

IBM Spectrum Protect Plus Protecting Red Hat OpenShift Containerized Environments

Axel Westphal

Pia Nymann

Neil Patterson

Alexis Jolin

Julien Sauvanet

Gauthier Siri

Joerg Walter



Storage



IBM Redbooks

**IBM Spectrum Protect Plus: Protecting Red Hat
OpenShift Containerized Environments**

December 2022

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Second Edition (December 2022)

This edition applies to Version 10.1.12 of IBM Spectrum Protect Plus.

© Copyright International Business Machines Corporation 2021, 2022. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	ix
Now you can become a published author, too!	xi
Comments welcome	xi
Stay connected to IBM Redbooks	xii
Chapter 1. Introducing containers	1
1.1 Containers overview and concepts	2
1.1.1 Benefits of containers	3
1.1.2 Container concepts	3
1.2 Orchestrating containers	4
1.2.1 Red Hat OpenShift introduction	4
1.3 Protecting a Red Hat OpenShift environment	7
1.3.1 Protected items	8
1.3.2 IBM Spectrum Protect Plus Container Backup support	9
Chapter 2. IBM Spectrum Protect Plus architecture	11
2.1 Overview	12
2.2 Architecture	12
2.3 Components and terminology	15
2.3.1 IBM Spectrum Protect Plus server	15
2.3.2 vSnap Backup Storage server	15
2.3.3 VADP proxy server with vSnap backup storage server	16
2.3.4 Open Snap Store Manager storage server	16
2.3.5 VADP proxy with OSSM	16
2.3.6 Backup clients	16
2.3.7 Service-level agreement backup policies	17
2.3.8 Site	18
2.3.9 Extra copies	19
2.3.10 Replication	19
Chapter 3. Installing IBM Spectrum Protect Plus as a containerized application	21
3.1 Understanding the environment and requirements	22
3.1.1 Containers and pods	23
3.1.2 Persistent storage	24
3.1.3 Services	26
3.1.4 Routes	26
3.2 Preparing Red Hat OpenShift Container Platform	27
3.2.1 Requirements	27
3.2.2 Step 3: Adding the IBM registry to the IBM Spectrum Protect Plus operator	30
3.2.3 Steps 4 and 5: Creating project and image pull secret	32
3.3 Deploying IBM Spectrum Protect Plus by using an operator	38
3.3.1 Creating IBM Spectrum Protect Plus instance	38
3.3.2 Reviewing the IBM Spectrum Protect Plus deployment	42
3.4 Configuring an IBM Spectrum Protect Plus Server catalog backup	43
3.4.1 Defining a cloud storage to be used by IBM Spectrum Protect Plus	43

3.4.2	Defining an SLA policy to protect the IBM Spectrum Protect Plus catalog	45
3.5	Upgrading an IBM Spectrum Protect Plus containerized server	48
3.6	IBM Spectrum Protect Plus deployment troubleshooting	52
3.6.1	Pods information	52
3.6.2	Accessing the pods and container logs	53
3.6.3	Monitoring the IBM Spectrum Protect Plus Server resource usage	57
3.6.4	Changing the hostname of IBM Spectrum Protect Plus server	57
Chapter 4.	Container Backup Support	61
4.1	Overview	62
4.1.1	Version history	62
4.1.2	User roles	62
4.2	Components	64
4.2.1	Container Backup Support (BaaS) namespace	64
4.2.2	Data mover pods	66
4.2.3	Container Storage Interface	66
4.3	Red Hat OpenShift prerequisites and supported environments	68
4.3.1	Installation types	68
4.3.2	Supported Red Hat OpenShift architecture	69
4.3.3	Supported storage types	70
4.3.4	Red Hat OpenShift cluster prerequisites	72
4.4	IBM Spectrum Protect Plus and storage prerequisites	72
4.4.1	General prerequisites	72
4.4.2	Managing connection performance in Red Hat OpenShift	74
4.4.3	Vertical Pod Autoscaler	75
4.4.4	Cloud storage for direct backup operations	75
4.4.5	Communication ports	76
4.5	Container backup and restore types	77
4.5.1	Backup types	77
4.5.2	Restore types	78
4.6	Service-level agreement policies for container backup	80
4.7	Container Backup Support security features	82
4.8	IBM Spectrum Fusion and IBM Cloud Pak for Data	83
Chapter 5.	Implementing Container Backup Support	85
5.1	Validating the prerequisites	86
5.1.1	Preparing IBM Spectrum Protect Plus backup servers	86
5.2	Installing in an online environment	87
5.2.1	Obtaining the IBM Container Software Library Entitlement key	87
5.2.2	Setting up installation variables	88
5.2.3	Editing the cluster global pull secret	93
5.2.4	Creating the Container Backup Support namespace	96
5.2.5	Creating the image pull secret	97
5.2.6	Creating the general-purpose secret for Container Backup Support	99
5.2.7	Exchanging a certificate between IBM Spectrum Protect Plus and Red Hat OpenShift Container Platform	101
5.2.8	Adding the online catalog source	104
5.2.9	Installing the Container Backup Support (BaaS) operator	107
5.2.10	Creating a Container Backup Support (BaaS) instance	111
5.2.11	Registering the Red Hat OpenShift cluster in IBM Spectrum Protect Plus server manually	115
5.2.12	Updating Container Backup Support	120
5.3	Installing in an air-gapped environment	120

Chapter 6. Using Container Backup Support	121
6.1 Sample application	122
6.2 Application-consistent backups	127
6.2.1 Velero hooks	127
6.2.2 Configuring the sample application with backup hooks	128
6.3 Creating a service-level agreement policy	130
6.3.1 General parameters	131
6.3.2 Configuring a backup policy	132
6.3.3 Defining a replication policy	134
6.3.4 Additional Copies	134
6.3.5 Configuring object storage archival	135
6.4 Red Hat OpenShift Container Platform backups	136
6.4.1 Performing container backups	137
6.4.2 Monitoring backup jobs	141
6.5 Performing restores	142
6.5.1 Deciding which objects must be restored	142
6.5.2 Restoring the application	142
6.5.3 Application considerations for necessary rework after a restore	146
Chapter 7. Red Hat OpenShift cluster disaster recovery solution	147
7.1 General disaster recovery considerations	148
7.2 Protecting the environment	149
7.2.1 Protecting the infrastructure, and complementary systems and services	149
7.2.2 Protecting cluster resources and persistent data	150
7.3 Preparing for a DR restore with IBM Spectrum Protect Plus	152
7.4 Cluster-scoped resources recovery	153
7.4.1 Temporarily pausing the machine-config operator	153
7.4.2 Restoring cluster-scoped resources	154
7.4.3 Monitoring the progress of the restore job	156
7.5 Application namespaces recovery	158
Related publications	161
IBM Redbooks	161
Online resources	161
Help from IBM	162
Abbreviations and acronyms	163

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

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 provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, 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 <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Db2®	IBM Cloud®	IBM Spectrum®
DS8000®	IBM Cloud Pak®	Redbooks®
FlashCopy®	IBM FlashSystem®	Redbooks (logo)  ®
IBM®	IBM Services®	Tivoli®

The following terms are trademarks of other companies:

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 the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Ansible, Ceph, OpenShift, Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

VMware, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM® Redpaper publication describes support for Red Hat OpenShift Container Platform application data protection with IBM Spectrum® Protect Plus. It explains backup and restore operations for persistent volume data by using the Container Storage Interface (CSI) plug-in.

The paper starts with an introduction and overview of containers and Red Hat OpenShift, followed by a quick review of the IBM Spectrum Protect Plus architecture and functions. Readers who are familiar with any one of these topics can skip the corresponding chapters.

Chapter 3, “Installing IBM Spectrum Protect Plus as a containerized application” on page 21 reviews the specific case of deploying IBM Spectrum Protect Plus as a containerized application. Chapter 4, “Container Backup Support” on page 61 - Chapter 6, “Using Container Backup Support” on page 121 review the support for containers backup, the CSI, and several options for installing, configuring, and the use of the backup as a service (BaaS) components as implemented and supported by IBM Spectrum Protect Plus.

Chapter 6, “Using Container Backup Support” on page 121 also reviews the use case of backing up and restoring a containerized application in a consistent way by implementing pre- and post-backup tasks.

Finally, Chapter 7, “Red Hat OpenShift cluster disaster recovery solution” on page 147 describes a solution for a Red Hat OpenShift cluster recovery that combines Velero and IBM Spectrum Protect Plus restore operations.

Authors

This paper was produced by a team of IBM specialists from around the world.



Axel Westphal is an IBM Certified IT Specialist at the IBM ESCC in Kelsterbach, Germany. He joined IBM in 1996, working for IBM Global Services as a systems engineer. His areas of expertise include setting up and demonstrating IBM System Storage products and solutions in various environments. He has written several storage white papers and co-authored several IBM publications.



Pia Nymann is an IBM Certified Senior IT Specialist working with IT infrastructure in IBM technology services. She has more than 25 years of experience in the IT industry and has several years of IBM Spectrum Protect (formerly IBM Tivoli® Storage Manager) software experience, which includes designing and implementing backup and recovery solutions on various platforms and applications. Pia has worked with various areas of the storage management discipline, including IBM Spectrum Control, IBM Spectrum Scale, storage analysis, and data encryption of storage systems and applications. She also is the co-author of other IBM Redbooks® publications.



Neil Patterson is an executive architect in the World Wide SWAT organization for Cloud Paks. He also has over 30 years experience within IT. Currently, he provides thought leadership to customers around the globe that are implementing IBM Cloud Pak®.




Alexis Jolin is a Storage IT Specialist from IBM France in the Montpellier Client Engineering Center. He started in 2017 as apprentice engineer and became member of the Montpellier Storage Demo team in 2020. Alexis is focused on developing demonstrations of storage solutions to help IBM and Business Partner sales teams to show the value of IBM Storage solutions. More precisely, he is focused on storage solutions for containerized environments. such as Red Hat OpenShift.



Julien Sauvanet is an OpenGroup Certified Expert IT Specialist, working at IBM for more than 15 years. He spent more than 10 years working in IBM Services®, helping with IBM Spectrum Protect design and deployment, and helping customers to solve their challenges around resiliency. Julien is focusing on helping customers with their IBM Spectrum Protect and Plus deployment, working as a Technical Advisor in IBM Systems organization. He has co-authored several IBM Redbooks publications.



