown in the table and filter the resources that are shown by progress status.

   To customize the columns, click the settings icon to select the columns. By default, all columns are shown.

   To filter the resources by progress status, click the filter icon and select the status values that you want. For example, if you want to see only resources that are in the process of running, select the **Running** checkbox and clear the others.

Viewing job logs

For each job run, a log is provided that shows such information as the status of the job, the start and end time for the job, and a message that is associated with the job.

Procedure

To view job logs, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**
2. Click **Running Jobs** for jobs that are in progress or **Job History** for jobs that are complete.
3. Select a job, and click **Job Log**.

   The job log for the selected job is shown.

Viewing concurrent jobs

Jobs that overlap other jobs are referred to as concurrent jobs. You can view jobs that are running or ran concurrently with another job.

Procedure

To view jobs, that are running or ran concurrently with another job, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**
2. Click **Running Jobs** for jobs that are in progress or **Job History** for jobs that are complete.
3. Select a job, and click **Concurrent Jobs**.

   For jobs that are shown on the **Running Jobs** tab, a list of all jobs that are running concurrently with the selected job is shown. For jobs that are shown on the **Job History** tab, a list of all jobs that ran concurrently with the selected job are shown.

   Restriction: Multiple backup jobs cannot back up the same resource at the same time. If multiple jobs share a resource or resources, the job that processes the resource first will run and any other jobs that start during the same time period will fail.

Pausing and resuming jobs

You can pause and resume a scheduled job. When you pause a scheduled job, the job will not run until it is resumed.

Procedure

To pause and release job schedules, complete the following steps:

1. In the navigation panel, click **Jobs and Operations**, and click the **Schedule** tab.
2. Choose the job that you want to pause, and click the actions menu icon ⋮, and then click **Pause Schedule**.
3. To resume the job schedule, click ⋮, and then click **Release Schedule**.

Editing jobs and job schedules

You can edit the job options and schedule for some job types.

About this task

For restore jobs, you can edit the job options by using the **Restore** wizard.

For the following job types, you can edit the job schedule:

- Restore (recurring jobs)
- Inventory
- Report
- Maintenance

Procedure

To edit a job or a job schedule, complete the following steps:

1. In the navigation panel, click **Jobs and Operations** and then click the **Schedule** tab.
2. Click the edit or schedule icon.

Option	Description
	Click this edit icon to open the **Restore** wizard and change the options for the job. Follow the instructions for using the wizard in the applicable resource restore topic in Chapter 11, “Protecting virtualized systems,” on page 335 and Chapter 15, “Protecting databases,” on page 497.
	Click this edit icon to change the job schedule.

Canceling jobs

You can cancel a job that is running.

Procedure

To cancel a job, complete the following steps:

1. In the navigation panel, click **Jobs and Operations** and then click the **Running Jobs** tab.
2. Click the **Actions** menu that is associated with the job, and then click **Cancel**.

Deleting jobs

You can delete a restore or report job that has a status of IDLE.

About this task

This procedure applies only to restore and report jobs. To delete a backup job, you must delete the service level agreement (SLA) policy that is associated with that job.

Procedure

To delete a restore or report job, complete the following steps:

1. In the navigation panel, click **Jobs and Operations** and then click the **Schedule** tab.
2. Click the delete icon that is associated with the job.

Rerunning partially completed backup jobs

If the last instance of a backup job was partially completed, you can rerun the job to back up virtual machines and databases that were skipped.

About this task

A backup job can be rerun only in the same session ID as the original partially completed backup job. No successful backup of the same resource can have completed since the partial backup job you choose to rerun.

Tip: Backup jobs can be rerun only in response to a hypervisor or database backup failure. The following events do not qualify for backup job rerun operations:

- A VM backup was completed with an FLI failure.
- A snapshot condense failure occurred for a storage system.
- A backup job failed with an unknown issue such as a cataloging error.
- A resource is missing from the vCenter.

For applications for which log backups are supported, log backups are not disabled when using the rerun feature. Log backups will be disabled for the applicable databases when the job is next started without using the on-demand backup or rerun feature.

Procedure

Complete the following steps to rerun a partially completed backup operation:

1. In the navigation panel, click **Jobs and Operations** and then click the **Job History** tab.
2. Use the search function and filters to find the last instance of the backup job that was partially completed.
3. Select the job instance and then click **Rerun**.

   If the backup job cannot be rerun, the **Rerun** option is not available.

Results

All SLA options and any exclusions that are associated with the original job are included in the rerun operation. Any option or exclusion changes that you applied after the last partial backup operation are ignored. If the rerun job is completed successfully, the job summary is updated to show success.

Running an ad hoc backup job

With an ad hoc backup job, you can select one or more resources that are associated with an SLA policy and run an on-demand backup operation for those resources.

About this task

This feature associates the selected SLA policy and resources in an ad hoc job for the purposes of running an immediate, on-demand backup operation. It does not change SLA policy assignments for resources that are associated with scheduled jobs.

You can disable log backups using an ad hoc job for all application workloads that normally allow for the enablement of option. This is useful when an SLA policy has log backup enabled for several databases, but one or more specific databases must have a log backup disabled. After the backup job completes without log backup enabled, the log backup option can be re-enabled using an ad hoc backup job. For example, when attempting to clone an SQL database from an SQL Always-On Availability Group primary instance to secondary instance, the log backup on the database being cloned must be disabled prior to backing up that database. This must be done so that the cloned database in the secondary instance can be brought up in a synchronized state.

- Disable the option for log backups on the database that is to be cloned. Click on **Manage Protection** > **Databases**. Select the appropriate workload type.
- If necessary, change the **View** setting to make Instances visible. Select the database that is to be cloned.
- Click on **Select Options**. Deselect the **Enable Log Back up** option and then click on **Save**.
- Run an ad hoc backup using the procedure in this topic for the database.
- After the backup completes, clone the database to the secondary instance.
- Upon completion of the clone operation, re-enable the log backup option for the database using the same process.
- Finally, run a back-up job of the SLA that contains database.

Procedure

To run an ad hoc backup job, complete the following steps:

1. In the navigation panel, click **Jobs and Operations** > **Create Job**.
2. Select **Ad hoc backup** to open the backup wizard.

   Tips:

   - You can also open the wizard from the individual hypervisor or application management pages by clicking **Manage Protection** > **Hypervisors** or **Manage Protection** > **Applications**.
   - For a running summary of your selections in the wizard, click **Preview Backup** in the navigation panel in the wizard.
3. On the **Source type** page, click the hypervisor or application for the resources that you want to include in the job.
4. On the **Select SLA policy** page, select the SLA policy and then click **Next**.
5. On the **Select source** page, take the following actions:
   a) Review the available resources.

      You can enter all or part of a name in the filter box to locate resources that match the search criteria. You can use the wildcard character (*) to represent all or part of a name. For example, `vm2*` represents all resources that begin with "vm2".

   b) Click the plus icon next to the resource that you want to add to the job.

      To remove a resource from the list, click the minus icon next to the resource.

   c) Click **Next**.
6. On the **Review** page, review the job settings and then click **Submit** to create and run the job.

What to do next

To view the status and other information about the job, click **Jobs and Operations** in the navigation panel and click the job on the **Running Jobs** tab.

Configuring scripts for backup and restore operations

Prescripts and postscripts are scripts that can be run before or after backup and restore jobs run at the job level. Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts are created locally, uploaded to your environment through the **Script** page, and then applied to job definitions.

Before you begin

Review the following considerations for using scripts with hypervisors:

- The user who is running the script must have the **Log on as a service** right enabled, which is required for running prescripts and postscripts. For more information about this right, see Add the Log on as a service Right to an Account.
- Windows Remote Shell (WinRM) must be enabled.

Uploading a script

Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts must be created using the associated file format for the operating system.

Procedure

Complete the following steps to upload a script:

1. In the navigation panel, click **System Configuration** > **Script**.
2. In the **Scripts** section, click **Upload Script**.
   The **Upload Script** pane is displayed.
3. Click **Browse** to select a local script to upload.
4. Click **Save**.
   The script is displayed in the **Scripts** table and can be applied to supported jobs.

What to do next

After you upload the script, complete the following action:

Action	How to
Add the script to a server from which it will run.	See “Adding a script to a server” on page 629.

Adding a script to a server

You can add a script to the server from which the script will run.

Procedure

Complete the following steps to add a script to a server:

1. In the navigation panel, click **System Configuration** > **Script**.
2. In the **Script Servers** section, click **Add Script Server**.
   The **Script Server Properties** pane displays.
3. Set the server options.

   Host Address

   Enter the resolvable IP address or a resolvable path and machine name.

   Use existing user

   Enable to select a previously entered user name and password for the provider.

   Username

   Enter your username for the provider. If entering a SQL server, the user identity follows the default *domain**name* format if the virtual machine is attached to a domain. The format *local_administrator* is used if the user is a local administrator.

   Password

   Enter your password for the provider.

   OS Type

Select the operating system of the application server.

4. Click **Save**.

Chapter 18. Managing reports and logs

IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Types of reports

You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

Reports are based on the data that is collected by the most recent inventory job. You can generate reports after all cataloging jobs and subsequent database condense jobs are completed. You can run the following types of reports:

- Backup storage utilization reports
- Protection reports
- System reports
- Virtual machine environment reports

Reports include interactive elements, such as searching for individual values within a report, vertical scrolling, and column sorting.

Backup storage utilization reports

IBM Spectrum Protect Plus provides backup storage utilization reports that display the storage utilization and status of your backup storage, such as vSnap servers.

To view backup storage utilization reports, complete the following steps:

1. In the navigation panel, click **Reports and Logs** > **Reports**.
2. Click on the **Reports** tab.
3. Select **Backup Storage Utilization** in the **Filter by category** drop-down menu.
4. Run the report by clicking the **Run Report** ( ) icon beside the desired report.

The following reports are available:

VM Backup Utilization report

Virtual machines may be narrowed through the use of the **Hypervisor type**, **Hypervisor**, and **VM tags** selection boxes. The default value is **All**, which shows data for all VM backups.

The VM Backup Utilization report includes the VM name, its location, the hypervisor type, the SLA policy that is used to protect the VM, and the location of the backup storage used. This back storage may be the host name or IP address of a disk, the name of the cloud server, or the name of the repository server. The backup size of each VM, and the number of recovery points that are available for each VM that is displayed. Finally, the total number of virtual machines protected appears at the bottom of the report. The **Search** box may be used to further filter report results.

vSnap Storage Utilization report

Use the report options to filter specific vSnap servers to display through the **vSnap Storage** selection box. To filter out replica destination volumes, select **Exclude Replica Destination Volumes**. For a detailed view of the individual virtual machines and databases that are protected on each vSnap server, select **Show Resources protected per vSnap Storage**. This area of the report displays the names of the virtual machines, associated hypervisor, location, and the compression and deduplication ratio of the vSnap server.

© Copyright IBM Corp. 2017, 2021

The vSnap Storage Utilization report displays the vSnap servers, the site, status, total space, free space, and used space. When expanded, the deduplication and compression ratios, if applicable, are displayed for each vSnap server. The vSnap Storage Utilization report displays both an overview of your vSnap servers and a detailed view of the individual virtual machines and databases that are protected on each vSnap server. The **Search** box may be used to further filter report results.

Note: Storage capacity and usage values that are displayed by IBM Spectrum Protect Plus might vary between those that appear on the dashboard versus those that appear on the vSnap Storage Utilization report. The dashboard displays live information, while the report reflects data from the last inventory job run. Variations are also due to differing rounding algorithms.

Related concepts

"Report actions" on page 638
You can run, save, or schedule reports in IBM Spectrum Protect Plus.

"Types of reports" on page 631
You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

Protection reports

IBM Spectrum Protect Plus provides reports that display the protection status of your resources. By viewing the reports and taking any necessary action, you can help to ensure that your data is protected through user-defined recovery point objective parameters.

To view protection reports, complete the following steps:

1. In the navigation panel, click **Reports and Logs** > **Reports**.
2. Click on the **Reports** tab.
3. Select **Protection** in the **Filter by category** drop-down menu.
4. Run the report by clicking the **Run Report** ( ) icon beside the selected report.

The following reports are available:

Container Persistent Volume Backup History report

The Container Persistent Volume Backup History report displays the history of persistent container volume back jobs. Use the report options to filter by Persistent Volume Claim (PVC) type and to select specific **PVCs** to display. The report can be further filtered by failed jobs or successful jobs in the **Status** field and by specific service level agreement (SLA) policies using the **SLA Policy** field. Set an integer value in the **Backup History for Past Number of Days** field to show the backup history for a specified number of days.

Database Backup History report

Run the Database Backup History report to review the protection history of specific databases. To run the report, at least one database must be specified in the **Databases** option. You can select multiple databases. Use the report options to filter **Status** by failed or successful jobs. The report can be further filtered by specific service level agreement (SLA) policies using the **SLA Policy** field. An integer value can be specified for the **Backup History for Past Number of Days** field to limit results.

In the detail view of the report, expand an associated job to view further job details, such as the reason why a job failed or the size of a successful backup. The **Search** box may be used to further filter report results.

Database SLA Policy RPO Compliance report

Use the report options to filter by **Application Type** and to select a specific **Application Server** to display. The report can be further filtered by databases that are in compliance or not in compliance with the defined RPO through the **Display Databases That Are** field, or by **Protection Type**, including data that was backed up to vSnap, using replication, using object storage copy, or using archive.

The Database SLA Policy RPO Compliance report displays databases in relation to recovery point objectives as defined in SLA policies. The quick view displays a pie chart of a count of backups to

vSnap that are in compliance and those that are not in compliance. The summary view displays the SLA policy, the SLA schedule, the number of backups to vSnap that are in compliance and the number that are not in compliance, and the replications that are in compliance and not in compliance. Also displayed are databases not in compliance for the protection types which includes the database names, application servers, application types, the last successful protection time, and the out of compliance reason.

File System Backup History report

Run the File System Backup History report to review the protection history of specific file systems. To run the report, at least one server must be specified in the **Server** option and one file system must be selected for the **File System** option. Use the report options to filter **Status** by failed or successful jobs. The report can be further filtered by specific service level agreement (SLA) policies using the **SLA Policy** field. The default setting for all four options is **All**. An integer value can be specified for the **Backup History for Past Number of Days** field to limit results.