Gauthier Siri is a Storage IT Specialist from IBM France who is based at IBM Systems Center, Montpellier. He started 13 years ago as an administrator and is now part of a worldwide pre-sales team, whose goal is to help IBM and Business Partner sales teams demonstrate IBM Storage added value. With a strong block storage background, his natural curiosity led him to new technologies, from automation to containerization, and to support solutions, such as IBM Spectrum Discover and IBM Storage solution for Red Hat OpenShift Container Platform.



Joerg Walter is an IBM Certified IT Specialist at the IBM EMEA Storage Competence Center (ESCC) in Kelsterbach, Germany. Since he joined IBM in 2000, he has worked in various technical positions, starting with network engineering and design, but changed focus to storage, data protection, and retention in 2005. Since 2011, Joerg is a member of IBM's Systems Storage Lab Services team. In this role, he supports customers with the planning and implementation of DP&R related hardware and software products, such as IBM Spectrum Protect, IBM Spectrum Protect Plus, IBM Spectrum Protect Snapshot, and physical and virtual tape.

Thanks to the following people for their contributions to this project:

Wade Wallace and Erica Wazewski
IBM Redbooks

Frank Lautenbach, Dominic Mueller-Wicke, Jim Smith, Christian Burns, Markus Fehling, Kai Ruth, Adam Young
IBM

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- Send your comments in an email to:

redbooks@us.ibm.com

- Mail your comments to:

IBM Corporation, IBM Redbooks
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>



Introducing containers

This chapter presents the conceptual foundations of containers, the evolution of computing from bare metal computers to virtual machines (VMs), and then, containers. It also describes the Red Hat OpenShift Container Platform that is based on Kubernetes.

The chapter also addresses the concept of persistent data, storage for containers, the approach that is used to back up containers, and the challenges that can be encountered when backing up a container environment with IBM Spectrum Protect Plus.

This chapter includes the following topics:

- ▶ 1.1, “Containers overview and concepts” on page 2
- ▶ 1.2, “Orchestrating containers” on page 4
- ▶ 1.3, “Protecting a Red Hat OpenShift environment” on page 7

1.1 Containers overview and concepts

For decades, most users are familiar with the physical (bare-metal) server that they rely on to run a single monolithic application by using the compute and storage resources of one computer in a beige steel box. Often times, a rack was used that was connected to a network switch and storage-area network.

Later on came the VM, where the compute and storage resources of the physical server were abstracted by software. This configuration enabled multiple virtually running computer systems with their own CPU, memory, network interface, storage, and operating systems. Although they worked independently of each other, they always shared hardware. They are still used and relevant, but now are evolving to a new way of abstracting the compute and storage hardware into what are called *containers*.

These “containers” are an executable unit of software in which application code is packaged, along with its libraries and dependencies, and stored in a repository as an image. By doing so, it can be run anywhere, whether it be on desktop, traditional IT, or the cloud. Containers are small, fast, and portable because unlike a VM, containers do not need to include a guest operating system in every instance. Instead, they can use the features and resources of the host operating system.

The three deployment model’s characteristics are shown in Figure 1-1. Traditional deployment workloads examples include physical and bare metal deployments, files, databases, and applications. The virtualized deployment consists of multiple VMs, each with a separate operating system on a single physical server’s CPU cores. The container-based deployment allows multiple applications to be within the same operating system, but with their own independent sets of dependencies.

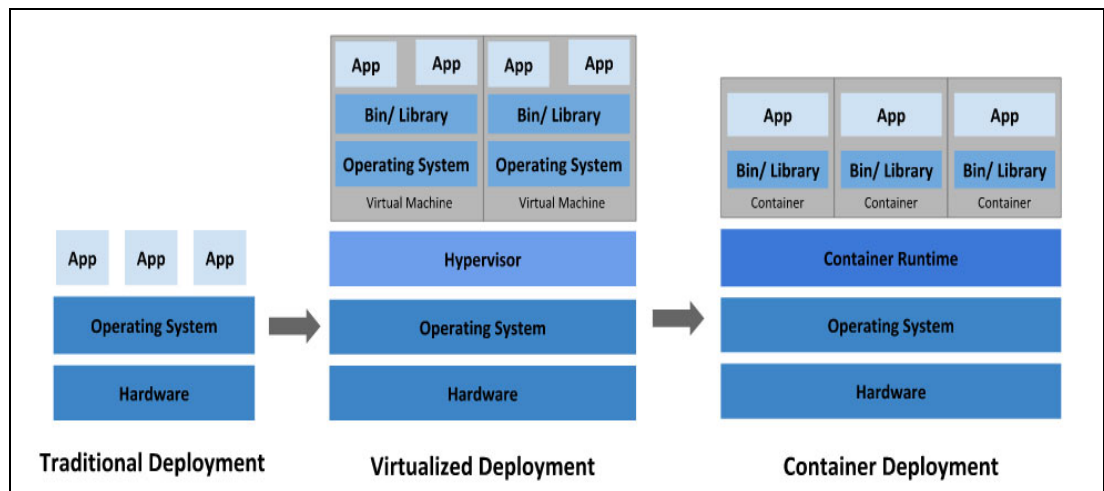


Figure 1-1 Application deployment models

1.1.1 Benefits of containers

The primary advantage of containers, especially compared to a VM, is providing a level of abstraction that makes them lightweight and portable:

- ▶ Lightweight containers share the machine operating system kernel, which eliminates the need for a full operating system instance per application and makes container files small and easy on resources. Their smaller size, especially compared to VMs, means that they can spin up quickly and better support cloud-native applications that scale horizontally.
- ▶ Portable and platform independent: Containers carry all their dependencies with them; that is, software can be written once and then, run without needing to be reconfigured across laptops, cloud, and on-premises computing environments.
- ▶ Supports modern development and architecture: Because of a combination of their deployment portability and consistency across platforms and their small size, containers are an ideal fit for modern development and application patterns (such as DevOps, serverless, and microservices) that are built as regular code deployments in small increments.
- ▶ Improves utilization: Like VMs before them, containers enable developers and operators to improve CPU and memory usage of physical machines. Where containers go even further is that application components can be deployed and scaled more granularly because containers also enable microservice architectures. This ability is an attractive alternative to having to scale up an entire monolithic application because a single component is struggling with load.

1.1.2 Container concepts

A standardized approach is used to developing and deploying containerized applications, as shown in Figure 1-2.

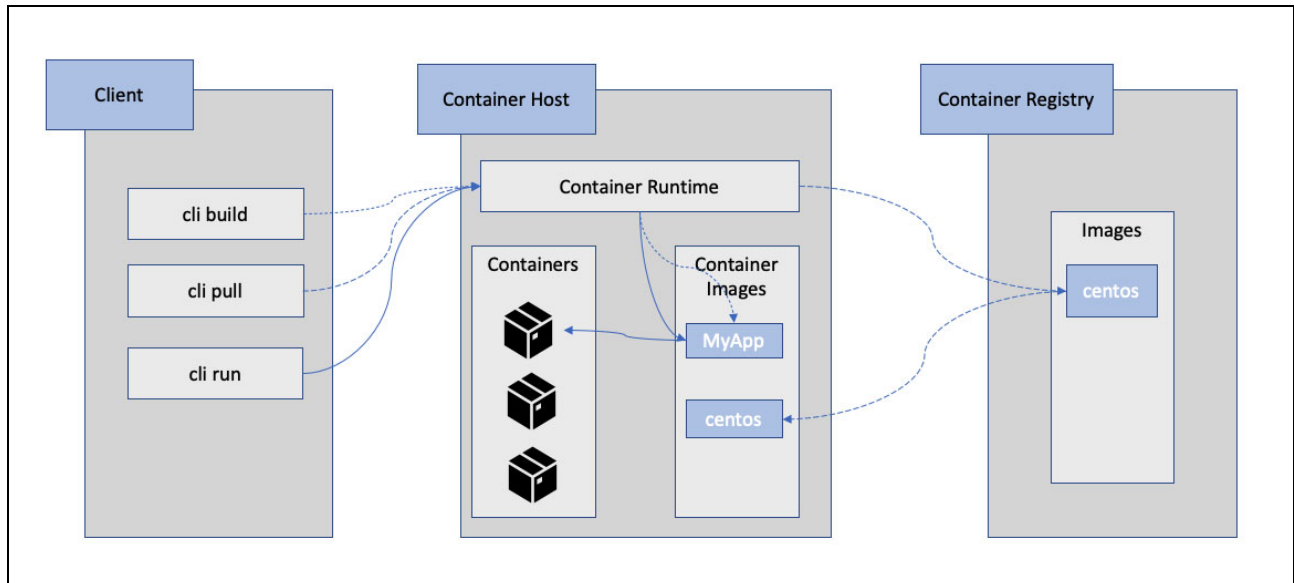


Figure 1-2 Container concepts

Containers feature the following main components:

- ▶ Container command-line interface (CLI)
Provides an interface to interact with the container run time. The CLI is used to interact with the container run time to build new images or move images between the host and the container registry.
- ▶ Container run time
This run time manages the lifecycle of a container, along with how a container runs code. Several container run times are available, including Docker and Podman.
- ▶ Container image
A container image is a read-only snapshot of a container that is stored in public or private repository or on the host for running by the container run time. Images also are used as a template for building containers.
- ▶ Container
This unit of software packages code with its required dependencies to run in an isolated, controlled environment.
- ▶ Container registry
Container images are stored in registries. Several large public registries are used in which images for “off the shelf” applications are published. Most enterprises that write in-house applications also include private registries that are hosted by a cloud provider (IBM Cloud®) or run internally.

1.2 Orchestrating containers

Although it is possible to run containers on an individual machine, the use of containers at scale is practical only with a container orchestration platform. As containers proliferated (today, an organization might have hundreds or thousands of them), operations teams are needed to schedule and automate container deployment, networking, scalability, and availability.

To support these orchestration needs, several platforms exist. However, in this IBM Redpaper publication, we focus on Red Hat OpenShift Container Platform, which is based on Kubernetes. This orchestration platform is available on most cloud providers (IBM Cloud, Microsoft Azure, Amazon Web Services (AWS), and others) and also is available for installation on-premises.

1.2.1 Red Hat OpenShift introduction

This section is not meant to be a guide to the Red Hat OpenShift Container Platform. Instead, the intention is to introduce the key components of the Red Hat OpenShift platform that are relevant to backup and restore by using the IBM Spectrum Protect tools.

For more information about available documentation, see this Red Hat OpenShift Documentation [web page](#).

In 1.1.2, “Container concepts” on page 3, we reviewed the concept of containers and the container run time that runs the containers on a single host. Red Hat OpenShift extends these capabilities by using a cluster, as shown in Figure 1-3 on page 5.

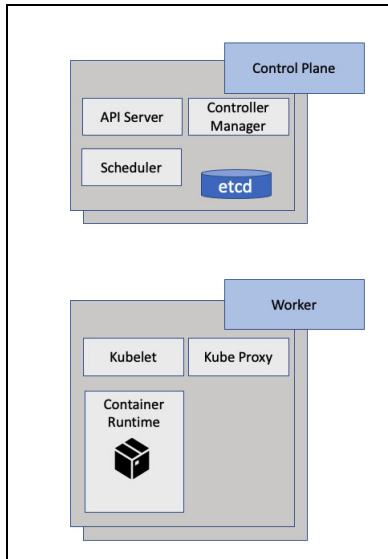


Figure 1-3 Red Hat OpenShift cluster topology

A Red Hat OpenShift cluster consists of the following key components:

- ▶ Control Plane

Consists of one or more nodes that host the control components for the cluster.

- ▶ Apiserver

The entry point to the cluster. The central point for all users, automation, and Kubernetes resources to interact with the cluster. It is responsible for storing the application state in the etcd persistent storage for the cluster.

- ▶ etcd

Highly available data store of the configuration of the cluster. Contains the details of the cluster and the applications that are hosted on the cluster.

- ▶ Controller manager

A daemon that manages the key control loops of the platform that ensure the state of the cluster matches the state as defined in the etcd database.

- ▶ Scheduler

Controls the placement of application pods on the worker nodes.

- ▶ Worker

A set of one or more nodes where the workloads on the cluster run.

- ▶ Kubelet

Primary node agent of the node that communicates with the apiserver.

- ▶ kube proxy

The network proxy for the cluster.

- ▶ Container run time

The run time that runs the containers, as with the single host solution. It can be any container run time that supports the Open Container Initiative (OCI) specification.

Red Hat OpenShift provides the capability to deploy an application as a set of containers. The description of the containers to run and the configuration for the application are packaged as a deployment. A high-level view of a deployment is shown Figure 1-4.

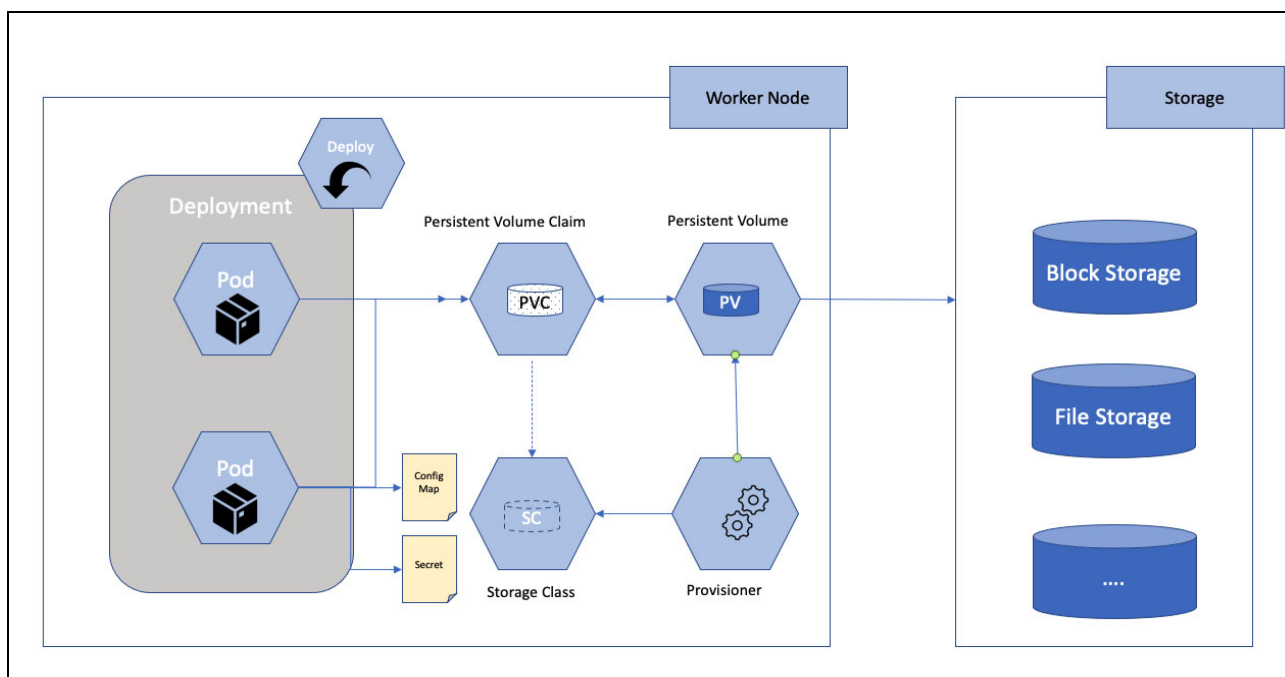


Figure 1-4 A deployment

A deployment is a Red Hat OpenShift artifact that defines the pods that are to be deployed. Deployments include the following key components:

- Namespaces

The container management system provides several constructs to make it easier to manage larger-scale deployments where multiple teams share an instance. Namespaces provide a way to create separate islands; for example, to keep two teams from seeing each other's containers. Deployments are deployed to namespaces.

- Pods

A pod is the unit of work that the scheduler schedules onto the worker nodes. Pods contain one or more containers, and describe the persistent volume claims (PVCs) that the containers need for storage.

- PVC

A request to a storage system for a place to store data. PVCs can be static or dynamic. With dynamic provisioning, a developer can deploy an application with a PVC, and a volume is automatically created and assigned. A PVC defines the type of storage class that it needs.

- Storage classes

Storage classes define the characteristics of a persistent volume. They indicate which storage provider must provision the volume and the specifics of the volume that is required (retention, priority, and expansible, and so on).

- Persistent volumes

When containers must store data that exists through restarts, they use persistent volumes. These storage volumes are available until deleted.

- Provisioner

Each persistent volume is created by a provisioner that uses a plug-in to interface with different types of back-end storage. The latest specification for provisioners is the Container Storage Interface (CSI). CSI is an abstraction layer for underlying storage technologies.

Storage vendors can write a CSI driver and then, Kubernetes dynamically provisions persistent volumes, as needed. CSI also supports other storage functions, such as snapshots, although not all CSI enabled storage includes support for the snapshot functions.

The CSI standard replaces an established approach that is known as *in-tree storage plug-ins* that Kubernetes is deprecating.

- Configuration maps and secrets

Extra data that is used to fully describe the deployment. Configuration maps provide a way to pass configuration data to a running container without hardcoding that information in the container image that is stored in the registry or the YAML file definition.

Secrets provide similar functions for sensitive information that requires extra privacy and security.

1.3 Protecting a Red Hat OpenShift environment

When implementing a solution to protect Red Hat OpenShift containers, it is important to understand the components of the environment and their specific protection requirements.

Some components do not need to be protected; for example, because they can be easily reinstalled, re-created from images, or are part of the underlying infrastructure (for example, physical hosts, storage systems, and network devices).

Other entities might be important to be protected because they contain unique customer-specific information, such as passwords, configuration information, or application data that was created by users.

1.3.1 Protected items

Figure 1-5 shows a basic Red Hat OpenShift environment.

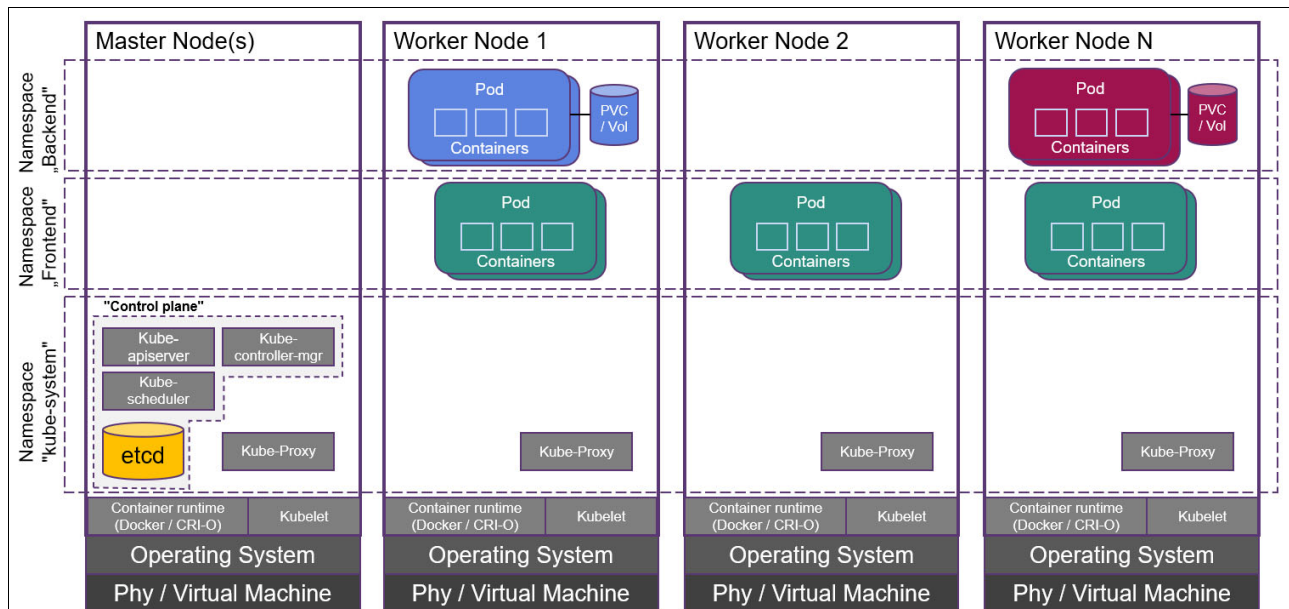


Figure 1-5 Components to be protected in a Red Hat OpenShift cluster

The disabled items in gray that are shown in Figure 1-5 do not need to be backed up because they can be rebuilt if a failure occurs. If for example, a worker-node fails, the Red Hat OpenShift-Scheduler relocates application containers to other nodes in the cluster. The infrastructure administrators can deploy a new worker-node to reestablish the wanted level of redundancy.