Report properties displays the creation date and the account that was used to generate the report. Also included are the report filters used when the report was generated. In the detail view of the report, the file system is listed with the server and the total number of runs. The SLA policy, time of the job, and that status of the job is displayed. The information can be expanded of an associated job to view further job details, such as the reason why a job failed and the size of a successful backup. The **Search** box may be used to further filter report results.

File System SLA Policy RPO Compliance report

Use the report options to select a specific **Server** to display. The report can be further filtered by **Protection Type**, including data that was backed up to vSnap, using replication, using object storage copy, or using archive. The default setting for these two filters is **All**. File systems that are in compliance or not in compliance with the defined RPO can be filtered through the **Display File Systems That Are** field.

The File System SLA Policy RPO Compliance report displays file systems in relation to recovery point objectives as defined in SLA policies. Report properties display the creation date and the account that was used to generate the report. Also included are the report filters used when the report was generated. The quick view displays a pie chart of a count of backups to vSnap that are in compliance and those backups that are not in compliance. The summary view displays the SLA policy, the SLA schedule, the number of backups to vSnap and jobs using replication. Non-compliant file system SLA policy jobs are included if the not in compliance filter is selected. Information that is displayed are non-compliant SLA jobs using: backup to vSnap, replication, object storage copy, and archive. For non-compliant file system SLA policy jobs, the SLA policy and SLA schedule is listed with each file system, server, the last successful protection time, and the out of compliance reason.

Microsoft 365 Backup History report

Run the Microsoft 365 Backup History report to review the protection history for specific tenants. To run the report, specify at least one tenant in the **Tenants** filter. To filter reports based on a particular jobs, specify a value in the **Status** field. Reports can be further filtered based on specific service level agreement (SLA) policies by specifying a value in the **SLA Policy** field. The default setting for all four options is **All**. To limit the report to a specified number of days, enter an integer in the **Backup History for Past Number of Days** field.

Report properties list the creation date and the account that was used to generate the report. Also included are the filters that were used to generate the report. In the detail view of the report, the file system is listed with the server and the total number of runs. The SLA policy, the job run time, and the job status are displayed. You can expand the report to view further details, such as the reason why a job failed and the size of a successful backup. The **Search** box can be used to further filter report results.

Microsoft 365 SLA Policy RPO Compliance report

The Microsoft 365 SLA Policy RPO Compliance report displays tenants that meet or do not meet recovery point objectives (RPOs) as defined in SLA policies. Report properties display the creation date and the account that was used to generate the report. Report properties also include any filters that were applied.

Use the report options to select a specific tenant for the report. The report can be filtered by specifying a value in the **Protection Type** field. For example, you can include data that was backed up to a vSnap server, was replicated, was copied to IBM Spectrum Protect, or was archived. You can filter the report to show either policy settings.

The **Quick View** displays pie charts of the backup copies written to the vSnap server and show the levels of compliance and noncompliance. The **Summary View** displays the SLA policy, the SLA schedule, the number of backup operations to the vSnap server, and the number of replication jobs. Tenants that are not in compliance are listed with the number of applications that are affected. The last successful backup is listed with the reason why the tenant is not compliant.

Protected and Unprotected Databases report

Run the Protected and Unprotected Databases report to view the protection status of your databases. The report displays the total number of databases added to the IBM Spectrum Protect Plus inventory before backup jobs are started. Use the report options to filter by **Application Type**, **Application Server**, and **Application Server Type**to display. To exclude databases that are protected through hypervisor-based backup jobs, select **Hide Databases Protected as part of Hypervisor Backup**. To exclude unprotected databases in the report, select **Hide Unprotected Databases**.

The summary view displays an overview of your application server protection status, including the number of unprotected and protected databases, as well as the front end capacity of the protected databases. The front end capacity is the used capacity of a database. The detail view is displayed for each database type and provides further information including database names, application server, and hosting VM. The detail view also provides this information about unprotected databases in the detail view - unprotected databases section. The **Search** box may be used to further filter report results.

Protected and Unprotected File Systems report

Run the Protected and Unprotected File Systems report to view the protection status of your file systems. The report displays both the protected and unprotected file systems added to the IBM Spectrum Protect Plus inventory before backup jobs are started. Use the report options to filter by **Server**, **Operating System Type**, and **File System Type** to display. To exclude file systems that are protected through hypervisor-based backup jobs, select **Hide File Systems protected as part of Hypervisor Backup**. To exclude unprotected file systems in the report, select **Hide Unprotected File Systems**.

Report properties displays the creation date and the account that was used to generate the report. Also included are the report filters used when the report was generated. Summary view displays the protection status of registered file systems. Two detailed views are displayed, one for protected file systems and the other for unprotected file systems. Information is organized by File System, Path, File System Type, OS Type, and Server with the total number of protected and unprotected file systems displayed. The **Search** box may be used to further filter report results.

Protected and Unprotected Microsoft Office Users report

Run the Protected and Unprotected Microsoft Office Users report to view the protection status of the user accounts in your Microsoft 365 tenants. The report displays both the protected and unprotected Microsoft 365 tenants that were added to the IBM Spectrum Protect Plus inventory before the backup jobs ran. To filter a report based on tenants or applications, specify values in the **Tenants** or **Application** fields. To exclude unprotected accounts from the report, select **Hide Unprotected Accounts**.

If you decide to view all accounts, the report includes two areas, one for protected user accounts and the other for unprotected user accounts.

Within each of these sections, depending on whether you filtered results by application, there is a subsection for OneDrive and another for Outlook. For each of these applications, tenant proxies with details of each user account at that proxy site are listed.

You can use the **Search** option to locate information in the report.

Protected and Unprotected VMs report

Run the Protected and Unprotected VMs report to view the protection status of your virtual machines. The report displays the total number of virtual machines added to the IBM Spectrum Protect Plus inventory before backup jobs are started.

Use the report options to filter by **Hypervisor Type** and to select specific **Hypervisor/Accounts** to display. To exclude unprotected virtual machines in the report, select **Hide Unprotected VMs**. To exclude virtual machines that are not backed up to secondary backup storage, select **Show only the VMs with Object Storage Copy Backups**. **Tags** may also be used to filter reports.

The protected VMs displays an overview of your protected virtual machines, including the total number of VMs protected, the VM name, the hypervisor/account, type of hypervisor, location, and the managed capacity. The managed capacity is the used capacity of a virtual machine. The unprotected VMs provides the same information for virtual machines that are not protected. The **Search** box may be used to further filter report results.

VM Backup History report

Run the VM Backup History report to review the protection history of specific virtual machines. To run the report, at least one virtual machine must be specified in the **VMs** option. You can select multiple virtual machine names. Use the report options to filter **Status** by failed or successful jobs. The report can be further filtered by specific service level agreement (SLA) policies using the **SLA Policy** field. An integer may be specified for the **Backup History for Past Number of Days** field to limit results and **Tags** may also be used to filter the report.

The detail view displays the SLA policy used listed under the VM, account, and total number of runs. Information for each run can be expanded to list the backup data size. The protection time, status, and the backup storage used is also displayed. The **Search** box may be used to further filter report results.

VM SLA Policy RPO Compliance report

Use the report options to filter by **Type**, **Hypervisor/Account**, the **Protection Type**which includes data that was backed up to vSnap, using replication, using object storage copy, using archive or using snapshot, and to display virtual machines that are in compliance or not in compliance with the defined RPO through the **Display VMs That Are** field. There is also a filter for **Tags**.

The VM SLA Policy RPO Compliance report displays virtual machines in relation to recovery point objectives as defined in SLA policies. The quick view displays a pie chart of a count of backups to vSnap that are in compliance and those that are not in compliance. Also is a pie chart of snapshots that are in compliance and those that are not in compliance. A summary view displays the SLA policy used, the SLA schedule, the ratio of backups to vSnap for in compliance and not in compliance and the snapshot ratio of in compliance and not in compliance. Also displayed are VMs not in compliance view for each protection type. The **Search** box may be used to further filter report results.

Related concepts

“Types of reports” on page 631
You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

System reports

IBM Spectrum Protect Plus provides system reports that display an in-depth view of the status of your configuration, including storage system information, jobs, and job status.

To view system reports, complete the following steps:

1. In the navigation panel, click **Reports and Logs** > **Reports**.
2. Click on the **Reports** tab.
3. Select **System** in the **Filter by category** drop-down menu.
4. Run the report by clicking the **Run Report** ( ) icon beside the desired report.

The following reports are available:

Configuration report

Use the option, **Configuration Type**, to filter the configuration types to display. The Configuration report displays the configuration of the application servers, virtualized systems, backup storage for disk, object storage, and repository servers, VADP proxies, LDAP servers, and SMTP servers. Included

in the report is the name of the resource, resource type (OS or application), provider, associated site, state and the TLS connection status. Not all options are displayed for each component in the Configuration report. The **Search** box may be used to further filter report results.

Job report

Use the report options to filter the job types by selecting the **Job Type** selection box and to display jobs that ran successfully over a period of time in the **Days Since Successful Run** selection box. The quick view displays a pie chart with the number of completed jobs, failed jobs, and other jobs. Summary view for jobs that have been run at least once displays the type of job, the number of jobs associated with that type, the number of runs, the number of completed jobs, failed jobs, and other jobs. The detail view for jobs run at least once includes the job, type, number of runs, and the number of completed jobs, failed jobs, and other jobs, the last successful run, and the success percentage. In all cases, other jobs are jobs that are aborted, partially run, are currently running, skipped, or stopped.

In the detail view, click the plus () icon next to an associated job to view further job details such as the job ID, the average run time, last run time status, last run time, and the next scheduled run time if the job is scheduled, and the protected resources. At the end of the report is a detail view for jobs that have never run.

License report

Review the configuration of your IBM Spectrum Protect Plus environment in relation to licensed features. The following sections and fields display in this report:

Virtual Machine Protection

The **Total Number of VMs** field displays the total number of virtual machines protected through hypervisor backup jobs, plus the number of virtual machines hosting application databases protected through application backup jobs (not hypervisor backup jobs). The **Front End Capacity** field displays the used size of these virtual machines.

Physical Machine Protection

The **Total Number of Physical Servers** field displays the total number of physical application servers hosting databases that are protected through application backup jobs. The **Front End Capacity** field displays the used size of these physical application servers.

Office 365 Protection

The **Office 365 Protection** field displays the users protected through the Office 365 application backup job. The **Front End Capacity** field displays the total used size of the protected users.

Container Persistent Volume Protection

The **Container Persistent Volume Protection** field displays the protected container persistent volumes. The **Front End Capacity** field displays the used size of these protected container persistent volumes.

Backup Storage Utilization (vSnap)

The **Total Number of vSnap Servers** field displays the number of vSnap servers that are configured in IBM Spectrum Protect Plus as a backup destination. The **Target Capacity** field displays the total used capacity of the vSnap servers, excluding replica destination volumes.

Related concepts

“Types of reports” on page 631

You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

Running a VM environment report

You can run reports for your Virtual Machine (VM) environment in IBM Spectrum Protect Plus. Reports can help you to monitor the amount of free space on each hypervisor, the storage usage of logical unit numbers (LUNs), and the status of all VMs.

Procedure

1. In the navigation panel, click **Reports and Logs** > **Reports**.
2. Click on the **Reports** tab.

3. Select **VM Environment** in the **Filter by category** drop-down menu.
4. Run the report by clicking the **Run Report** () icon beside the desired report.

   The following reports are available:

   VM Datastore report

   Choose this to review the storage utilization of the datastores in your VM environment. Information that this report provides can be filtered using the **Hypervisor Type** and **Hypervisor**. The **Detail View Filter** controls the datastores to display in the detail view based on the percentage of space used. Use the **Show Only Orphaned Datastores** filter to view datastores that do not have any virtual machines assigned to them, or virtual machines that are in an inaccessible state. The reason for a datastore to be in an orphaned state is displayed in the **Datastore** field in the detail view.

   Quick view displays a pie chart with the storage utilization of free and used space. The summary view displays the hypervisor, datastore count, the capacity and the free space. The detail view shows the datastores and displays orphaned datastores that have no VMs registered. Also displayed is the associated hypervisor, hypervisor type, datastore type, the capacity, the free space, and the percentage used. All three views contain the total datastores, total capacity, and the total free space. The **Search** box may be used to further filter report results.

   VM LUNs report

   Review the storage utilization of your virtual machine logical unit numbers (LUNs). Filters for this report type include **Hypervisor Type**, and **Hypervisors**. Use the **Show Only Orphaned Datastores** filter to view datastores that do not have any virtual machines assigned to them, or virtual machines that are in an inaccessible state.

   In the report, the summary view displays the hypervisor, the number of LUNs associated with the hypervisor, and the capacity. In the detail view, the LUN name, LUN ID, storage vendor, hypervisor, the datastore or volume, the capacity, transport type, and the raw device mapping for each LUN is displayed. Both views display the total LUN count and total capacity. The **Search** box may be used to further filter report results.

   VM Snapshot Sprawl report

   This snapshot sprawl report displays the age, name, and number of snapshots that are used to protect your Hypervisor resources. The available report options to filter are by **Hypervisor Type**, **Hypervisor**, and **Tags**. Use the **Snapshot Creation Time** filter to display snapshots from specific periods of time.

   The report contains a detail view which displays the snapshot name and snapshot creation time. Each snapshot appears under the associated VM, hypervisor, and hypervisor type. The total number of VMs and snapshots are displayed at the end of the view. The **Search** box may be used to further filter report results.

   VM Sprawl report

   Review the status of your virtual machines, including virtual machines that are powered off, powered on, or suspended. Run this report to view unused virtual machines, the date and time when they were powered off, and virtual machine templates. The available report options to filter are by **Hypervisor Type**, **Hypervisor**, **Days Since Last Powered Off**, **Dayces Since Last Suspended**, **Days Since Last Powered On**, and **Tags**.

   The report contains the quick view which is a pie graph that displays the storage utilization based on the virtual machine's power state: powered off VMs, powered on VMs, templates, and suspended VMs. There are also detail views for each of the power states. The detail view - powered off VMs displays the VM name, the date and number of days since powered off, the associated hypervisor, the type of hypervisor, the provisioned space, and the datastore or volume. The total powered off VMs are displayed at the bottom of this view along with the total provisioned space. The detail view - suspended VMs contains the VM name, the date and number of days since the VM was suspended, the associated hypervisor, the type of hypervisor, the provisioned space, and the datastore of volume. The total number of suspended VMs and the total provisioned

space is displayed at the bottom of the view. The detail view - templates contains the template names, associated hypervisor, the hypervisor type, the provisioned space, and the datastore or volume. The total templates and total provisioned space appears at the bottom of the view. The detail view - powered on VMs contains the VM name, the date and number of days that the VM has been powered on, the associated hypervisor, the hypervisor type, the provisioned space, and the datastore or volume. At the end of the view is the total number of powered on VMs and the total provisioned space. The **Search** box may be used to further filter report results.

VM Storage report

Review your virtual machines and associated datastores in this report. View associated datastores and provisioned space of the datastores. Use the report options to filter by **Hypervisor Type** and to select which **Hypervisor** to display.

The report contains a detail view which displays the VM name and the provisioned space. Each VM appears under the associated datastore or volume, hypervisor, and hypervisor type. The total number of datastores/volumes and VMs are displayed at the end of the view. The **Search** box may be used to further filter report results.

Related concepts

"Types of reports" on page 631
You can customize predefined reports to monitor the utilization of backup storage and other aspects of your system environment.

Report actions

You can run, save, or schedule reports in IBM Spectrum Protect Plus.

Running a report

You can run IBM Spectrum Protect Plus reports with default parameters or run customized reports with custom parameters.

Before you begin

Custom roles that are assigned to users that run reports require that the appropriate permissions be set on that role so that the report can be viewed. For more information about roles, permission types, and permissions, see "Managing roles" on page 648.

Procedure

To run a report, complete the following steps:

1. In the navigation panel, click **Reports and Logs** > **Reports**.
2. Click on the **Reports** tab.
3. Run the report by clicking the **Run Report** ( ) icon beside the desired report.
   - To run the report with custom parameters, set the parameters in the **Run Report** window, and click **Run**. Parameters are unique to each report.
   - To run the report with default parameters, click **Run**.

What to do next

Review the report in the **Reports** pane.

Related concepts

"Managing reports and logs" on page 631

IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Creating a custom report

You can modify predefined reports with custom parameters in IBM Spectrum Protect Plus and save the customized reports.

Procedure

To create a report, complete the following steps:

1. In the navigation panel, click **Reports and Logs** > **Reports**.
2. Click on the **Reports** tab.
3. Click the **Create Custom Report** (+) icon beside the desired report to be customized.
4. On the **Create Custom Report** window, select the **Parameters** tab. Enter a name for the report in the **Name** field, and enter a description for the custom report in the **Description** field. Set your customized parameters that relate to the selected report.

   Note: Report names can include alphanumeric characters and the following symbols: $-_.+!*'(). Spaces are not allowed in the report name.
5. Optionally, on the **Schedule** tab, check the **Define Schedule** box. If a schedule is to be defined, provide this information:
   - For **Repeats**, enter an integer value and select **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.
   - For **Start Time**, enter a date and time, and select the appropriate timezone. The default timezone that is displayed is based on browser settings.
   - Enter the e-mail address of the recipient that is to receive a copy of the report in the e-mail address field. At least one recipient must be added. If more addresses are required, click on the **Add a receipient** plus (+) icon.
6. Click the **Save Report** button.
7. To locate a custom report, click on the **Custom Reports** tab.
8. Click on the **Run Custom Report** (▶) icon to run the report.
9. Optionally, to update a custom report, click the **Update Custom Report** ( ) icon. To remove a custom report, click the **Remove Report** ( ) icon.

What to do next

Run the custom report and review the report results.

Related concepts

“Managing reports and logs” on page 631

IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Scheduling a report

You can schedule reports in IBM Spectrum Protect Plus to run at specific times.

Procedure

To schedule a report, complete the following steps:

1. In the navigation panel, click **Reports and Logs** > **Reports**.
2. Click on the **Reports** tab.
3. Define a schedule for a report by clicking the **Schedule Report with default parameters** () icon beside the desired report.

   Note: To schedule a report with non-default parameters, create a custom report. For more information, see "Creating a custom report" on page 639.
4. The **Schedule Report with default parameters** window will appear.
   - For **Repeats**, enter an integer value and select **Subhourly**, **Hourly**, **Daily**, **Weekly**, **Monthly**, or **Yearly**. When **Weekly** is selected, you may select one or more days of the week. The **Start Time** will apply to the selected days of the week.
   - For **Start Time**, enter a date and time, and select the appropriate timezone. The default timezone that is displayed is based on your web-browser's settings.
   - Enter the e-mail address of the recipient that is to receive a copy of the report in the e-mail address field. At least one recipient must be added. If more addresses are required, click on the **Add a receipient** plus () icon.
5. Click the **Schedule** button.

What to do next

After the report runs, the recipient can review the report, which is delivered by email.

Related concepts

"Managing reports and logs" on page 631

IBM Spectrum Protect Plus provides a number of predefined reports that you can customize to meet your reporting requirements. A log of actions that users complete in IBM Spectrum Protect Plus is also provided.

Collecting audit logs for actions

You can collect audit logs and search for actions that are completed in IBM Spectrum Protect Plus.

Procedure

To collect audit logs:

1. In the navigation panel, click **Reports and Logs** > **Audit Logs**.
2. Review a log of actions that were completed in IBM Spectrum Protect Plus. Information includes the users who completed the actions and descriptions of the actions.
3. To search for the actions of a specific user in IBM Spectrum Protect Plus, enter the user name in the user search field.
4. Optional: Expand the **Filters** section to further filter the displayed logs. Enter specific action descriptions and a date range in which the action was completed.
5. Click the search icon  .

6. To download the audit log as a `.csv` file, click **Download**, and then select a location to save the file.

Related concepts

“Managing user accounts” on page 653
Before a user can log on to IBM Spectrum Protect Plus and use the available functions, a user account must be created in IBM Spectrum Protect Plus.

Chapter 19. Managing user access

By using role-based access control, you can set the resources and permissions available to IBM Spectrum Protect Plus user accounts.

You can tailor IBM Spectrum Protect Plus for individual users, giving them access to the features and resources that they require.

Once resources are available to IBM Spectrum Protect Plus, they can be added to a resource group along with high-level IBM Spectrum Protect Plus items such as a hypervisor and individual screens.

Roles are then configured to define the actions that can be performed by the user associated with the resource group. These actions are then associated with one or more user accounts.

Use the following sections of the **Accounts** pane to configure role-based access:

Resource Groups
A resource group defines the resources that are available to a user. Every resource that is added to IBM Spectrum Protect Plus can be included in a resource group, along with individual IBM Spectrum Protect Plus functions and screens. By defining resource groups, you can fine tune the user experience. For example, a resource group could include an individual hypervisor, with access to only backup and reporting functionality. When the resource group is associated with a role and a user, the user will see only the screens that are associated with backup and reporting for the assigned hypervisor.

Restriction: Do not assign a role-based access control (RBAC) user to more than one VMware resource group. Users that have been assigned to the Tag and Categories resource group and then are also assigned to either Hosts and Clusters or VMs and Templates will result in data not being displayed for the Hosts and Clusters view or the VMs and Templates view. Only information for Tags and Categories will be displayed when that is selected as a view when performing operations.

Roles
Roles define the actions that can be performed on the resources that are defined in a resource group. While a resource group defines the resources that will be made available to a user account, a role sets the permissions to interact with the resources defined in the resource group. For example, if a resource group is created that includes backup and restore jobs, the role determines how a user can interact with the jobs.

Permissions can be set to allow a user to create, view, and run the backup and restore jobs that are defined in a resource group, but not delete them. Similarly, permissions can be set to create administrator accounts, allowing a user to create and edit other accounts, set up sites and resources, and interact with all of the available IBM Spectrum Protect Plus features.

User accounts
A user account associates a resource group with a role. To enable a user to log in to IBM Spectrum Protect Plus and use its functions, you must first add the user as an individual user (referred to as a native user) or as part of an imported group of LDAP users, and then assign resource groups and roles to the user account. The account will have access to the resources and features that are defined in the resource group as well as the permissions to interact with the resources and features that are defined in the role.

© Copyright IBM Corp. 2017, 2021

Managing user resource groups

A resource group defines the resources are made available to a user. Every resource added to IBM Spectrum Protect Plus can be included in a resource group, along with individual IBM Spectrum Protect Plus functions and screens.

Creating a resource group

Create a resource group to define the resources that are available to a user.

Before you begin

You may not assign more than one application per machine as an application server to a resource group. For example, if SQL and Exchange occupy the same machine and both are registered with IBM Spectrum Protect Plus, only one of those can be added as an application server to a given resource group.

Procedure

To create a resource group, complete the following steps:

1. In the navigation panel, click **Accounts** > **Resource Group**.
2. Click **Create Resource Group**. The **Create Resource Group** pane displays.
3. Enter a name for the resource group.
4. From the **I would like to create a resource group** menu, select one of the following options:

Option	Actions
New	a. Select a resource type from the **Choose a resource type** menu. b. Select resource subtypes, and then click **Add Resources**. Resources are added to the **Selected Resources** view.
From template	a. Select a resource group from the **Which resource group would you like to use as a template?** list. Resources from the selected template are added to the **Selected Resources** view. b. You can add resources by using the **Choose a resource type** list and its associated lists. To view available resource types and their usage, see “Resource types ” on page 645.

   If you want to delete resources from the group, click the delete icon that is associated with a resource or click **Delete All** to delete all resources.
5. When you are finished adding resources, click **Create resource group**.

Results

The resource group displays in the resource group table and can be associated with new and existing user accounts.

What to do next

After you add the resource group, complete the following action:

Action	How to
Create roles to define the actions that can be performed by the user account that is associated with the resource group. Roles are used to define permissions to interact with the resources that are defined in the resource group.	See “Creating a role” on page 650.

Resource types

Resource types are selected when resource groups are created and determine the resources that are available to a user assigned to a group.

The following resource types and subtypes are available:

Table 106. Resource types and subtypes that can be selected for a resource group

Resource Type	Subtype	Description
Accounts	• Role • User • Identity	Used to grant access to roles and users through the **Accounts** pane.
Database	• Db2 • Exchange Standalone/Failover Cluster • Exchange Database Availability Groups • MongoDB • Oracle • SQL Standalone/Failover Cluster • SQL Always On	Used to grant access to viewing individual application databases on an application server in IBM Spectrum Protect Plus.
Cloud	Microsoft 365	Used to grant access to cloud system resources.
Container	• Kubernetes • OpenShift	Used to grant access to container resources.
File System	Windows	Used to grant access to file system resources.
Server	• All • Db2 • Exchange • File systems • Kubernetes • OpenShift • MongoDB • Microsoft 365 • Oracle • SQL	Used to grant access to application servers in IBM Spectrum Protect Plus without access to individual databases.

Table 106. Resource types and subtypes that can be selected for a resource group (continued)

Resource Type	Subtype	Description
Job	None	Used to grant access to Inventory, Backup, and Restore jobs. The Job resource group is mandatory for all Backup and Restore operations, including assigning SLA Policies to resources.
Report	• Backup Storage Utilization • Protection • System • VE Environment	Used to grant access to report types and individual reports.
Screen	None	Used to grant or deny access to screens in the IBM Spectrum Protect Plus interface. If certain screens are not included in a resource group for a user, the user will not be able to access the functionality provided on the screen, regardless of the permissions granted to the user.
SLA Policy	None	Used to grant access to SLA Policies for Backup operations.

Table 106. Resource types and subtypes that can be selected for a resource group (continued)

Resource Type	Subtype	Description
System Configuration	Certificates	Used to grant access to TLS certificates to access cloud servers.
	Object Storage	Used to grant access to object storage that is defined as backup storage for copy operations.
	Disk	Used to grant access to vSnap backup storage servers.
	Keys	Used to grant access to the credentials required to access your resources. Identity functionality is available through the **Accounts** > **Identities** pane.
	LDAP	Used to grant access to LDAP servers for user registration.
	Logs	Used to grant access to viewing and downloading Audit and System logs.
	Repository Servers	Used to grant access to a repository server.
	Script	Used to grant access to uploaded prescripts and postscripts.
	Script Server	Used to grant access to script servers, where scripts are run during a Backup or Restore job.
	Site	Used to grant access to sites, which are assigned to vSnap backup storage servers.
	SMTP	Used to grant access to SMTP servers for job notifications.
	VADP Proxy	Used to grant access to VADP proxy servers.
Virtualized System	• VMware • Hyper-V • Amazon EC2	Used to grant access to virtualized system resources.

Editing a resource group

You can edit a resource group to change the resources and features that are assigned to the group. Updated resource group settings take effect when user accounts that are associated with the resource group log in to IBM Spectrum Protect Plus.

Before you begin

Note the following considerations before editing a resource group:

- If you are signed in when the permissions or access rights for your user account are changed, you must sign out and sign in again for the updated permissions to take effect.
- You can edit any resource group that is not designated as **Cannot be modified**.

You may not assign more than one application per machine as an application server to a resource group. For example, if SQL and Exchange occupy the same machine and both are registered with IBM Spectrum Protect Plus, only one of those can be added as an application server to a given resource group.

Procedure

To edit a resource group, complete the following steps:

1. In the navigation panel, click **Accounts** > **Resource Group**.
2. Select a resource group and click the options icon ••• for the resource group. Click **Modify resources**.
3. Revise the resource group name, resources, or both.
4. Click **Update Resource Group**.

Deleting a resource group

You can delete any resource group that is not designated as **Cannot be modified**.

Procedure

To delete a resource group, complete the following steps:

1. In the navigation panel, click **Accounts** > **Resource Group**.
2. Select a resource group and click the options icon ••• for the resource group. Click **Delete resource group**.
3. Click **Yes**.

Managing roles

Roles define the actions that can be completed for the resources that are defined in a resource group. While a resource group defines the resources that are available to an account, a role sets the permissions to interact with the resources.

For example, if a resource group is created that includes backup and restore jobs, the role determines how a user can interact with the jobs. Permissions can be set to enable a user to create, view, and run the backup and restore jobs that are defined in a resource group, but not delete them.

Similarly, permissions can be set to create administrator accounts, enabling a user to create and edit other accounts, set up sites and resources, and interact with all of the available IBM Spectrum Protect Plus features.

The functionality of a role is dependent on a properly configured resource group. When selecting a predefined role or configuring a custom role, you must ensure that access to necessary IBM Spectrum Protect Plus operations, screens, and resources align with the proposed usage of the role.

About the SUPERUSER role: The SUPERUSER role provides the user with access to all IBM Spectrum Protect Plus functions. The SUPERUSER role can be assigned to only one account and that account is referred to as the superuser account. This superuser account and the SUPERUSER role are discussed in “Managing the superuser account” on page 655.