The colored items that are shown in Figure 1-5 indicate entities that must be protected because they contain specific user data that cannot be rebuilt from other sources.

The following items include the cluster metadata (which is stored in the etcd database) and the application data from stateful pods (which is stored in persistent volumes):

- Cluster metadata

Red Hat OpenShift stores most of the cluster's operational state in the etcd cluster database on the control nodes. A backup of etcd is always considered a best practice for production clusters. However, etcd is not the only consideration. It is also essential to have backups of some of the static assets that are created during deployment or updated later; for example, PKI assets or certificates that are used by the Red Hat OpenShift apiserver and any secrets or configuration maps.

- Application data

In addition to all of the metadata that is necessary to reconstitute Red Hat OpenShift, recovering stateful pods is useless, unless the persistent data that is associated with those pods also is recovered. This process involves protecting the stateful data on persistent volumes. Containerized applications also need crash or application-consistent backups, depending on the requirements of the application.

1.3.2 IBM Spectrum Protect Plus Container Backup support

IBM Spectrum Protect Plus offers support to protect Red Hat OpenShift environments.

IBM Spectrum Protect Plus uses the open source tool Velero to facilitate Kubernetes metadata (etcd) backups. These backups can be stored locally in the cluster (for example, for fast recovery of older states) or copied to external object or disk storage (for example, for disaster protection or to enable cross-cluster recovery).

Also, IBM Spectrum Protect Plus implements mechanisms to protect user data that is stored in persistent volumes. This backup data can be stored locally in the form of storage snapshots or copied to external object or disk storage.

For more information about the general IBM Spectrum Protect Plus architecture, see in Chapter 2, “IBM Spectrum Protect Plus architecture” on page 11.

For more information about container backup support, such as components, prerequisites, supported Red Hat OpenShift environments, and storage providers, see Chapter 4, “Container Backup Support” on page 61.



IBM Spectrum Protect Plus architecture

This chapter provides a general architectural overview of the IBM Spectrum Protect Plus solution. Readers that are familiar with the product can skip this chapter.

The information that is included in this chapter is intentionally kept brief. For more information about this topic, see [IBM Spectrum Protect Plus documentation](#) or read the IBM Redbooks publication *IBM Spectrum Protect Plus Practical Guidance for Deployment, Configuration, and Usage*, REDP-5532.

This chapter includes the following topics:

- ▶ 2.1, “Overview” on page 12
- ▶ 2.2, “Architecture” on page 12
- ▶ 2.3, “Components and terminology” on page 15

2.1 Overview

IBM Spectrum Protect Plus is used for backing up virtualized systems, databases, file systems, cloud-managed applications, and containerized applications. Incremental backups are realized by using snapshot technology to provide rapid backup, data reuse, recovery, and (self-service) data management.

Data availability retention compliance is provided through automated service-level agreements (SLAs).

IBM Spectrum Protect Plus is a zero touch data protection solution; that is, no manual agent roll-out or installation is needed.

2.2 Architecture

Figure 2-1 shows an overview of the IBM Spectrum Protect Plus architecture with key components and the data transfer options from the data sources.

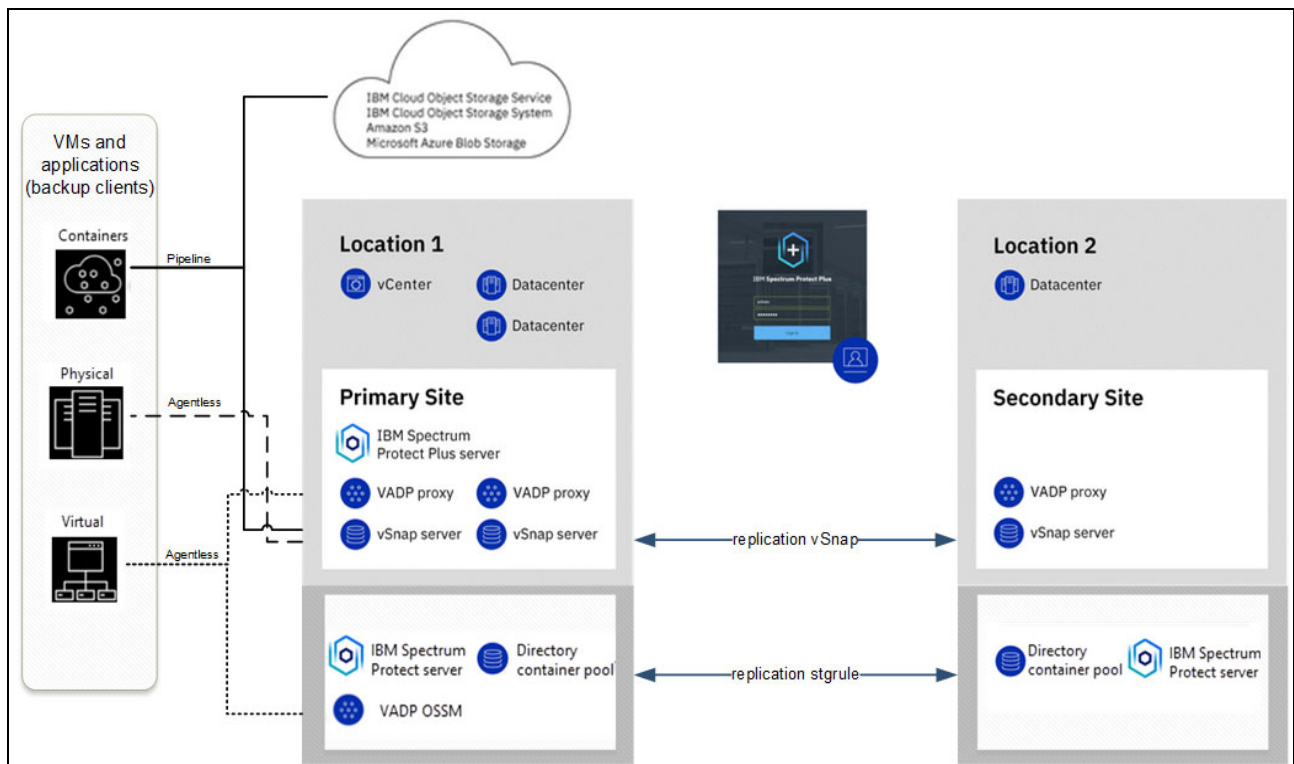


Figure 2-1 IBM Spectrum Protect architecture overview

Repositories can be a vSnap storage server or an IBM Spectrum Protect directory container pool, or S3 cloud storage. The data source type defines which options are available and is shown in Figure 2-1 with the lines from the “VMs and application” box.

The IBM Spectrum Protect Plus server is the control plane of all backup data, regardless of where it is stored.

A common deployment approach to copy snapshots to secondary backup storage is shown in Figure 2-2 (copy to IBM Cloud Object Storage) and Figure 2-3 (copy to a repository IBM Spectrum Protect server).

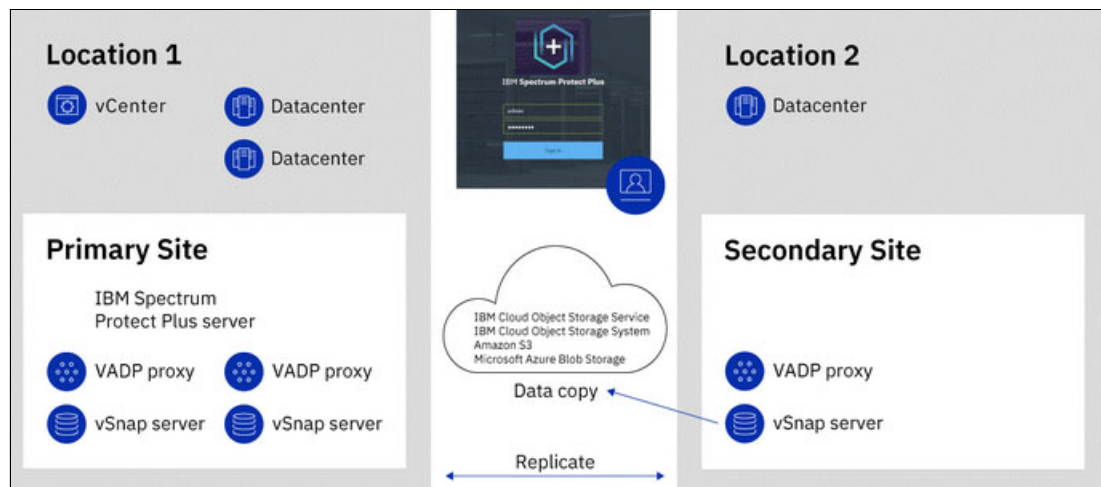


Figure 2-2 Deployment across two geographical locations with copy to cloud storage

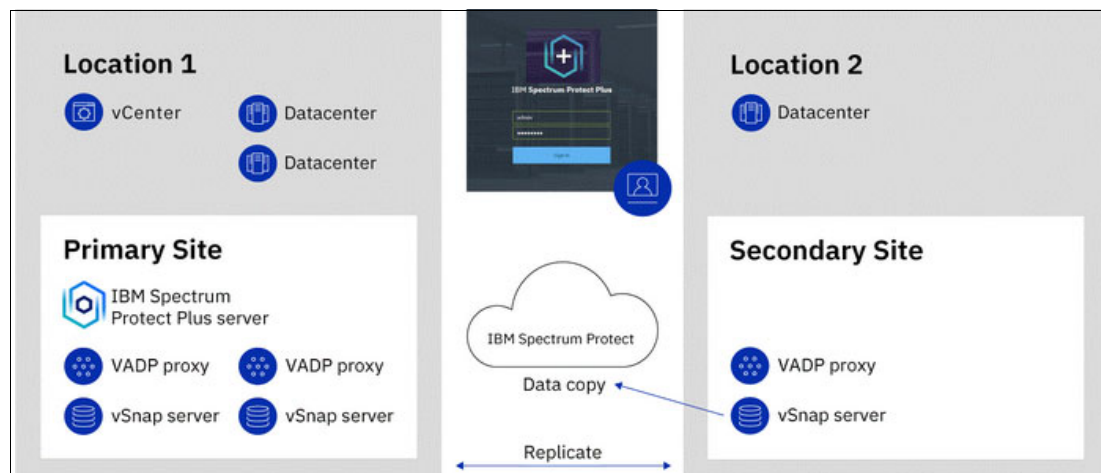


Figure 2-3 Deployment across two geographical locations with copy to IBM Spectrum Protect

Backup of VMware workloads directly to an IBM Spectrum Protect server was introduced with IBM Spectrum Protect server 8.1.16 and IBM Spectrum Protect Plus 10.1.12. This configuration does not require a vSnap server. Instead, an Open Snap Store Manager (OSSM) agent is used to provide an extra layer of data protection by backing up the data to IBM Spectrum Protect directory-container storage pools.