The following user account roles are available:

Application Admin
Users with the Application Admin can complete the following actions:

- Register and modify application database resources that are delegated by an administrator
- Associate application databases to assigned SLA policies

- Complete backup and restore operations
- Run and schedule reports to which the user has access

Access to resources must be granted by an administrator through the **Accounts** > **Resource Groups** pane.

Backup Only

Users with the Backup Only role can complete the following actions:

- Create, view, and run backup operations
- View, create, and edit SLA policies to which the user has access

Access to resources, including specific backup jobs, must be granted by an administrator by clicking **Accounts** > **Resource Groups**.

OC_MONITOR_ROLE

The OC_MONITOR_ROLE is created when an OC_MONITOR user is created by the IBM Spectrum Protect Operations Center. This role and user are required by the Operations Center to connect to the IBM Spectrum Protect Plus environment. The OC_MONITOR_ROLE is used only by the OC_MONITOR user and provides permissions that are required to connect the Operations Center to IBM Spectrum Protect Plus. Do not edit this role.

Restore Only

Users with the Restore Only role can complete the following actions:

- Run, edit, and monitor restore operations.
- View, create, and edit SLA Policies to which the user has access.

Access to resources, including specific restore jobs, must be granted by an administrator through the **Accounts** > **Resource Groups** pane.

Self Service

Users with the Self Service role can monitor existing backup and restore operations that are delegated by an administrator.

Access to resources, including specific jobs, must be granted by an administrator through the **Accounts** > **Resource Groups** pane.

SYSADMIN

The SYSADMIN role is the administrator role. This role provides access to all resources and privileges.

Users with this role can add users and complete the following actions for all users other than the user who is assigned the SUPERUSER role:

- Modify and delete user accounts
- Change user passwords
- Assign user roles

VM Admin

Users with the VM Admin role can complete the following actions:

- Register and modify hypervisor resources to which the user has access
- Associate hypervisors to SLA policies
- Complete backup and restore operations
- Run and schedule reports to which the user has access

Access to resources must be granted by an administrator through the **Accounts** > **Resource Groups** pane.

Creating a role

Create roles to define the actions that can be completed by the user of an account that is associated with a resource group. Roles are used to define permissions to interact with the resources that are defined in the resource group.

Procedure

To create a user role, complete the following steps:

1. In the navigation panel, click **Accounts** > **Role**.
2. Click **Create Role**. The **Create Role** pane displays.
3. From the **I would like to create a role** list, select one of the following options:

Option	Actions
New	Select permissions to apply to the role. By default, none of the permissions are pre-selected.
From template	a. Select a role from the **Which role would you like to use as a template?** menu. Permissions that are associated with the template role are selected by default. b. Select additional permissions to apply to the role, and delete permissions that are not required. To view available permissions and their usage, see “Permission types ” on page 650.

4. Enter a name for the role, and then click **Create Role**.

Results

The new role is displayed in the roles table and can be applied to new and existing user accounts.

Permission types

Permission types are selected when user accounts are created and determine the permissions that are available to the user.

The following permissions are available:

Name	Permissions	Description
Application	View	Used to view individual application databases on an application server in IBM Spectrum Protect Plus.
Application Server	Register, view, edit, deregister	Used to interact with application servers, such as SQL or Oracle servers, without access to individual databases.
Certificate	Create, view, edit, delete	Used to interact with TLS certificates to access cloud servers.
Object Storage	Register, view, edit, deregister	Used to interact with object storage that is defined as backup storage for copy operations.

Name	Permissions	Description
Cloud	Register, view, edit, deregister	Used to interact with cloud servers that are defined as backup storage for copy operations.
Hypervisor	Register, view, edit, deregister, options	Used to interact with hypervisor virtual machines, such as VMware or Hyper-V virtual machines.
Identity and Keys	Create, view, edit, delete	Used to interact with the credentials required to access your resources. Identity functionality is available through the Accounts > Identities pane.
LDAP	Register, view, edit, deregister	Used to interact with LDAP servers for user registration.
Log	View	Used to view Audit and System logs.
Job	Create, view, edit, run, delete	Used to interact with Inventory, Backup, and Restore jobs. **Note:** If the user has permission to **Run** a job, then they also can **Hold**, **Release**, and **Perform custom restore actions** for the job.
VADP Proxy	Register, view, edit, deregister	Used to interact with VADP.
Report	Create, view, edit, delete	Used to interact with reports.
Resource Group	Create, view, edit, delete	Used to interact with resource groups, which define the IBM Spectrum Protect Plus resources that are made available to a user.
Role	Create, view, edit, delete	Used to interact with roles, which define the actions that can be performed on the resources defined in a resource group.
Script	Upload, view, replace, delete	Used to interact with prescripts and postscripts that are added to IBM Spectrum Protect Plus and run before or after a job.
Script Server	Register, view, edit, deregister	Used to interact with the server on which prescripts and postscripts run.
Site	Create, view, edit, delete	Used to interact with sites, which are assigned to vSnap backup storage servers.
SMTP	Register, view, edit, deregister	Used to interact with SMTP servers for job notifications.
Backup Storage	Register, view, edit, deregister	Used to interact with vSnap backup storage servers.

Name	Permissions	Description
SLA Policy	Create, view, edit, delete	Used to interact with SLA Policies, which allow users to create customized templates for Backup jobs.
User	Create, view, edit, delete	Used to interact with users, associate a resource group with a role, and provide access to the IBM Spectrum Protect Plus user interface.

Editing a role

You can edit a role to change the resources and permissions that are assigned to the role. Updated role settings take effect when user accounts that are associated with the role log in to IBM Spectrum Protect Plus.

Before you begin

Note the following considerations before editing a role:

- If you are signed in when the permissions or access rights for your user account are changed, you must sign out and sign in again for the updated permissions to take effect.
- You can edit any role that is not designated as **Cannot be modified**.

Procedure

To edit a user role, complete the following steps

1. In the navigation panel, click **Accounts** > **Role**.
2. Select a role and click the options icon ••• for the role. Click **Modify Role**.
3. Revise the role name, permissions, or both.
4. Click **Update role**.

Deleting a role

You can delete a role that is not designated as **Cannot be modified**.

Procedure

To delete a role, complete the following steps:

1. In the navigation panel, click **Accounts** > **Role**.
2. Select a role and click the options icon ••• for the role. Click **Delete role**.
3. Click **Yes**.

Managing user accounts

Before a user can log on to IBM Spectrum Protect Plus and use the available functions, a user account must be created in IBM Spectrum Protect Plus.

Creating a user account for an individual user

Add an account for an individual user in IBM Spectrum Protect Plus. If you are upgrading from a version of IBM Spectrum Protect Plus that is earlier than 10.1.1, permissions assigned to users in the previous version must be reassigned in IBM Spectrum Protect Plus.

Before you begin

If you want to use custom roles and resource groups, create them before you create a user. See "Creating a resource group" on page 644 and "Creating a role" on page 650.

Procedure

To create an account for an individual user, complete the following steps:

1. In the navigation panel, click **Accounts** > **User**.
2. Click **Add User**. The **Add User** pane is displayed.
3. Click **Select the type of user or group you want to add** > **Individual new user**.
4. Enter a name and password for the user.
5. In the **Assign Role** section, select one or more roles for the user.
6. In the **Permission Groups** section, review the permissions and resources that are available to the user, and then click **Continue**.
7. In the **Add Users - Assign Resources** section, assign one or more resource groups to the user, and then click **Add resources**.

   The resource groups are added to the **Selected Resources** section.
8. Click **Create user**.

Results

The user account is displayed in the users table. Select a user from the table to view available roles, permissions, and resource groups.

Creating a user account for an LDAP group

With IBM Spectrum Protect Plus, you can use a Lightweight Directory Access Protocol (LDAP) server to manage users. When you create an LDAP user account, you can add the user account to a user group.

Before you begin

Complete the following tasks:

- Ensure that you have registered an LDAP provider with IBM Spectrum Protect Plus. To register an LDAP provider, follow the instructions in "Adding an LDAP server" on page 294.
- If you want to use custom roles and resource groups, ensure that the roles or groups are available. For instructions about creating roles and groups, see "Creating a role" on page 650 and "Creating a resource group" on page 644.

Procedure

To create a user account for an LDAP group, complete the following steps:

1. In the navigation panel, click **Accounts** > **User**.
2. Click **Add User**. The **Add User** pane is displayed.

3. Click **Select the type of user or group you want to add** > **LDAP Group**.
4. In the **Group Name** field of the **Select LDAP Group** section, specify the LDAP group by taking one of the following actions:
   - Enter the LDAP group name.
   - Search for the LDAP group name by entering partial text, an asterisk (*) as a single wildcard character, or a question mark (?) for pattern matching. To view all LDAP groups, click the **View All** button.
   - Optionally, a relative distinguished name (RDN) can be provided by filling out the **Group RDN** field.
5. LDAP Groups are displayed in **LDAP Groups** table. Select an LDAP Group.
6. In the **Assign Role** section, select one or more roles for the user.
7. In the **Permission Groups** section, review the permissions and resources that are available to the user, and then click **Continue**.
8. In the **Add Users - Assign Resources** section, assign one or more resource groups to the user, and then click **Add resources**.

   The resource groups are added to the **Selected Resources** section.
9. Click **Create user**.

Results

The user account is displayed in the users table. Optionally, to view available roles, permissions, and resource groups, select a user in the users table.

Editing a user account

You can edit the user name, password, associated resource groups, and roles for a user account, with the exception of the user who is assigned the SUPERUSER role. For information about managing the super user, see “Managing the superuser account” on page 655.

Before you begin

If you are signed in when the permissions or access rights for your user account are changed, you must sign out and sign in again for the updated permissions to take effect.

Procedure

Complete the following steps to edit the credentials of a user account:

1. In the navigation panel, click **Accounts** > **User**.
2. Select one or more users. If you select multiple users with different roles, you can modify only their resources and not their roles.
3. Click the options icon ••• to view available options. The options that are shown depend on the selected user or users.

   Modify settings
   Edit the user name and password, associated roles, and resource groups.

   Modify resources
   Edit the associated resource groups.
4. Modify the settings for the user, and then click **Update user** or **Assign resources**.

Deleting a user account

You can delete any user account, with the exception of the user who is assigned the SUPERUSER role.

Procedure

To delete a user account, complete the following steps:

1. In the navigation panel, click **Accounts** > **User**.
2. Select a user.
3. Click the options icon ***, and then click **Delete user**.

Managing the superuser account

The superuser is the user who is assigned the IBM Spectrum Protect Plus SUPERUSER role. The SUPERUSER role provides the user with access to all IBM Spectrum Protect Plus functions.

When you log on to IBM Spectrum Protect Plus for the first time, you must log on with the username `admin` and the password `password`. You are then prompted to change this username and password. This process creates the superuser who is assigned the SUPERUSER role.

The following considerations apply to the superuser:

- There is only one superuser account.
- You cannot delete the superuser account.
- You cannot assign the SUPERUSER role to any additional user accounts that you create. You can create other accounts for administration purposes and assign the `SYSADMIN` role to those accounts.

Changing the superuser password

If you are logged on as the IBM Spectrum Protect Plus superuser, you can change the password for the superuser at any time. Only the superuser can change this password.

Procedure

To change the password for the superuser, complete the following steps:

1. In the navigation panel, click **Accounts** > **User**.
2. Select the superuser, and then click the options icon ***.
3. Click **Change Password**.
4. Enter the new password in the **Password** field and current password in the **Old Password** field, and then click **Update user**.

   The **Password** field is populated with the current password by default.

Resetting the superuser credentials

If you have forgotten the password for the superuser, you can reset the password. You can also reset the superuser name or maintain the existing name.

Before you begin

This procedure applies only if IBM Spectrum Protect Plus is installed as a virtual appliance. If IBM Spectrum Protect Plus is installed as a set of OpenShift containers, contact IBM Software Support for instructions about how to reset the password for the superuser.

Procedure

To reset the password for the superuser, complete the following steps:

1. Open the administrative console by entering the following URL from a supported browser:

   ```
   https://HOSTNAME:8090/
   ```

   where *HOSTNAME* is the IP address of the virtual machine where the application is deployed.
2. From the **Authorization type** list, select **System**.
3. Log on as the `serveradmin` user.

4. Click **System Management**, and then click **Reset SUPERUSER password**.
5. Click **Reset Password** to confirm the request.
6. Open the IBM Spectrum Protect Plus user interface by entering the following URL from a supported browser:

   ```
   https://host_name
   ```

   Where *host_name* is the IP address of the virtual machine where the application is deployed.
7. Log on by entering the superuser name and the password `password`.
8. At the prompt, enter a name for the superuser and a new password. You can enter the existing name or a new name.

   You cannot use the names `admin`, `root`, or `test`.
9. Click **Sign In**.

Managing identities

Some features in IBM Spectrum Protect Plus require credentials to access your resources. For example, IBM Spectrum Protect Plus connects to Oracle servers as the local operating system user that is specified during registration to complete tasks like cataloging, data protection, and data restore.

User names and passwords for your resources can be added and edited through the **Identity** pane. Then when utilizing a feature in IBM Spectrum Protect Plus that requires credentials to access a resource, select **Use existing user**, and select an identity from the drop-down menu.

Adding an identity

Add an identity to provide user credentials.

Procedure

To add an identity, complete the following steps:

1. In the navigation panel, click **Accounts** > **Identity**.
2. Click **Add Identity**.
3. Complete the fields in the **Identity Properties** pane:

   Name

   Enter a meaningful name to help identify the identity.

   Username

   Enter the user name that is associated with a resource, such as an SQL or Oracle server.

   Password

   Enter the password that is associated with a resource.
4. Click **Save**.

   The identity displays in the identities table and can be selected when you are using a feature that requires credentials to access a resource through the **Use existing user** option.

Editing an identity

You can revise an identity to change the user name and password used to access an associated resource.

Procedure

To edit an identity, complete the following steps:

1. In the navigation panel, click **Accounts** > **Identity**.

2. Click the edit icon that is associated with an identity.
   The **Identify Properties** pane displays.
3. Revise the identity name, user name, and password.
4. Click **Save**.
   The revised identity displays in the identities table and can be selected when utilizing a feature that requires credentials to access a resource through the **Use existing user** option.

Deleting an identity

You can delete an identity when it becomes obsolete. If an identity is associated with a registered application server, it must be removed from the application server before it can be deleted. To remove the association, navigate to the **Backup** > **Manage Application Servers** page associated with the application server type, then edit the settings of the application server.

Procedure

To delete an identity, complete the following steps:

1. In the navigation panel, click **Accounts** > **Identity**.
2. Click the delete icon that is associated with an identity.
3. Click **Yes** to delete the identity.

Chapter 20. Troubleshooting

Troubleshooting procedures are available for problem diagnosis and resolution.

For a list of known issues and limitations for each IBM Spectrum Protect Plus release, see technote 567387.

Troubleshooting installation issues for IBM Spectrum Protect Plus as a set of containers

Review the available information to resolve installation issues with the IBM Spectrum Protect Plus server in a container environment.

Table 107. Installation issues and potential resolutions for IBM Spectrum Protect Plus as a set of containers

Problem	Symptom	Resolution
Pods remain in the `ContainerCreating` state	After the IBM Spectrum Protect Plus instance is created, any of the following pods remain in the `ContainerCreating` state: • `virgo` • `nodejs` • `awsebs` • `awsec2`	Ensure that different zone label exists for each node by issuing the following command: `kubectl get nodes --show-labels | grep topology.kubernetes.io/zone` Set a zone label for each node if it is not already set by issuing the following command: `kubectl label nodes node_name topology.kubernetes.io/zone=label_value` where *node_name* is the name of a node in the OpenShift cluster, and *label_value* is the zone label for the node. For example, to set the zone label for all nodes on the cluster: `n=0; for node in $(kubectl get nodes --output NAME); do kubectl label $node topology.kubernetes.io/zone="us-local-$n"; (n++)); done`
The `virgo` pod is not ready	The IBM Spectrum Protect Plus `virgo` pod is taking too long to reach a ready state, or the liveness or readiness probes are failing for more than an hour.	Increase the CPU, memory, or both limits of the `virgo` pod. You can increase the limits by updating the **`spec.spp.virgo.limits`** field of the IBM Spectrum Protect Plus instance custom resource. The following values are the default limits for CPU and memory for the `virgo` pod: `limits:` `  cpu: 2000m` `  memory: 8Gi`

© Copyright IBM Corp. 2017, 2021

Collecting log files for troubleshooting

To troubleshoot the IBM Spectrum Protect Plus application, you can download an archive of log files that are generated by IBM Spectrum Protect Plus.

Procedure

To collect log files for troubleshooting, complete the following steps:

1. Click the user menu, and then click **Download System Logs**.

   The download process may take some time to complete.
2. Open or save the file log zip file, which contains individual log files for different IBM Spectrum Protect Plus components.

   For information about log files, see the protecting applications or protecting hypervisors backup sections.

What to do next

To troubleshoot issues, complete the following steps:

1. Analyze the log files and take appropriate actions to resolve the issue.
2. If you cannot resolve the issue, submit the log files to IBM Software Support for assistance.

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 278 and “Configuration for copying or archiving data to cloud” on page 271. 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 280.

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.

For information about how to set up and run a VMware restore job, see “Restoring VMware data” on page 350.

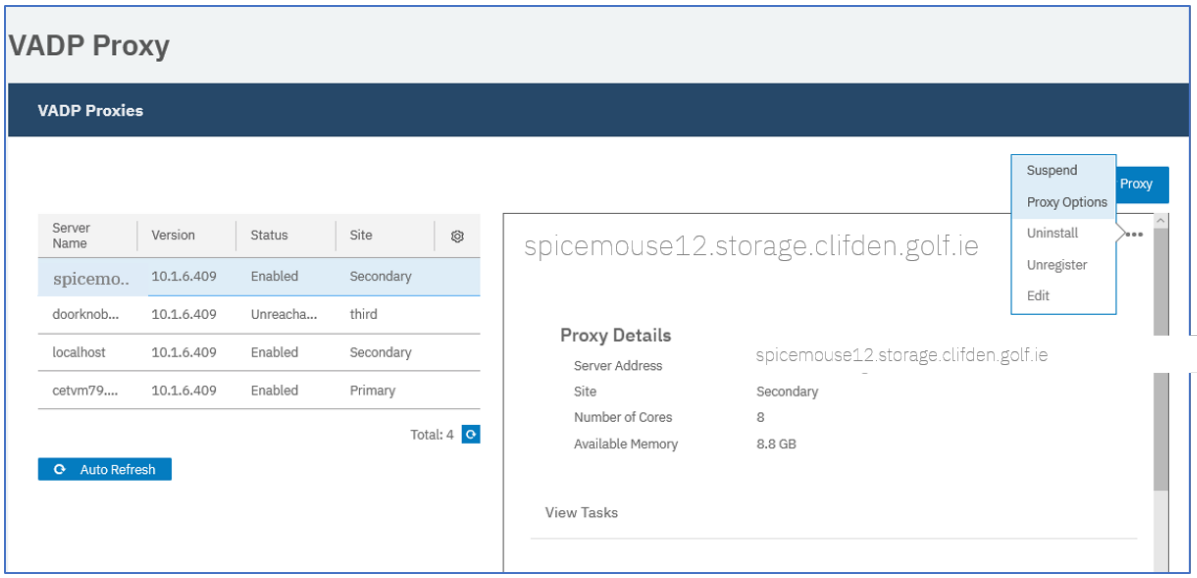
Communication

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.



4. From the **Transport Modes** list, select SAN.
5. Click **Save**.

Troubleshooting Container Backup Support

To help troubleshoot issues with Container Backup Support, you can collect debug log files and view trace logs. You can also follow procedures to diagnose problems.

Troubleshooting Container Backup Support installation issues

Review the available information to resolve Container Backup Support installation issues.

Possible installation issues and potential resolutions are provided.

An issue occurs when creating a secret

The following error occurs when a secret or other resource is created:

```
error: exactly one NAME is required, got 2
```

Resolution: Ensure that the values that you provided in the `baas-options.sh` file do not include spaces.

Pods remain in the pending state

After installation of Container Backup Support, some pods (for example, `zookeeper`) are in the pending state.

Resolution: Issue the following command and look for indications of insufficient CPU or other resources:

```
kubectl get events --namespace baas
```

If necessary, increase the value for the corresponding resource. For more information, see Resource requirements.

Issue occurs with the MinIO pod

The following error is displayed:

```
The pod:  baas-minio-0  is not in Running status
ERROR: Installation of product release baas version 10.1.8 to target cluster failed.
Please refer to log file at /tmp/baas_install.sh_20200804-165623.log for more information!
```

Resolution: Uninstall Container Backup Support, ensure that a default storage class is defined for your cluster, and reinstall Container Backup Support.

The `baas-minio` pod fails to start

The following error is issued for an installation on a cluster with IBM Spectrum Scale installed:

```
"CSINode <node> does not contain driver spectrumscale.csi.ibm.com":
  Warning  FailedMount           kubelet                  Unable to attach or mount
volumes: unmounted volumes=[storage], unattached volumes=[storage config tls-key-crt baas-
minio-token-sstqs]: timed out waiting for the condition
  Warning  FailedAttachVolume    attachdetach-controller  AttachVolume.Attach
failed for volume "pvc-7c3b64a7-b86d-42c8-a8c7-16b6293b776b" : CSINode
doorknob77.scaleocp.dkcloud.local does not contain driver spectrumscale.csi.ibm.com
```

Resolution: Complete the following steps:

1. Add the **`nodeSelector`** field to the `baas-minio` statefulset to schedule the `baas-minio` pod to the worker node that has the IBM Spectrum Scale CSI driver. Issue the following command:

   ```
   kubectl patch statefulset baas-minio -n baas -p '{"spec": {"template": {"spec":
   {"nodeSelector": {"scale": "true"}}}}}'
   ```

2. Restart the `baas-minio` pod by issuing the following command:

   ```
   kubectl delete pod baas-minio -n baas
   ```

Installing a second instance of Velero

You cannot complete the installation of a second instance of Velero if you are installing it in another namespace within the same cluster. When you attempt to install a second instance of Velero for namespace-scoped and cluster-scoped protection, a message indicates that the **clusterrolebinding** parameter already exists. To work around this issue, you can patch the **clusterrolebinding** parameter with a new service account for each version of Velero that is added to any namespace.

Procedure

To enable the installation of a second version of Velero, change to the directory where Velero was initially installed and run teh following command.

```
kubectl patch clusterrolebinding velero -p '{"subjects":[\
{"kind": "ServiceAccount", "name": "velero", "namespace": "<1st velero namespace>"},\
{"kind": "ServiceAccount", "name": "velero", "namespace": "<2nd velero namespace>"}]}
```

What to do next

When you install additional Velero packages in the cluster, repeat the steps in the procedure to complete the installation.

Collecting Container Backup Support log files for troubleshooting

You can generate debugging log files in the Kubernetes or OpenShift environment to troubleshoot the deployment of Container Backup Support and Container Backup Support operations on the IBM Spectrum Protect Plus server.

About this task

All logs are collected in the `~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/problemdetermination` directory on the local system and packaged into a `tar.gz` archive file. The archive file is typically named `baas_debug_logs_`*`timestamp`*`.tar.gz`.

Procedure

Use one of the following methods to collect logs for troubleshooting:

- To collect only Kubernetes or OpenShift logs for debugging purposes, issue the following command:

  ```
  cd ~/installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/
  problemdetermination
  ```

  This command collects debugging logs for the Container Backup Support deployment. The current state information and logs of the Container Backup Support components in the Kubernetes or OpenShift cluster are collected. The logs are structured based on the Kubernetes basic logging architecture. For more information, see Basic logging in Kubernetes.
- To collect the log package that includes the debugging logs for the Container Backup Support deployment and IBM Spectrum Protect Plus server, issue the following command:

  ```
  ./collect-logs.sh -x
  ```

What to do next

To troubleshoot issues, complete the following steps:

1. Analyze the log files and take appropriate actions to resolve the issue.
2. If you cannot resolve the issue, submit the log files to IBM Software Support for assistance.

Related tasks

“Setting the trace level of log files” on page 664

You can set the trace level of local log files to help troubleshoot issues that you might encounter in Container Backup Support.

Related reference

"Troubleshooting quick reference" on page 667
Solutions to basic Container Backup Support problems are provided.

"Troubleshooting Container Backup Support operations" on page 672
Troubleshooting procedures are available to help you diagnose and resolve Container Backup Support issues.

Setting the trace level of log files

You can set the trace level of local log files to help troubleshoot issues that you might encounter in Container Backup Support.

About this task

You can set the trace levels to troubleshoot issues with the Container Backup Support transaction manager, controller, and scheduler components. The trace level that you set also applies to log levels for the Container Backup Support agent, as well as the log levels in the IBM Spectrum Protect Plus job logs and `command.log` file.

The data mover component is not affected by this setting.

To set the trace level, you must update the `baas-values.yaml` file and then update the Container Backup Support deployment.

Tip: The default trace level is INFO. If you are experiencing issues that require troubleshooting, set the trace level to DEBUG.

Procedure

To set the trace level, complete the following steps at the Kubernetes command line:

1. Log in to the operating system on the master node of the Kubernetes cluster that is used as the installation node.
2. Go to the directory where the Container Backup Support installation package was unpacked (`installer/ibm-spectrum-protect-plus-prod/ibm_cloud_pak/pak_extensions/install`).
3. Edit the `baas-values.yaml` file with a text editor and modify the value for the **`productLoglevel`** parameter.

   The following trace options are available:

   `DEBUG`
   Display debugging-level messages in the transaction manager, controller, and scheduler log files.

   `INFO`
   Display all user messages in the transaction manager, controller, and scheduler log files, including information, warning, and error messages. This value is the default.

   `WARNING`
   Display warning and error messages in the transaction manager, controller, and scheduler log files.

   `ERROR`
   Display only error messages in the transaction manager, controller, and scheduler log files.

   For example, to set the trace level to debugging mode, set the **`productLoglevel`** parameter as follows:

   ```
   productLoglevel: DEBUG
   ```

4. Update the Container Backup Support deployment by issuing the following command:

```
./baas-upgrade.sh
```

5. Optional: To verify the status of the update, issue the following command:

```
helm3 status release_name -n baas
```

where *release_name* is the name of the Helm release for Container Backup Support, such as `ibm-spectrum-protect-plus-prod`.

Results

After the update is completed, verify the values that you specified by issuing the following command:

```
helm3 get values -n baas
```

You can also show the revision history by issuing the following command:

```
helm3 history release_name -n baas
```

where *release_name* is the name of the Helm release for Container Backup Support, such as `ibm-spectrum-protect-plus-prod`.

What to do next

You can collect Container Backup Support log files for troubleshooting or use a visualization tool such as Kibana to view and query data in the transaction manager, controller, and scheduler log files. For instructions, see:

- "Collecting Container Backup Support log files for troubleshooting" on page 663
- "Viewing trace logs for Container Backup Support" on page 665

Viewing trace logs for Container Backup Support

You can optionally use the Elasticsearch, Fluentd, and Kibana (EFK) stack to view and analyze trace logs that are produced by Container Backup Support.

Elasticsearch is a distributed full-text search engine. Fluentd is a tool that collects logs from cluster nodes and sends the logs to the Elasticsearch engine. Kibana is a visualization tool for Elasticsearch with a web user interface and development tool that is used for querying data.

Before you begin

Complete the following steps:

1. Deploy the EFK stack to your Kubernetes cluster:
   a. Deploy the Elasticsearch search engine. For instructions, see Installing Elasticsearch.
   b. Deploy the Fluentd log collector on each cluster node. For instructions, see the Fluentd documentation.
   c. Deploy the Kibana visualization tool. For instructions, see the Kibana Guide.
2. Complete the EFK stack deployment by adding a `logstash` index in Kibana:
   a. Access the Kibana user interface by opening a web browser and entering the URL of the computer where Kibana is running and specify the port number. For example, specify one of the following URLs in your web browser:

```
https://localhost:5601
```

or

```
http://your_domain.com:5601
```

where *your_domain* specifies the domain name for the computer.

b. If you are prompted with options to explore data, select **Explore on my own**.

c. Click the **Discover** > **Create Index Pattern** and create the `logstash-*` index pattern.

About this task

When you use the EFK stack, the logs from all container components are merged and shown in the same view. Any logs for stopped pods are preserved in Elasticsearch persistent data storage. You can apply filters to display specific errors or messages. You can also apply a time filter to show events that occurred in a specific time period.