Also, by using OSSM, you can replicate the data from one IBM Spectrum Protect server to another IBM Spectrum Protect server.

The OSSM components are shown in Figure 2-4.

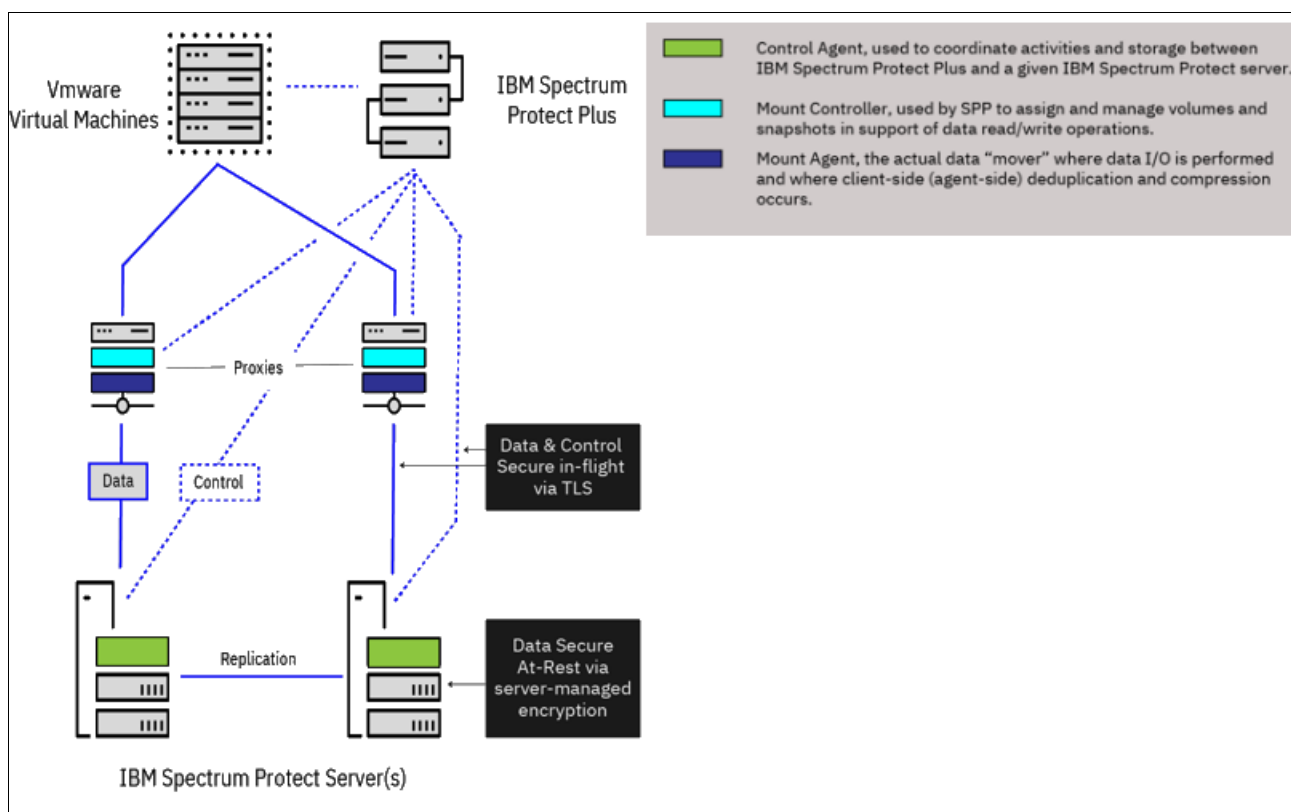


Figure 2-4 Configuration with Open Snap Store Manager

Three components are introduced. On the VADP proxies, the Mount Controller and the Mount Agent are installed. On the IBM Spectrum Protect server, the Control Agent can be selected as part of the server installation. These components make data transfer transparent to IBM Spectrum Protect Plus and IBM Spectrum Protect.

Note: As of this writing, IBM Spectrum Protect server that also serves as OSSM control agent is supported on only Linux x86_64 that is protecting VMware workloads.

Container workloads, such as Kubernetes and Red Hat OpenShift clusters backup (snapshot) copies, can be stored directly into an S3 object store. This feature makes the solution independent of a noncontainerized infrastructure.

Moving the objects to or from the object store is based on the restic tool that is part of the backup as a service (BaaS) client installation. The restic tool also handles the expiration of data in the object store. The following object stores are supported:

- ▶ Amazon S3
- ▶ IBM Cloud Object Storage
- ▶ Microsoft Azure Blob storage
- ▶ S3 compatible storage

For more information about other validated vendors, see IBM Support [technote 1087149](#).

For more information about the BaaS installation, see Chapter 5, “Implementing Container Backup Support” on page 85.

The key components and terms are described next.

2.3 Components and terminology

This section includes a short description of some of the most important IBM Spectrum Protect Plus components and terms in the context of this publication. For more information, especially about sizing aspects, see the IBM Spectrum Protect Plus Blueprints that are available at [this web page](#).

2.3.1 IBM Spectrum Protect Plus server

The IBM Spectrum Protect Plus server is the component that manages and orchestrates the entire system. It also is the “brain” that provides the web interface portal with role-based access. It is used for configuring and operating the backup and restore solution, and performing centralized scheduling of activities.

Although the IBM Spectrum Protect Plus server orchestrates the entire system, it only “triggers” actions; that is, it is *not* part of the backup processing data path. Backup data flow is handled between other components directly to prevent a bottleneck at the IBM Spectrum Protect Plus server. The IBM Spectrum Protect Plus server is responsible for scheduling, starting, and cataloging the backups.

Figure 2-1 on page 12 shows the IBM Spectrum Protect Plus server in the middle, which is deployed as a virtual machine (VM). The IBM Spectrum Protect Plus server communicates with the backup clients by using standardized API calls. This communication is used to manage the inventory and backup processes.

Note: One IBM Spectrum Protect Plus server instance is used per deployment.

2.3.2 vSnap Backup Storage server

The vSnap Backup Storage server (vSnap server) is the component that is responsible for storing and processing the backup data that is received from production systems (backup clients).

Backups that are taken as block-level incremental forever backups from VMs, databases, and applications are stored according to the backup policies (SLA) as a read/write snapshot in the vSnap server. A vSnap server can run on virtual or physical systems, and can be replicated to a second vSnap server as a disaster recovery (DR) option.

If required, data can be copied to IBM Spectrum Protect, onto multiple different storage options (such as Object Storage), or onto a dedicated Object Storage (such as IBM Cloud Object Storage) to increase the protection level.

The vSnap server manages the backup traffic and workload. The backup metadata is sent to the IBM Spectrum Protect Plus server. To achieve local backup data placement by using SAN or LAN networks and to scale out, multiple vSnap servers often are used within an IBM Spectrum Protect Plus solution. For efficient performance, correct sizing of the vSnap server is crucial.

For more information, see the IBM Spectrum Protect Plus Blueprints that are available at [this web page](#).

Note: One or more vSnap servers are used per deployment, usually at least one per site. Any vSnap server can belong to a single site only.

2.3.3 VADP proxy server with vSnap backup storage server

To protect VMware VMs, a VADP proxy server must be used. This proxy server or data mover is responsible for moving data from vSphere data stores (through the proxy server) to the vSnap server. Each site, with VMware VMs that are required to be backed up, needs at least one VADP proxy server to be deployed. Based on sizing, more VADP proxies can be added.

Note: A common practice is to deploy the VADP proxy together with the vSnap server. In this case, IBM Spectrum Protect optimizes data movement by eliminating extra transfer overhead (NFS mount over LAN), and handling this traffic inside the same (virtual or physical) machine.

2.3.4 Open Snap Store Manager storage server

As an alternative to vSnap servers, backup data can be ingested directly into an IBM Spectrum Protect server directory container pool. At the time of this writing, VMware workloads are supported. The OSSM component is installed on the IBM Spectrum Protect server that is running Linux x86_64 and on at least one VADP proxy node.

A new concept of a central control agent that is installed on the IBM Spectrum Protect server controls the backup and restore data flow. Data can be replicated to a secondary IBM Spectrum Protect target server.

For more information about how to implement OSSM, see IBM Support [technote 6612979](#).

2.3.5 VADP proxy with OSSM

The proxy agent is controlled by the central control agent to start the backup and restore process.

The mount agent handles the file operations for backup and restore operations. A local MariaDB holds temporary metadata during the backup process. To reduce network traffic data, deduplication queries ensure that extents are not resent to the IBM Spectrum Protect server.

The VADP proxy OSSM components cannot be configured on the same operating system as the IBM Spectrum Protect server.

2.3.6 Backup clients

Systems that are protected by IBM Spectrum Protect Plus often are referred to as *backup clients*. Backup clients are in one of the following categories:

- ▶ Virtualized: VMware, Hyper-V, and Amazon EC2
- ▶ Databases: IBM Db2®, Oracle, MongoDB, Microsoft SQL, Microsoft Exchange, and SAP HANA

- Container: Kubernetes and Red Hat OpenShift Container Platform
- File Systems: Microsoft Windows

This publication focuses mainly on the use of IBM Spectrum Protect Plus to protect containerized application data that is deployed in a Red Hat OpenShift environment.

For more information about IBM Spectrum Protect Plus Container Backup Support (BaaS), see Chapter 4, “Container Backup Support” on page 61.

For more information about backup and restore requirements for the different supported systems, see this [IBM Support web page](#).

IBM Spectrum Protect Plus Online Services

IBM Spectrum Protect Online Services replaced the former way of protecting M365 workloads. Today, the service also supports sales force data protection. The service runs in the Microsoft Azure cloud and is operated by IBM.