In addition to error and debugging messages, you can view trace logs for the following Container Backup Support components:

- Transaction manager
- Controller
- Scheduler

Procedure

To view transaction logs for Container Backup Support, complete the following steps:

1. Open the Kibana user interface and click the **Discover** icon.
2. Click the `logstash-*` index.
3. To view logs for Container Backup Support, add a filter by taking the following actions:
   a) Click **Add filter** and specify the following filter values:
      - Field: `kubernetes.container_image`
      - Operator: `is`
      - Value: `baas-`
   b) Enter a name for the search and click **Save**.
      The trace logs for the `baas-transaction-manager`, `baas-controller`, and `baas-scheduler` containers are displayed.
4. You can create additional filters to show more granular views of Container Backup Support trace logs.

Table 108. Filters for viewing Container Backup Support trace logs

Type of data to show	Filter 1	Filter 2
Transaction manager logs	`kubernetes.container_image is baas-transaction-manager`	None
Controller logs	`kubernetes.container_image is baas-controller`	None
Scheduler logs	`kubernetes.container_image is baas-scheduler`	None
Error messages	`kubernetes.container_image is baas-`	`log is ERROR`
Debugging messages	`kubernetes.container_image is baas-`	`log is DEBUG`

Troubleshooting quick reference

Solutions to basic Container Backup Support problems are provided.

Use the solutions in the following table to resolve basic problems that might occur with Container Backup Support operations. If you still cannot resolve a problem, see "Troubleshooting Container Backup Support operations" on page 672 for more detailed troubleshooting procedures.

Table 109. Solutions to basic problems

Problem	Solution
The Container Backup Support request is invalid. For example, the **Backupstatus** or **Restorestatus** field is listed as `Invalid` when you run the following command: `kubectl describe baasreq request_name -n namespace` where: ***request_name*** The name of the backup or restore request. For backup requests, the value is the name of the persistent volume claim (PVC). For restore requests, the name must be unique, and must not be the same as the name of the PVC. ***namespace*** The namespace in which the PVC exists.	Ensure that the request is structured correctly by verifying the following elements in the YAML file: • Ensure that no typographical errors exist. • Ensure that the correct case is used in the statements. Kubernetes is case-sensitive. For example, ensure that the API version declaration is listed as `apiVersion` and not `apiversion`. • For restore requests: – Ensure that the timestamp for a restore point is specified correctly in the **restorepoint** field. – Ensure that the restore type is specified correctly in the **restoretype** field. For more information, see "Restoring container data by using the command line" on page 481.
The snapshots are failing.	Take one or more of the following actions: • Verify the CSI configuration to ensure that your containers are running correctly. The CSI software is required for snapshot backups. • Ensure that a volume snapshot class is defined for the PVCs that are being backed up. • Ensure that the secret is in the correct namespace (the namespace for the PVC). • Ensure that the configurations are correct in the configmap (`baas-configmap`). For more information, see "Troubleshooting issues with snapshot backup jobs" on page 672.

Table 109. Solutions to basic problems (continued)

Problem	Solution
The data mover fails to start.	Take one or more of the following actions: • Ensure that the PVC or resource is mounted. You can verify whether the PVC or resource is failing to mount by issuing the **kubectl describe** command on the data mover pod. • In the output of the **kubectl describe** command, check the events to ensure that the volume has been initialized by running the PVC as part of another pod in read/write mode. • In the output of the **kubectl describe** command, check for authentication failure events. To resolve authentication errors, ensure that you are running a secure Docker registry. Ensure that the pull secret is in the namespace of the PVC. For instructions, see Pull an Image from a Private Registry. • Ensure that the data mover is started in the namespace in which the PVC exists. For resources, ensure that the data mover is started in the baas namespace.
Access is denied or the connection fails while mounting NFS volumes from the vSnap server.	Take one or more of the following actions: • Check the data mover network policy. Ensure that the vSnap server addresses match the IBM Spectrum Protect Plus server addresses. • Ensure that a direct connection from the Kubernetes cluster to the IBM Spectrum Protect Plus vSnap server exists. Connection by proxies is not supported.
The scheduler, transaction manager, and controller pods have started but each pod continues to restart. In the output of the **kubectl describe** command for the transaction manager pod, the events indicate that the liveness probe failed.	Verify that the values for the **clusterAPIServerips** and **clusterAPIServerport** parameters are correctly specified in the `baas-values.yaml` file. If you update the values in the `baas-values.yaml` file, issue the following command to update the configuration from the directory where the installation was performed: `./baas-upgrade.sh` Alternatively, you can uninstall and reinstall Container Backup Support to clear the previous log files. For instructions, see "Uninstalling Container Backup Support" on page 231 and Chapter 6, "Installing Container Backup Support," on page 197.

Table 109. Solutions to basic problems (continued)

Problem	Solution
A Kubernetes object persists in the terminating state.	Issue the following command: `kubectl delete object object_name --force --grace-period=0` If the object continues to be in the terminating state, issue the following command: `kubectl patch object -n namespace object_name -p '{"metadata":{"finalizers":null}}'` Where: • *object* is a type of object in Kubernetes, such as a deployment, pod, persistent volume (PV), or PVC. • *object_name* is the name of the object. • *namespace* is the name of the namespace that the object is in.
Container Backup Support did not uninstall cleanly.	Manually clean up your environment by issuing the following commands: `kubectl delete namespace baas` `kubectl delete clusterrole baas-controller` `kubectl delete clusterrole baas-scheduler` `kubectl delete clusterrole baas-spp-agent` `kubectl delete clusterrole baas-transaction-manager` `kubectl delete clusterrole aggregate-basreqs-admin-edit` `kubectl delete clusterrolebinding baas-controller` `kubectl delete clusterrolebinding baas-scheduler` `kubectl delete clusterrolebinding baas-spp-agent` `kubectl delete clusterrolebinding baas-transaction-manager` `kubectl delete customresourcedefinition baasreqs.baas.io`

Table 109. Solutions to basic problems (continued)

Problem	Solution
Canceling a copy backup job results in resources being left behind.	Clean up the leftover resources by completing the following steps: 1. Delete the data mover deployment by issuing the following commands: `kubectl get deploy -n namespace` `kubectl delete deploy --all -n namespace` 2. Delete the service account by issuing the following commands: `kubectl get serviceaccount -n namespace` `kubectl delete serviceaccount --all -n namespace` 3. Delete the network policy by issuing the following commands: `kubectl get networkpolicy -n namespace` `kubectl delete networkpolicy --all -n namespace` 4. Delete the PVC and PV: A PVC that is created during a copy backup operation has the following naming convention: `pvc-backup-pvcname-jobid-job_timestamp` Issue the following commands: `kubectl get pvc -n namespace | grep pvc-backup` `kubectl get pvc -n namespace | grep pvc-backup | awk '{print $1}' | xargs kubectl delete pvc -n namespace` If any PVs remain, issue the following commands: `kubectl get pv | grep pvc-backup` `kubectl get pv | grep pvc-backup | awk '{print $1}' | xargs kubectl delete pv` 5. If needed, remove the `volumesnapshot` and `volumesnapshotcontent` objects by issuing the following commands: `kubectl get volumesnapshot -n namespace` `kubectl get volumesnapshotcontent`

Table 109. Solutions to basic problems (continued)

Problem	Solution
On a cluster with IBM Spectrum Scale installed, a snapshot times out after 15 minutes. The following error messages are issued: • `CTGGK3050 Snapshot creation timed out after` *`timeout_duration`* `seconds.` • `CTGGK3060 Restore Volume timed out after` *`timeout_duration`* `seconds.` • `CTGGK3062 Snapshot of PersistentVolumeClaim` *`PersistentVolumeClaim`* `timed out after` *`timeout_length`* `minutes.`	If the load on the CSI driver or storage system is causing snapshots to time out, edit the `baas-configmap` configmap to change the number of snapshots that are requested simultaneously: 1. By using the KUBE_EDITOR variable to set your editor of choice (for example the Vim editor): • On Kubernetes: `KUBE_EDITOR=vim kubectl edit configmap baas-configmap --namespace baas` • On OpenShift: `KUBE_EDITOR=vim oc edit configmap baas-configmap --namespace baas` 2. Edit the **`pvcSnapshotMaxParallel`** field by changing the value from the default of "20" to some other whole number. If your storage system and CSI driver can process more snapshots, you can increase this value instead. The value must be surrounded by double-quotation marks. 3. Run the snapshot job again.
The `baasreq` request for a job remains in the "Requested" state and does not update to the "Ready" state long after the job has completed.	This issue occurs because of a time difference between the Kubernetes or OpenShift cluster and the IBM Spectrum Protect Plus server. Container Backup Support marks a backup request (`baasreq`) as completed by checking the creation time of the `baasreq` on the Kubernetes or OpenShift cluster against the completion time of the job on the IBM Spectrum Protect Plus server. If the system clock on the Kubernetes or OpenShift cluster is ahead of the system clock on the IBM Spectrum Protect Plus server, the completion time of the job always lags behind the `baasreq` creation time. Hence, the `baasreq` is never marked as completed. To resolve this issue: 1. On the Kubernetes or OpenShift cluster and the IBM Spectrum Protect Plus server, enable the Network Time Protocol (NTP) service to synchronize the system time to keep both the cluster and the IBM Spectrum Protect Plus server in sync. The command that you issue depends on the Linux distribution or packages that are installed on the systems. For example, issue the following command: `timedatectl set-ntp true` 2. Run a new backup job. Upon the completion of the backup job, the `baasreq` goes into the "Ready" state.

Related tasks

"Collecting Container Backup Support log files for troubleshooting" on page 663

You can generate debugging log files in the Kubernetes or OpenShift environment to troubleshoot the deployment of Container Backup Support and Container Backup Support operations on the IBM Spectrum Protect Plus server.

Troubleshooting Container Backup Support operations

Troubleshooting procedures are available to help you diagnose and resolve Container Backup Support issues.

The following instructions are provided:

- "Viewing log files" on page 672
- "Troubleshooting issues with snapshot backup jobs" on page 672
- "Troubleshooting issues with copy backup jobs" on page 673
- "Troubleshooting restore jobs" on page 675

Viewing log files

To troubleshoot Container Backup Support issues, start by viewing information in the log files. Log files are available for the transaction manager, controller, and scheduler components of Container Backup Support.

You can view the log files for multiple transaction manager components. For example, view the log file for one of the transaction manager components, issue the following command:

```
kubectl logs -f $(kubectl get pods -n baas | awk '/baas-transaction-manager/ {print $1;exit}')
-n baas -c baas-transaction-manager -f
```

To view the log file for the transaction manager worker, issue the following command:

```
kubectl logs -f $(kubectl get pods -n baas | awk '/baas-transaction-manager/ {print $1;exit}')
-n baas -c baas-transaction-manager-worker -f
```

To view the log file for the controller component, issue the following command:

```
kubectl logs -f $(kubectl get pods -n baas | awk '/baas-controller/ {print $1;exit}') -n baas -f
```

To view the log file for the scheduler component, issue the following command:

```
kubectl logs -f $(kubectl get pods -n baas | awk '/baas-scheduler/ {print $1;exit}') -n baas -f
```

Tip: To help speed up the display of log files, you can add the **--since=***duration* flag to the **kubectl logs** command to return only logs that are newer than a relative duration. You can specify the duration in seconds (*N*s), minutes (*N*m), or hours (*N*h).

For example, to view the log files for the scheduler component that are newer than 3 hours, issue the following command:

```
kubectl logs -f $(kubectl get pods -n baas | awk '/baas-scheduler/ {print $1;exit}') -n baas -f
--since=3h
```

Troubleshooting issues with snapshot backup jobs

If a snapshot backup operation is unsuccessful, you can take a series of actions to diagnose the problem.

Before you begin, ensure that the trace level is set to DEBUG. For instructions on setting the trace level of log files, see "Setting the trace level of log files" on page 664.

Complete the following steps to troubleshoot snapshot backup problems:

1. Ensure that the Container Backup Support log files are available. For instructions about viewing the log files, see "Viewing log files" on page 672.

2. If IBM Spectrum Protect Plus is sending the snapshot request, check the `baas-transaction-manager` container log in the `baas-transaction-manager` pod. In the log file, look for text that is similar to the following example:

```
/createvolumesnapshot/demo/demo-vol01 Begin
Received parameters {'metadata.name': 'k8s18-1004-2222-1727b1c0828',
'spec.snapshotClassName':
'cirrus-csi-rbdplugin-snapclass', 'metadata.labels': {'storage.kubernetes.io/pvc': 'demo-
vol01'}}
```

   The expected snapshot name is value of the `metadata.name` key.

   Next, look for the `createsnapshot` call in the following example:

```
2020-06-03 16:55:43,579[MainThread][kubernetes_api:createsnapshot Line 1056][INFO] -
{'apiVersion':
'snapshot.storage.k8s.io/v1alpha1', 'kind': 'VolumeSnapshot', 'metadata': {'annotations':
{}, 'name':
'k8s18-1004-2222-1727b1c0828', 'namespace': 'demo', 'labels': {'app.kubernetes.io/
component': 'snapshot',
'app.kubernetes.io/managed-by': 'baas', 'app.kubernetes.io/name': 'baas', 'app.kubernetes.io/
version': '10.1.6',
'storage.kubernetes.io/pvc': 'demo-vol01'}, 'spec': {'snapshotClassName': 'cirrus-csi-
rbdplugin-snapclass',
'source': {'kind': 'PersistentVolumeClaim', 'name': 'demo-vol01'}}}
```

3. If an exception is found in Step 2, you might find the following exception in the `createsnapshot` call.

Table 110. Possible snapshot backup exception

Exception	Action
The snapshot does not exist. The snapshot might not be created properly.	Run the following command to determine whether the snapshot was created correctly: `kubectl describe volumesnapshots snapshotname -n namespace`

4. Troubleshoot IBM Spectrum Protect Plus issues by taking the following actions:

   a. In the IBM Spectrum Protect Plus user interface, verify whether any inventory jobs are hung that are preventing all other jobs from being recorded in IBM Spectrum Protect Plus.

   b. Look for the hung job in the list of running jobs or in the job history. Look for job names with the following naming convention:

      - On Kubernetes: k8s_*sla_name*
      - On OpenShift: openshift_*sla_name*

      where the *sla_name* is the name of the SLA policy that is assigned to the PVC.

   c. Check the job logs and resolve any reported issues. For information about how to view and download IBM Spectrum Protect Plus log files, see “Viewing job logs” on page 462.

      Download the package of log files and expand the package. The downloaded package has the following naming convention: `JobLog_job_name_job-timestamp.zip`.

      For detailed information about a job, review the following log files:

      - On Kubernetes, review `command.log` and `JobLog_k8s_sla_name_job_timestamp.csv` files.
      - On OpenShift, review the `command.log` and `JobLog_openshift_sla_name_job_timestamp.csv` file.

Troubleshooting issues with copy backup jobs

If a copy backup job is unsuccessful, you can take a series of actions to diagnose the problem.

Before you begin, ensure that the trace level for log files is set to DEBUG. For instructions on setting the trace level of log files, see “Setting the trace level of log files” on page 664.

Complete the following steps to troubleshoot copy backup problems:

1. Ensure that the Container Backup Support log files are available. For instructions about viewing the log files, see "Viewing log files" on page 672.
2. Verify whether the IBM Spectrum Protect Plus agent is sending a request to the IBM Spectrum Protect Plus scheduler. Open the scheduler log file and look for text that is similar to the following example:

```
Schedule data copy for snapshot: demo:pvc-backup-demo-vol01-1004-1591203980176
```

   The naming convention for the copy backup of the PVC is:

```
namespace:pvc-backup-pvcname-jobid-job_timestamp
```

   Look for the call to the transaction manager to deploy a data mover, such as the following example:

```
url tmCopyBackupRequest: https://baas-transaction-manager:5000/datamover/demo/pvc-backup-
demo-vol01-1004-1591203980176"
```

   If the scheduler is not sending copy backup requests, investigate and resolve the scheduler issues.
3. If the scheduler is sending the snapshot request, check the `baas-transaction-manager` container log in the `baas-transaction-manager` pod. In the transaction manager log file, look for the create data mover call in the text that is similar to the following example:

```
/datamover/demo/pvc-backup-demo-vol01-1004-1591203980176 method=POST
2020-06-03 17:11:26,455[MainThread][main:createdatamover Line 1187][DEBUG] - Creating
deployment
backup-demo-vol01-k8s-k8s18-copy2-1591203980176 for PVC demo:pvc-backup-demo-
vol01-1004-1591203980176
```

   In the `baas-transaction-manager-worker` log in the `baas-transaction-manager` pod, the beginning of the request shows the task ID, COPYBACKUP request, deployment name or data mover name, and the volume name:

```
2020-06-03 17:11:26,589: DEBUG/MainProcess] TaskPool: Apply <function _fast_trace_task at
0x7ff1707ac268> (args:('main.backgroundprocess',
'29606e23-b6e3-4965-8156-930b42c12a25', {'lang': 'py', 'task': 'main.backgroundprocess',
'id': '29606e23-b6e3-4965-8156-930b42c12a25', 'shadow': None, 'eta':
None, 'expires': None, 'group': None, 'retries': 0, 'timelimit': [None, None], 'root_id':
'29606e23-b6e3-4965-8156-930b42c12a25', 'parent_id': None, 'argsrepr': "(
'COPYBACKUP', {'command': 'backup', 'namespace': 'demo', 'deploymentName': 'backup-demo-
vol01-k8s-k8s18-copy2-1591203980176', 'volumename':
'pvc-backup-demo-vol01-1004-1591203980176', 'vSnapIPAddresses': ['9.11.62.84'],
'vSnapMountPath': '/vsnap/vpool1/fs489', 'kafkaAddress':
'baas-kafka-svc.baas:9092', 'kafkaStatusLog': 'backup-demo-vol01-k8s-k8s18-
copy2-1591203980176-status', 'kafkaCommandLog':
'backup-demo-vol01-k8s-k8s18-copy2-1591203980176-command', 'storageClass': None,
'sizeInBytes': None, 'pvclabels': {}})", 'kwargsrepr': '{}', 'origin':
'gen28@baas-transaction-manager-69cffc84fd-95kc4', 'reply_to': '38ff7ee8-718f-3b14-
bd70-8a3f866823f6', 'correlation_id':... kwargs:{})
[2020-06-03 17:11:26,593: DEBUG/MainProcess] Task accepted: main.backgroundprocess[29606e23-
b6e3-4965-8156-930b42c12a25] pid:24

Create datamover demo:backup-demo-vol01-k8s-k8s18-copy2-1591203980176 PVC=pvc-backup-demo-
vol01-1004-1591203980176 isBackup=True

[2020-06-03 17:11:27,127: INFO/ForkPoolWorker-1] Task main.backgroundprocess[29606e23-
b6e3-4965-8156-930b42c12a25] succeeded in 0.5342374939937145s: 0
```

   In the transaction manager log, the following trace statement shows whether the deployment succeeded or failed with the `Get deployment` call:

```
Get deployment backup-demo-vol01-k8s-k8s18-copy2-1591203980176 for PVC demo:backup-demo-
vol01-k8s-k8s18-copy2-1591203980176
```

4. In the scheduler log, verify whether a copy backup has completed by looking for traces that are similar to the following examples:

```
copyBackup volume:demo:pvc-backup-demo-vol01-1004-1591203980176 jobInfoId=1004
ipAddr=[9.11.62.84] fileLocation= volumeSize=1073.741824 nextRunTime=1591290380176"
```

```
Setting backup to complete for copyBackup: demopvc-backup-demo-
vol01-1004-1591203980176:1591203980176
```

5. If an exception is found, you might find the following exceptions in the COPYBACKUP request.

Table 111. Possible copy backup exceptions

Exception	Action
The snapshot does not exist. The snapshot might not be created properly.	Run the following command to determine whether the snapshot was created correctly: `kubectl describe volumesnapshots snapshotname -n namespace`
The deployment does not exist. The data mover might not be created properly.	For more information about the issue, get the data mover name from the error message and run the following command: `kubectl describe deploy backup-pvcname-jobname-job_timestamp -n namespace`

6. Troubleshoot IBM Spectrum Protect Plus issues by taking the following actions:
   a. In the IBM Spectrum Protect Plus user interface, verify whether any inventory jobs are hung that are preventing all other jobs from being recorded in IBM Spectrum Protect Plus.
   b. Look for the hung job in the list of running jobs or in the job history. Look for job names with the following naming convention:
      - On Kubernetes: k8s_*sla_name*
      - On OpenShift: openshift_*sla_name*

      where the *sla_name* is the name of the SLA policy that is assigned to the PVC.
   c. Check the job logs and resolve any reported issues. For information about how to view and download IBM Spectrum Protect Plus log files, see “Viewing job logs” on page 462.

      Download the package of log files and expand the package. The downloaded package has the following naming convention: JobLog_*job_name_job-timestamp*.zip.

      For detailed information about a job, review the following log files:
      - On Kubernetes, review command.log and JobLog_k8s_*sla_name_job_timestamp*.csv files.
      - On OpenShift, review the command.log and JobLog_openshift_*sla_name_job_timestamp*.csv file.

Troubleshooting restore jobs

If a restore job is unsuccessful, you can take the following actions to diagnose the problem.

Before you begin, ensure that the trace level is set to DEBUG. For instructions on setting the trace level of log files, see “Setting the trace level of log files” on page 664.

Complete the following steps to troubleshoot restore job problems:

1. Ensure that the Container Backup Support log files are available. For instructions about viewing the log files, see “Viewing log files” on page 672.
2. Check for errors in the IBM Spectrum Protect Plus server restore job log named onDemandRestore_*timestamp*.

If restore job was started from the **kubectl** command line, you can find the name of restore job in the BaasReq objects while the restore job is in progress by issuing the following command:

```
kubectl describe baasreq restore_request_name -n namespace | grep Inprogress
```

Look for output that is similar to the following example:

```
Inprogress: onDemandRestore_1591384200276
```

3. If you restored data from the **kubectl** command line, check whether the restore request was invalidated due to invalid parameters in the YAML configuration file. Use the **kubectl describe** command to check the restore status (Restorestatus) in the output.

   If the value in the Restorestatus field is Invalid, the Errmsg field will show the reason why the restore request was invalidated. In the following example, an incorrect value was specified in the **VolumeStorageClass** parameter in the YAML file.

   For example, to show the restore status of restore request copy-restore-pvc02 in namespace test, issue the following command:

   ```
   kubectl describe baasreq copy-restore-pvc02 -n test
   ```

   The output is similar to the following example:

   ```
   Name:          copy-restore-pvc02
   Namespace:     test
   Labels:        <none>
   Annotations:   <none>
   API Version:   baas.io/v1alpha1
   Backupstatus:  None
   Errmsg:        VolumeStorageClass invalid
   Kind:          BaaSReq
   Metadata:
     Creation Timestamp:  2020-06-05T19:51:29Z
     Generation:          2
     Resource Version:    4396987
     Self Link:           /apis/baas.io/v1alpha1/namespaces/test/baasreqs/copy-restore-pvc02
     UID:                 418cc8d5-7347-47ed-9436-9fe49f69b42a
   Restorestatus:         Invalid
   Spec:
     Inprogress:    None
     Origreqtype:   restore
     Pvcname:       pvc02
     Requesttype:   restore
     Restorepoint:  2020-06-05 17:22:35
     Restoretype:   copy
     Storageclass:  cirrus19-csi-rbd-sc
     Targetvolume:  pvc02-restored
   Volumename:     pvc02
   Events:          <none>
   ```

   To recover from this type of error, delete the invalid restore request, correct the YAML file, and re-create the restore request.

4. Review the error messages in IBM Spectrum Protect Plus server onDemandRestore_*timestamp* log. The error messages are usually sufficient to help you diagnose the problem.
5. To further troubleshoot a snapshot restore, you can look for traces in the application agent baas-spp-agent job log that are similar to the following example:

   ```
   DEBUG pid:3402 MainThread restoreDatabase: Starting restore of snapshot
   spp-1275-2213-17285db4b80 to test-snap-restore-pvc1
   DEBUG pid:3402 MainThread restoreDatabase: Restoring pvc labels {'department': 'sales',
   'team': 'green'}
   DEBUG pid:3402 MainThread restoreDatabase: Restoring snapshot named spp-1275-2213-17285db4b80
   DEBUG pid:3402 MainThread sendRestoreRequest: Sending restore request
   to https://baas-transaction-manager:5000/restorevolumebackup/test/test-snap-restore-pvc1?
   storageclass=cirrus-csi-rbd-sc&restoretype=FAST
   DEBUG pid:3402 MainThread sendRestoreRequest: Get restore response
   ```

Check the `baas-transactionmanager` container log in the `baas-transaction-manager` pod. In the log file, look for text that is similar to the following example:

```
/restorevolumebackup/test/test-snap-restore-pvc1 snapshot:spp-1275-2213-17285db4b80
restoretype:FAST
storageclass: cirrus-csi-rbd-sc
```

6. To further troubleshoot copy restore, you can look for traces in application agent `baas-spp-agent` job log that are similar to the following example:

```
JOBLOG_SUMMARY pid:4219 MainThread jobsummary: <CTGGK3005> Starting to restore a persistent
volume.
 DEBUG pid:4219 MainThread copyRestore: Starting restore of database cirrus19:test:pvc02
 DEBUG pid:4219 MainThread getPVC: PVC test:test-copy-restore-pvc02 not found.
 DEBUG pid:4219 MainThread copyRestore: PVC does not exist, the restore can continue.
 DEBUG pid:4219 MainThread createDatamover: PVC labels {'department': 'sales', 'team':
'green'}
 INFO pid:4219 MainThread createDatamover: Create datamover request to https://baas-
transaction-manager:5000/datamover/test/test-copy-restore-pvc02
```

Check the `baas-transactionmanager` container log in the `baas-transaction-manager` pod. In the log file, look for text that is similar to the following example:

```
main:createdatamover Line 1187][DEBUG] - Creating deployment
restore-pvc02-ondemandrestore-1591390864757-1591390865107 for PVC test:test-copy-restore-
pvc02
```

In the `transaction-manager-worker` log file, look for text that is similar to the following example:

```
DEBUG/ForkPoolWorker-1] Restore worker
DEBUG/ForkPoolWorker-1] Create datamover test:restore-pvc02-
ondemandrestore-1591390864757-1591390865107
PVC=test-copy-restore-pvc02 isBackup=False
```

Related tasks

“Setting the trace level of log files” on page 664
You can set the trace level of local log files to help troubleshoot issues that you might encounter in Container Backup Support.

Related reference

“Troubleshooting quick reference” on page 667
Solutions to basic Container Backup Support problems are provided.

Chapter 21. Product messages

IBM Spectrum Protect Plus components send messages with prefixes that help to identify which component they come from. Use the search option to find a particular message by using its unique identifier.

Messages consist of the following elements:

- A five-letter prefix.
- A number to identify the message.
- Message text that is displayed on screen and written to message logs.

Tip: Use your browser's search capability by using `Ctrl+F` to find the message code you are looking for.

The following example contains the Db2 agent prefix. When you click `More`, extra details that explain the reason for the message are shown.

```
Warning
Apr 16, 2019
9:14:37 AM
GTGGH0098
[myserver1.myplace.irl.ibm.com]
Database AC7 will not be backed up as it is ineligible for the backup operation. More
```

IBM Spectrum Protect Plus message prefixes

Messages have different prefixes to help you to identify the component that issues the message.

The following table identifies the prefix that is associated with each component.

Table 112. Messages prefixes by component

Prefix	Component
CTGGA	IBM Spectrum Protect Plus
CTGGE	IBM Spectrum Protect Plus for Microsoft SQL Server
CTGGF	IBM Spectrum Protect Plus for Oracle
CTGGG	IBM Spectrum Protect Plus for Microsoft Exchange Server
CTGGH	IBM Spectrum Protect Plus for IBM Db2
CTGGI	IBM Spectrum Protect Plus for MongoDB
CTGGK	IBM Spectrum Protect Plus for Containers
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.

© Copyright IBM Corp. 2017, 2021

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

© Copyright IBM Corp. 2017, 2021

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).

© Copyright IBM Corp. 2017, 2021

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

© Copyright IBM Corp. 2017, 2021

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.