Although no onsite infrastructure is required, it is possible to bring your own storage (BYOS) instead of what is provided by IBM. All backup and restore is done by using the [online portal](#). As of this writing, data centers are in the US, Canada, Germany, and the UK.

2.3.7 Service-level agreement backup policies

SLA backup policies (also known as *policies*) define the backup jobs parameters. This definition includes information about backup schedule, retention time, target (site), and extra protection (replication or extra copy). You assign backup clients to one or multiple SLAs to create corresponding backup jobs. Although an SLA cannot be renamed, the parameters can be modified for future backups.

Figure 2-5 shows the different options of target storage, depending on the data source.

Policy Type

Select policy type

i You will not be able to change the policy type after the policy has been created

Category
All protected

Tiered vSnap

This policy includes these functions:

- Back up to vSnap server
- Back up logs
- Replicate data
- Copy or archive data

The following workloads can use this policy:

- VMware
- Hyper-V
- Db2
- Exchange
- Microsoft 365
- MongoDB
- Oracle
- SQL
- SAP HANA
- Windows File System

OSSM

This policy includes these functions:

- Back up to OSSM
- Replicate data

The following workloads can use this policy:

- VMware

Tiered snapshot

This policy includes these functions:

- Back up to vSnap server
- Manage snapshot protection
- Back up logs
- Replicate data
- Copy or archive data

The following workloads can use this policy:

- Kubernetes
- OpenShift

Figure 2-5 Policy types available depending on workload type

2.3.8 Site

The term *site* is used as a policy construct to manage data placement in the IBM Spectrum Protect Plus environment. A site can have *physical* or *logical* meaning, such as a physical data center, or a logical organizational structure. IBM Spectrum Protect Plus components are assigned to sites to localize and optimize data paths.

Note: A default deployment includes two sites: primary and secondary. You can change names, add sites, and define a throttle rate to limit maximum throughput between sites that are used for replication or copy operations.

2.3.9 Extra copies

You can increase the level of protection for your backup data by creating copies. These copies can be stored in one of the following locations:

- ▶ On secondary backup storage (that is object storage)
- ▶ On multiple vSnap servers (assigning multiple SLA backup policies)
- ▶ As a replication between two vSnap servers (for more information, see “Replication”)
- ▶ As a tape or cloud archive storage (object archive cloud storage, physical tape or VTL)

Note: All types of extra copies establish new extra copies. None of those copies are moving or migrating any data from one storage tier to another.

2.3.10 Replication

Replication is a specific type of extra copy and refers to the replication of the backup storage data from one backup storage server to another backup storage server. The replication can be a vSnap-to-vSnap server or from a source-to-target IBM Spectrum Protect server by way of the OSSM agent.

The partner relationship configuration between sites dictates the data placement of the secondary copy.

Common use cases for replication are to protect your backup data against the following issues:

- ▶ Complete loss of a primary vSnap or IBM Spectrum Protect server
- ▶ Loss of an entire data center (site)
- ▶ Limited available space in source location; defining a longer retention time for the target

Note: Replication is possible between vSnap servers that are assigned to different sites only. Replication is the only way to recover a vSnap server.



Installing IBM Spectrum Protect Plus as a containerized application

This chapter discusses the deployment of IBM Spectrum Protect Plus as a set of containers within a Red Hat OpenShift Container Platform. Do not confuse this specific installation with the ability of protecting containerized application data, which also is supported by the traditional IBM Spectrum Protect Plus deployment.

For more information about this traditional deployment, see Chapter 5, “Implementing Container Backup Support” on page 85.

We assume that you have a basic knowledge of containers and the Red Hat OpenShift Container platform. You must understand the meaning of terms such as *Pods*, *storage class*, *persistent volume claim (PVC)*, and *operators*. For more information about these terms and concepts, see Chapter 1, “Introducing containers” on page 1.

It also is important to understand that when deploying the IBM Spectrum Protect Plus Server as a set of containers, the vSnap server is not part of this containerized deployment. Without a vSnap server, you can still perform backups of configuration resources and PVCs with IBM Spectrum Protect Plus server; however, these backups remain as local snapshots that are inside the Red Hat OpenShift Container Platform, without copy outside of the cluster.

To secure an extra copy outside of the containerized environment, starting IBM Spectrum Protect Plus version 10.1.9, you can store backup on a vSnap server (which must be installed and configured by using the .ova image for VMware Virtual environment) or store your backup directly to an object storage.

For more information, see this [IBM Documentation web page](#).

Note: For the installation that is described in this chapter, the Red Hat OpenShift environment must access the IBM registry, which presumes a working connection with the internet.

This chapter describes the installation process, starting from the prerequisites up to the IBM Spectrum Protect Plus instance running, and includes the following topics:

- ▶ 3.1, “Understanding the environment and requirements” on page 22
- ▶ 3.2, “Preparing Red Hat OpenShift Container Platform” on page 27
- ▶ 3.3, “Deploying IBM Spectrum Protect Plus by using an operator” on page 38
- ▶ 3.4, “Configuring an IBM Spectrum Protect Plus Server catalog backup” on page 43
- ▶ 3.6, “IBM Spectrum Protect Plus deployment troubleshooting” on page 52

3.1 Understanding the environment and requirements

Deploying IBM Spectrum Protect Plus server as a containerized application is different than deploying IBM Spectrum Protect Plus on VMware to back up containerized applications data. These installations involve different processes and result in distinct outcome.

The goal in the first case is to deploy the IBM Spectrum Protect Plus server as a set of containers whereas in the latter, various components are deployed to perform backups of containerized applications. For more information about how to protect containerized application data, see Chapter 5, “Implementing Container Backup Support” on page 85.

Before reviewing how to deploy IBM Spectrum Protect Plus Server as a containerized application, you must understand some key concepts to understand and complete some preparation work on the Red Hat OpenShift Container Platform.

Container deployment is made in a way that the IBM Spectrum Protect Plus server components are distributed across multiple containers, which are grouped as pods and are referred as a *deployment*. This deployment is managed by an operator that is made available by IBM.

IBM Spectrum Protect Plus server uses CPU, memory, disk space, and network resources, and is accessible through an HTTP interface. These concepts are defined and named next in a Red Hat OpenShift Container Platform environment.

In this section, we also present a typical IBM Spectrum Protect Plus Server deployment and the following resources that are created after the setup process is complete:

- ▶ Containers and pods (compute)
- ▶ Persistent volumes (PV) and PVC, which define persistent storage
- ▶ Services (network within the containers to make the application work)
- ▶ Routes (access the containerized application from outside the cluster)

These IBM Spectrum Protect Plus server resources are created and managed by the IBM Spectrum Protect Plus Operator. 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.

The IBM Spectrum Protect Plus Operator also can be used after the setup to change some of the IBM Spectrum Protect Plus Server deployment settings, which then are automatically applied by that operator to the IBM Spectrum Protect Plus related objects.

3.1.1 Containers and pods

IBM Spectrum Protect Plus Server is deployed as a set of 16 pods: Twelve pods for the server's core components and four pods for the managing components.

Each of these pods is composed of one or multiple containers. In an IBM Spectrum Protect Plus deployment (except the spp-operator pod), all pods are composed of a single container.

The following core component pods are defined:

- ▶ spp-awsebs-xxx
- ▶ spp-awsec2-xxx
- ▶ spp-plugins-mongo-xxx
- ▶ spp-plugins-redis-xxx
- ▶ sppdbmongo-xxx
- ▶ sppdbmongo2-xxx
- ▶ sppdbpostgres-xxx
- ▶ sppkc-xxx
- ▶ sppnodejs-xxx
- ▶ sppui-xxx
- ▶ sppvadm-xxx
- ▶ sppvirgo-xxx

Where the xxx is an auto-generated string that is unique to your deployment.

The following managing pods are defined:

- ▶ spp-operator-xxx is the pod where the operator is running. It is used to automatically deploy IBM Spectrum Protect Plus Server components.
- ▶ spp-manager-xxx is used to update the IBM Spectrum Protect Plus instance from IBM Spectrum Protect Plus (the `https://SPPUI:8090` interface).
- ▶ spp-proxy-xxx is used for internal communication between the virgo container and other containers.
- ▶ spp-ingress-xxx is responsible for routing requests to the spp-proxy container, which in turn routes the requests internally between the various containers that are hosting different components of the IBM Spectrum Protect Plus server.

Figure 3-1 shows an IBM Spectrum Protect Plus deployment in a Red Hat OpenShift Container Platform.

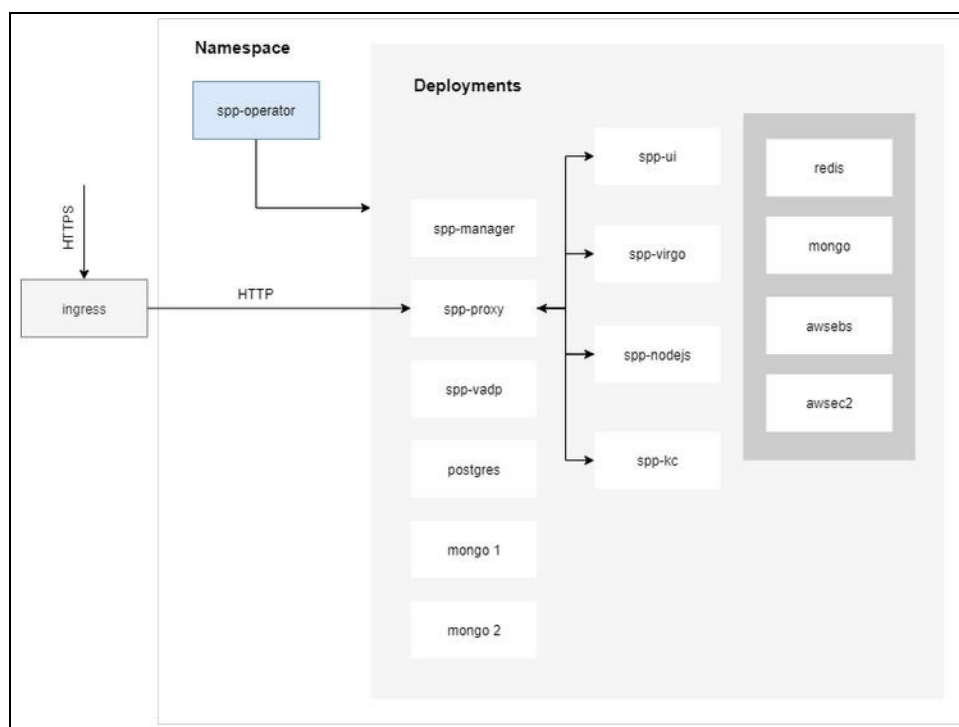


Figure 3-1 IBM Spectrum Protect Plus deployment as a container

Note: If you run the `oc get pods` command in your project or namespace, you notice a pod with a name starting with `spp-operator*`. This container is deployed for the IBM Spectrum Protect Plus Operator. For more information, see 3.2, “Preparing Red Hat OpenShift Container Platform” on page 27.