© Copyright IBM Corp. 2017, 2021

Index

A

Access control
  MongoDB 562
accessibility features 683
ad hoc jobs
  creating 627
Add Db2 partitions 500
adding
  Amazon EC2 account 375
  Hyper-V servers 360
  identities 656
  LDAP server 294
  Oracle application servers 588
  sites 291
  SMTP server 295
  SQL Server application servers 601
  vCenter Server instances 335
  virtual disks to a vCenter virtual machine 314
  vSnap servers 158
Adding a filesystem 384
Adding Db2 500
Adding MongoDB 564
Administrative Console, logging on to 312
Advanced backup options 165
Amazon EC2
  accounts
    adding 375
  backup job, creating 376
  detecting resources 376
  IAM user, creating 373
application server
  Db2 497
application-consistent backups
  namespace 440
application-consistent restore 454
AWS EC2
  restore job, creating 378

B

backing up
  container data 468
  Kubernetes containers 415
  OpenShift containers 436
Backing up
  Db2 504
  file system data 388
backing up container data
  application-consistent 440
  by label 475
  by namespace 478
  namespace 417, 438
  on demand 472, 474
  scheduling 415, 436, 468
backing up resources 478
backup jobs
backup jobs *(continued)*
  creating
    Amazon EC2 376
    Hyper-V 363
    IBM Spectrum Protect Plus 615
    Oracle 590
    SQL Server 602
    VMware 340
  excluding virtual disks from 367
  excluding VMDKs from 344
  rerunning
    on demand 627
  starting
    on demand 621
    on schedule 324, 328, 329
backup policies, *See* SLA policies
backup storage
  advanced options, managing 165
  storage options, managing disks 160
  storage options, managing partners 164
backup storage server
  storage options, managing 162, 165
backup types
  Container Backup Support 403
backup-by-label
  Container Backup Support 475
backup-by-namespace
  Container Backup Support 478
Beta program
  advantages xv
  overview xv

C

certificate
  adding 298
  deleting 299
cloud provider
  deleting 277
  editing 277
cloud server
  adding a Microsoft azure cloud resource 275
  adding an Amazon S3 272
  adding an IBM Cloud Object Storage resource 273
  adding an s3 compatible cloud resource 276
cold-data-cache storage pool 280
collecting debugging log filesContainer Backup Support 663
configuration parameters
  Container Backup Support 232
Configuring backup storage
  storage options, adding disks 161
container
  installing
    on OpenShift 131
Container Backup Support
  baas requests 466
  baas-options.sh 206

Container Backup Support *(continued)*
baas-values.yaml 206
backing up container data 468
backing up PVCs by label 475
backing up PVCs by namespace 478
backing up resources by label 475
backup and restore types 403
backup status 488
backup-by-label 475
backup-by-namespace 478
cascading actions 205
collecting debugging log files 663
configuration parameters 232
copy backup 468
copy restore 481
creating reports 462
deleting backups 489
deployment logs 663
destroy request 489
displaying log files 672
enable tracing 664
encryption 407
expiring jobs 429, 461
fast restore 481
installation prerequisites 197
installing 197
installing from IBM Helm Charts Repository 218
installing in an airgap environment 214
installing using Helm 3 205
managing jobs 486
monitoring jobs 462
multitenancy 407
overview 401
request types 466
restore status 488
restoring data 481
running reports 462
scheduling backups 468
security 407
set up installation variables 206
setting trace levels 664
SLA policies 405, 468
snapshot backup 468, 472
system requirements 75
troubleshooting 662
troubleshooting backup jobs 672
troubleshooting restore jobs 672
uninstalling 231
user roles 406
verifying Metrics Server 204
viewing backup history 463
viewing backup status 486
viewing backup utilization 464
viewing job logs 462
viewing restore status 486
viewing trace logs 665
copy backup
Container Backup Support 468
copy backups
Kubernetes 415
namespace 417, 438
OpenShift 436
copy restore
application-consistent 454
copy restore *(continued)*
container data 481
OpenShift Virtualization VMs 444
copying data to tape 280
creating
reports 639
resource groups 644
roles 650
SLA policies 324, 328, 329
users
individual 653
LDAP group 653
VADP proxies 346
creating reports
container backups 462

D

data copy to tape
configuring 280
data protection 286, 287
Db2
system requirements 83
Db2 log backup 510
DEFINE STGPOOL command 280
defining SLA backups
Kubernetes 415
OpenShift 436
deleting
identities 657
jobs 626
LDAP server 297
resource groups 648
roles 652
sites 293
SLA policies 334
SMTP server 297
users 654
deleting backups
Container Backup Support 489
deployment log filesContainer Backup Support 663
destroying backupsContainer Backup Support 489
Detailed process logs
Microsfot 365 494
Detecting
Db2 502
file system resources 386
disability 683

E

early availability updates, obtaining and applying 253
editing
identities 656
jobs and job schedules 625
LDAP server 296
resource groups 647
roles 652
settings 296
sites 292
SLA policies 333
SMTP server 296
users 654

efix 253
enable tracing
    Container Backup Support 664
Exchange Server
    system requirements 89
expire job session 616
expiring jobs
    Container Backup Support 429, 461

F

fast restore
    container data 481
fenced network, creating 357
file systems
    system requirements 71
files
    restoring 380
    searching for 681
Finding Db2 502
Finding file system drives 386
firewalls 122

G

global preferences
    configuring 304

H

Hyper-V
    adding 360
    backup job, creating 363
    backup job, excluding virtual disks from an SLA policy 367
    installing on virtual appliance 128
    restore job, creating 367
    servers
        detecting resources for 362
        enabling WinRM 362
        testing connection to 362
    virtual appliance
        accessing 313

I

IBM Cloud Pak for Multicloud Management, integrating IBM Spectrum Protect Plus 25
IBM Knowledge Center xi
IBM Spectrum Protect Operations Center
    Accessing from IBM Spectrum Protect Plus 20
    adding IBM Spectrum Protect Plus to 20
    monitoring IBM Spectrum Protect Plus from 20, 24
    starting from IBM Spectrum Protect Plus 24
    URL, setting 23
IBM Spectrum Protect Plus
    installing
        on OpenShift 133
    updating
        on OpenShift 244
IBM spectrum protect server
    adding a repository server 289
    registering a repository server 289
identities
    adding 656
    deleting 657
    editing 656
installation prerequisites
    Container Backup Support 197, 198
installing
    as a container
        on OpenShift 131
    as a virtual appliance 125
    Container Backup Support 197
    download packages, obtaining 125
    IBM Spectrum Protect Plus
        on OpenShift 133
    post installation tasks 121
    Velero 198
    virtual appliance
        on Hyper-V 128
        on VMware 126
    vSnap servers
        Hyper-V environment 155
        physical environment 153
        VMware environment 154
installing from IBM Helm Charts Repository
    Container Backup Support 218
installing in a airgap environment
    Container Backup Support 214
installing using Helm 3
    Container Backup Support 205
inventory
    file systems 386
iSCSI utilities
    installing 176

J

jobs
    canceling 626
    concurrent, viewing 625
    creating 620
    deleting 626
    editing 625
    logs
        downloading 625
        viewing 625
    names of 619
    pausing 625
    progress, viewing 624
    releasing 625
    rerunning 627
    schedules, editing 625
    starting
        on demand 621
        on schedule 324, 328, 329
    types of 619
    viewing 622
Jobs and Operations 619

K

key
    adding 297, 299

key *(continued)*
deleting 298, 301
keyboard 683
keys and certificates
for IBM Spectrum Protect Plus user interface 302
for secondary storage and resources 297, 301
Knowledge Center xi
Kubernetes
clusters
manually register 409
modify properties 409
detecting namespace resources 411
Kubernetes cluster
detecting resources 411
testing connection to 414
Kubernetes containers
backing up 415

L

LDAP
group, creating a user account for 653
server
adding 294
deleting 297
settings, editing 296
Log archiving
Db2 510
logs
audit
downloading 640
viewing 640
system
downloading 660
viewing 660

M

managing jobs
container backups and restores 486
manually registering
Kubernetes clusters 409
OpenShift clusters 431
message
prefixes 679
messages 679
Microsoft 365 491
Microsoft 365 log files
Detailed 494
modifying properties
Kubernetes clusters 409
OpenShift clusters 431
MongoDB
system requirements 96
MongoDB application server 561
monitoring
container backup jobs 462
multitenancy
Container Backup Support 401, 407

N

network
network *(continued)*
testing 303, 304
Network configuration 162
New in IBM Spectrum Protect Plus Version Version 10.1.8 xiii
NICs 162

O

object client 286, 287
Object Storage
Amazon S3 272
offline updates, virtual appliance 242
on-demand backup
containers 472, 474
online updates, virtual appliance 242
OpenShift
clusters
manually register 431
modify properties 431
detecting project resources 433
installing as a container 131
installing IBM Spectrum Protect Plus 133
preparing to restore PVCs and resources 443
updating IBM Spectrum Protect Plus server 244
OpenShift cluster
detecting resources 433
testing connection to 436
OpenShift containers
backing up 436
Operations Center
Accessing from IBM Spectrum Protect Plus 20
adding IBM Spectrum Protect Plus to 20
monitoring IBM Spectrum Protect Plus from 20, 24
starting from IBM Spectrum Protect Plus 24
URL, setting 23
Ops Manager
MongoDB 566
Oracle
application servers
adding 588
detecting resources for 590
testing connection to 590
backup job, creating 590
multithreaded databases 588
restore job, creating 593
system requirements 106
overview
Container Backup Support 401

P

password
superuser
changing 655
preferences
global
configuring 304
preparing to restore
OpenShift PVCs and resources 443
prerequisites
Db2 497
file systems 383
MongoDB 561

Prerequisites
MongoDB 562
product overview 19
publications xi

Q

quick start 255, 264, 268

R

RBAC
MongoDB 562
registering
Kubernetes clusters 409
OpenShift clusters 431
vSnap servers 158
removing SLA policy assignments
Kubernetes 415
OpenShift 436
repair vSnap 176
Replication partners 164
reports
custom, creating 639
file systems 632
Microsoft 365 632
running
on demand 638
on schedule 640
running VM 636
types of
backup storage utilization 631
protection 632
system 635
repository server provider
deleting 291
editing 290
request types
Container Backup Support 466
rerunning
jobs
on demand 627
Rescan
After expanding storage 162
Rescan vSnap 162
resource groups
creating 644
deleting 648
editing 647
types of 645
restore jobs
creating
AWS EC2 378
Hyper-V 367
IBM Spectrum Protect Plus 615
Oracle 593
SQL Server 607
VMware 350
running
AWS EC2 378
Hyper-V 367
Oracle 593
SQL Server 607
restore jobs *(continued)*
running *(continued)*
VMware 350
restore points, deleting 617
restore points, managing 616
restore types
Container Backup Support 403
restoring
OpenShift 449, 454
Restoring
Db2 511, 516, 519
file system 395
restoring container data
Container Backup Support 481, 484
restoring data 420, 424, 443, 445
Restoring Db2
Alternate instance 519
Original instance 516
restoring persistent volumes
Kubernetes 420
OpenShift 443, 445
restoring VMs
OpenShift Virtualization support 444
roles
creating 650
deleting 652
editing 652
permission types 650
running reports
container backup jobs 462

S

Schedule jobs
Backup 506, 527, 569
scheduling backups
Container Backup Support 468
Kubernetes 415
OpenShift 436
scripts for backup and restore operations
uploading 629
security features
Container Backup Support 407
service level agreement, *See* SLA policies
service level agreements
Container Backup Support 405
set up installation variables
Container Backup Support 206
Setting Db2
SLA options 509
setting trace levels
Container Backup Support 664
sites
adding 291
deleting 293
editing 292
throttling 291, 292
SLA 506, 527, 569
SLA options
Db2 509
SLA policies
adding 324, 328, 329
Container Backup Support 405, 468
deleting 334

SLA policies *(continued)*
editing 333
SLA policy 471
SMTP
server
adding 295
deleting 297
settings, editing 296
snapshot backup
Container Backup Support 468
containers 472, 474
snapshot backups
Kubernetes 415
namespace 417, 438
OpenShift 436
snapshot restore
application-consistent 454
OpenShift Virtualization VMs 444
snapshot retention 616
sponsor user program
advantages xv
overview xv
SQL Server
application servers
adding 601
detecting resources for 602
testing connection to 602
backup job, creating 602
requirements for data protection 600
restore job, creating 607
system requirements 113
SSL certificate, uploading 302
starting
IBM Spectrum Protect Plus 257
jobs
on demand 621
on schedule 324, 328, 329
superuser
changing password 655
changing password and name 655
system requirements
components 27
Container Backup Support 75
Db2 83
Exchange Server 89
file index and restore 64
file systems 71
hypervisors 57
MongoDB 96
Oracle 106
SQL Server 113

T

Testing connection
Db2 503
Testing connection file systems 387
time zone, setting 313
troubleshooting
Container Backup Support 662
Container Backup Support operations 672
displaying Container Backup Support logs 672

U

Uninstalling
Container Backup Support 231
updating
IBM Spectrum Protect Plus
on OpenShift 244
Updating
vSnap server 248
user access 12, 643
user roles
Container Backup Support 406
users
deleting 654
editing 654
individual, creating 653
LDAP group, creating 653
resource groups
creating 644
deleting 648
editing 647
types of 645
roles
creating 650
deleting 652
editing 652
permission types 650

V

VADP proxies
creating 346
options, setting 348
uninstalling 349
updating 251
Velero
installing 198
verifying Metrics Server
Container Backup Support 204
viewing backup history
container backups 463
viewing backup status
Container Backup Support 486, 488
viewing backup utilization
container backups 464
viewing job logs
container backups 462
viewing restore status
Container Backup Support 486, 488
viewing trace logs
Container Backup Support 665
virtual appliance
accessing
in Hyper-V 313
in VMware 313
adding a disk to 314
adding storage capacity 315
installing
on Hyper-V 128
on VMware 126
Virtual appliance
updating 242
virtual environments 286, 287
VMware

VMware *(continued)*
backup job, creating 340
backup job, excluding VMDKs from SLA policy 344
installing on virtual appliance 126
restore job
creating a fenced network 357
restore job, creating 350
vCenter Server instances
adding 335
vCenter Server, detecting resources 339
vCenter Server, testing connection to 339
virtual appliance
accessing 313
virtual machine privileges, required 336
vSnap
updating 250
vSnap recovery 176
vSnap server
administering
kernel headers
kernel tools 182
network administration 181
storage administration 178
user administration 177
change throughput 175
editing 159
initializing
advanced 170
simple 170
storage pools, expanding 174
Unregistering 159
vSnap servers
adding 158
installing
Hyper-V environment
155
physical environment 153
VMware environment 154
registering 158
uninstalling 157

W

WinRM, enabling for connection to Hyper-V servers 362

Y

YAML files
Container Backup Support 466



Product Number: 5737-F11