3.1.2 Persistent storage

Eight of the 16 pods use persistent storage. For more information about the required space, see this [IBM Documentation web page](#).

For IBM Spectrum Protect Plus Server version 10.1.12, the persistent volumes require approximately 400 GiB and are distributed as listed in Table 3-1.

Table 3-1 Persistent storage size

Persistent volume	Size (GiB)	Mount path	Permissions	Container that access the PVC
Virgo logs	10	/data/log	drwxrwsr-x	sppvirgo
plug-in logs	10	/data/platform/log	drwxrwsr-x	awsev2 and awsebs
MongoDB	100	/var/lib/mongodb/data	drwxrwsr-x	sppdbmongo
MongoDB catalog	100	/var/lib/mongodb/data	drwxrwsr-x	sppdbmongo2
Postgres	10	/var/lig/pgsql/data	drwxrwsr-x	sppdbpostgres

Persistent volume	Size (GiB)	Mount path	Permissions	Container that access the PVC
Apache Lucene	150	/data/lucene	drwxrwsr-x	sppvirgo
Nodejs logs	10	/data/log/ode-cdm-service	drwxrwsr-x	sppnodejs
VMware vStorage API for Data Protection proxy (VADP proxy) logs	10	/data/log/vmdkbackupproxy	drwxrwsr-x	sppvadb

Note: All PVC access is set to read_write_once (RWO), which means that the volume can be mounted as read/write by a single node.

Example 3-1 shows, from the Red Hat OpenShift console, the persistent storage that is created as part of the spp-redb IBM Spectrum Protect Plus server deployment. In this example, the persistent volumes are used through a PVC that belongs to the project spp-redb, which is the namespace that is used for our IBM Spectrum Protect Plus server deployment.

Example 3-1 Listing the persistent volumes that are associated to an IBM Spectrum Protect Plus deployment

```
[root@ocp-helper ~]# oc get pv |grep -E "NAME|spp"
NAME                                CAPACITY  ACCESS MODES  RECLAIM POLICY  STATUS  CLAIM
pvc-11dbc09d-30b0-4856-b5fe-12ad5e3fd135  10Gi      RWO           Delete          Bound   spp-redb/sppnodejs-spp-redb-claim
pvc-23484be9-5279-4c89-bc7d-9c33c2e3573a  150Gi     RWO           Delete          Bound   spp-redb/virgo-lucene-spp-redb-claim
pvc-3c9a3f79-a328-45e7-851a-b011a3ddc9f5  10Gi      RWO           Delete          Bound   spp-redb/sppvadb-spp-redb-claim
pvc-429384d8-8dc6-44f7-9b2b-5800a04c94cb  100Gi     RWO           Delete          Bound   spp-redb/sppdbmongo-spp-redb-claim
pvc-5bc7ad53-5e24-4835-9b45-ec8b2f0c6c86  20Gi      RWO           Delete          Bound   spp-redb/spp-log-spp-redb-claim
pvc-b5fe690f-d83b-4313-b583-07bdf9516ff  10Gi      RWO           Delete          Bound   spp-redb/sppdbpostgres-spp-redb-claim
pvc-d2aa0447-d1a5-4962-8bb3-848f229e44b7  100Gi     RWO           Delete          Bound   spp-redb/sppdbmongo2-spp-redb-claim
```

Example 3-2 shows the association between the containers and their PVCs.

Example 3-2 PVC that is used for IBM Spectrum Protect Plus server containers

```
[root@ocp-helper ~]# oc project spp-redb
Now using project "spp-redb" on server "https://api.ocp4.isv.escc.lab:6443".
[root@ocp-helper ~]# oc get pods -n spp-redb | awk '/spp/ { print $1}' | while read line; do oc describe pods $line
-n spp-redb | grep -E "^Name:|^Node:|^ClaimName"; echo ; done
Name:      spp-awsebs-6b7f7ffc6f-f9z2b
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230

Name:      spp-awsec2-9489c5d49-9zp55
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230

Name:      spp-ingressproxy-54d7dcc67c-6m21r
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230

Name:      spp-manager-586c57559b-84pd7
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230

Name:      spp-operator-664f985744-z9qmh
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230

Name:      spp-plugins-mongo-66f478dbd7-6d6vq
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230

Name:      spp-plugins-redis-6cbb55b57-w6cg1
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230
```

```

Name:      spp-proxy-755d6bf5c9-w5zgd
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230

Name:      sppdbmongo-5d94fdf7f9-kbdbx
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230
ClaimName: sppdbmongo-spp-redb-claim

Name:      sppdbmongo2-6695ffcf4c-zcdp9
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230
ClaimName: sppdbmongo2-spp-redb-claim

Name:      sppdbpostgres-64b9d68989-tprtz
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230
ClaimName: sppdbpostgres-spp-redb-claim

Name:      sppkc-77698479dd-5nfcg
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230

Name:      sppnodejs-69bbd6b7f6-trmfv
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230
ClaimName: sppnodejs-spp-redb-claim

Name:      sppui-58c76fd85b-gthjj
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230

Name:      sppvadb-699f4954b6-hznhg
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230
ClaimName: sppvadb-spp-redb-claim

Name:      sppvirgo-7459bd485f-7qx5g
Node:      ocp-worker-04.ocp4.isv.escc.lab/10.0.240.230
ClaimName: virgo-lucene-spp-redb-claim
ClaimName: spp-log-spp-redb-claim

```

3.1.3 Services

A *service* is an abstract way to make available an application that is running on a set of pods as a network service.

When you deploy an IBM Spectrum Protect Plus Server, Red Hat OpenShift creates two services for the inter-pods communication and possibly external communication when a route is built on top of that service.

As shown in Example 3-3, a service was created for the operator, which is fully integrated in this deployment because the operator is used to instantiate or upgrade the IBM Spectrum Protect Plus server containers.

Example 3-3 List of services that were created for an IBM Spectrum Protect Plus server redeployment

```

[root@ocp-helper ~]# oc get services | grep spp
spp                ClusterIP  172.30.29.120  <none>      8082/TCP,5672/TCP,5671/TCP  7h10m
spp-operator-metrics ClusterIP  172.30.85.140  <none>      8383/TCP,8686/TCP          24h
sppingressproxy    ClusterIP  172.30.1.84    <none>      80/TCP          7h9m
sppmanager          ClusterIP  172.30.78.219  <none>      80/TCP          7h10m
sppproxy            ClusterIP  172.30.106.193 <none>      443/TCP         7h9m

```

3.1.4 Routes

A Red Hat OpenShift *route* is a way to make available a service by giving it an externally reachable hostname, such as `sppname-spp.apps.cluster311.ocpstorage.icc`, as shown in Example 3-4 on page 27.

A defined route and the endpoints that are identified by its service can be used by a router to provide named connectivity. This feature enables external clients to reach the application; in this case, the IBM Spectrum Protect Plus server GUI through the `sppproxy` service.

Example 3-4 Listing the route that is created

```
[root@ocp-helper ~]# oc get routes
```

NAME	HOST/PORT	PATH	SERVICES	PORT
TERMINATION	WILDCARD			
spp-p9lwx	sppredb.apps.ocp4.isv.escc.lab	/	sppingressproxy	<all>
edge/Redirect	None			

3.2 Preparing Red Hat OpenShift Container Platform

Note: The Red Hat OpenShift Container Platform operations that are described in this section require Red Hat OpenShift Cluster administrator privileges.

This section reviews the Red Hat OpenShift Container Platform preparation steps that are required before the IBM Spectrum Protect Plus server is deployed as a set of containers.

The preparation process includes the following tasks:

- ▶ Review requirements.
- ▶ Get Entitlement key and create image pull secret.
- ▶ Register the IBM Container registry in Red Hat OpenShift Container Platform to access IBM Spectrum Protect Plus Operator.
- ▶ Create a project that is dedicated to IBM Spectrum Protect Plus deployment.

3.2.1 Requirements

At the time of this writing, IBM Spectrum Protect Plus 10.1.12 supports Red Hat OpenShift Container Platform versions 4.8 and beyond.

For more information about IBM Spectrum Protect Plus 10.1.12, requirements, see [IBM Documentation](#). The requirements and support matrix that are available at this IBM Documentation web page are changing rapidly; therefore, be sure you check the page often for the latest information.

In addition to the software requirements and product compatibilities, you must perform several preparation tasks in your environment, install the IBM Spectrum Protect Plus operator, and create an instance of the IBM Spectrum Protect Plus server to install and use IBM Spectrum Protect Plus on a Red Hat OpenShift cluster.

Complete the following steps in Red Hat OpenShift cluster before IBM Spectrum Protect Plus server's containers are deployed:

1. Log in to the IBM Container software library to get the entitlement key.
2. Create an image pull secret in the openshift-marketplace project/namespace to enable the Red Hat OpenShift cluster to authenticate with the IBM Entitled Registry.
3. Add the catalog source that contains the IBM Spectrum Protect Plus operator to the Red Hat OpenShift web console.
4. Create a project (namespace) in the Red Hat OpenShift web console, to where the IBM Spectrum Protect Plus server is deployed.

5. Create the image pull secret for IBM Spectrum Protect Plus in the project that was created in step 4 on page 27.
6. Install the IBM Spectrum Protect Plus operator.

The following sections describe these steps.

Steps 1 and 2: Preparing for the registry access

Complete the following steps:

1. Connect to the [IBM Container software library](#).
2. Log in with your IBMid and password, which is associated with the entitled software.
3. The Entitlement Key window opens. Click **Copy Key** to get the entitlement key (as shown in Figure 3-2) and save it in a secure place. This key is needed later in the setup process to create an image pull secret.

Note: If you do not have a key that is defined or you want to create one, click **Add New key** and a new key is generated automatically and presented to you on that window.

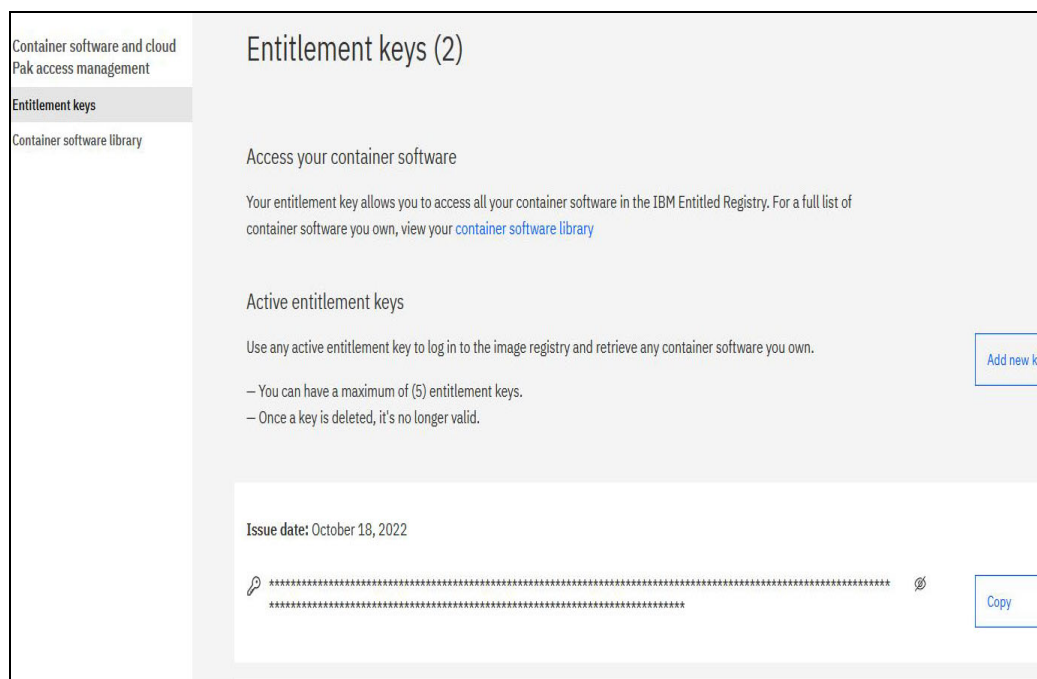


Figure 3-2 IBM Container Software registry key

The image pull secret provides the credentials for pulling Docker images from the IBM Entitled Registry. The image pull occurs at deployment time, and requires an internet connection from the Red Hat OpenShift environment.

4. After the key is downloaded from IBM Container software library, log in at the Red Hat OpenShift web console as cluster administrator and create the image pull secret.

To create an image pull secret, complete the following steps:

- a. Log on to the Red Hat OpenShift web console as the cluster administrator.
- b. In the navigation window, click **Workloads** → **Secrets**.

- c. Ensure that the project is created in the openshift-marketplace project by clicking **Project** → **openshift-marketplace**.
- d. On the Secrets window, click **Create** → **Image Pull Secret**.
- e. On the Create Image Pull Secret window, enter ibmspp-image-secret for the name of the secret.
- f. Create the image pull secret by entering the credentials for the IBM Entitled Registry:
 - i. In the Authentication Type list, click **Image Registry Credentials**.
 - ii. In the Registry Server Address field, enter the address for the IBM Entitled Registry: cp.icr.io/cp/sppserver.
 - iii. In the Username field, enter cp.
 - iv. In the Password field, enter the entitlement key that you obtained.
 - v. Click **Create**.

Figure 3-3 shows the Create Image Pull Secret window.

The screenshot displays the 'Create Image Pull Secret' interface in the Red Hat OpenShift Container Platform. On the left, a dark sidebar contains navigation links: Administrator, Home, Operators, Workloads (with a dropdown arrow), Secrets (highlighted), Config Maps, Cron Jobs, Jobs, Daemon Sets, Replica Sets, Replication Controllers, Horizontal Pod Autoscalers, and Networking. The main content area is titled 'Create Image Pull Secret' and includes a sub-header: 'Image pull secrets let you authenticate against a private image registry.' The form contains several input fields: 'Secret Name' with the value 'ibmspp-image-secret', 'Authentication Type' set to 'Image Registry Credentials', 'Registry Server Address' with 'cp.icr.io/cp/sppserver', 'Username' with 'cp', 'Password' (masked with dots), and an empty 'Email' field. A blue '+ Add Credentials' link is positioned below the password field. At the bottom, there are 'Create' and 'Cancel' buttons. A red circle highlights the 'Project: openshift-marketplace' dropdown menu at the top of the form.

Figure 3-3 Creating an image pull secret to access an IBM registry in openshift-marketplace

Adding the pull secret stores the credentials in Red Hat OpenShift that are required to access the IBM registry. The registry is accessed to pull the Docker image that is needed to deploy the IBM Spectrum Protect Plus Operator.

The next step is to add the IBM registry as a source in the Red Hat OpenShift Catalog.

3.2.2 Step 3: Adding the IBM registry to the IBM Spectrum Protect Plus operator

1. To add the IBM registry, click **Administrator** → **Administration** → **Cluster Settings** and then, select the **Configuration** tab, as shown in Figure 3-4.

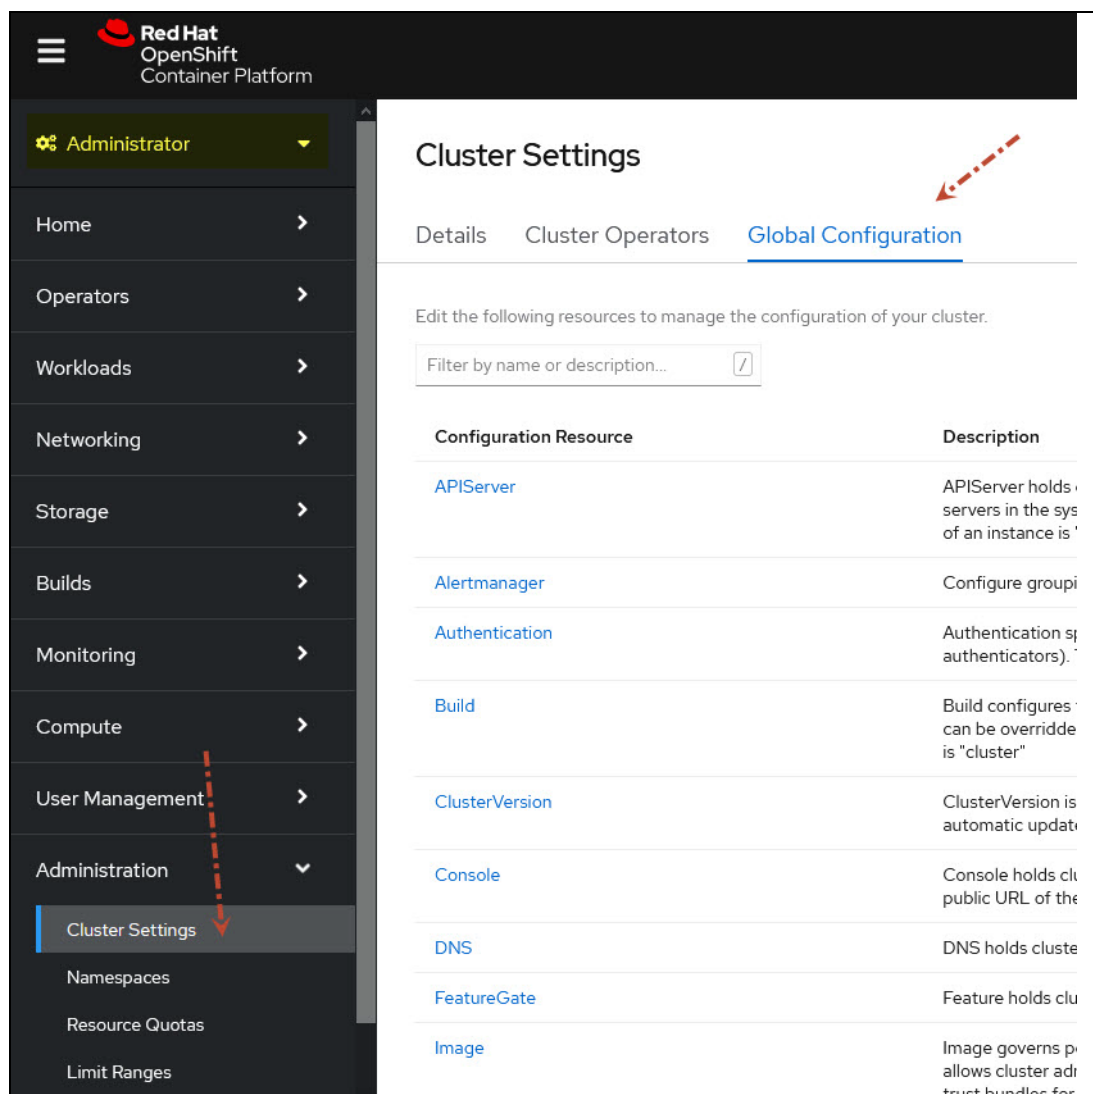


Figure 3-4 Browsing the cluster settings to add an IBM registry as an Operator Hub resource

2. In this window, look for the OperatorHub (see Figure 3-5 on page 31) configuration resource and click it to open it.

OperatorHub	OperatorHub is the Schema for the operatorhubs API. It can be used to change the state of the default hub sources for OperatorHub on the cluster from enabled to disabled and vice versa.
-------------	---

Figure 3-5 OperatorHub configuration resource item

3. From the OperatorHub Details window, in the Sources tab, click **Create Catalog Source** and specify the following values, as shown in Figure 3-6:

- Catalog Source Name: ibm-spp-operator
- Display Name: IBM SPP Operator
- Publisher Name: IBM
- Image: icr.io/cpopen/spp-operator-catalog:latest
- Availability: Namespaced catalog source
- Namespace: openshift-marketplace

Create CatalogSource

Create a CatalogSource in order to make operators available in OperatorHub.

CatalogSource name *

Display name

Publisher name

Image (URL of container image) *

URL of container image hosted on a registry

Availability

☐ Cluster-wide CatalogSource
Catalog will be available in all namespaces

☒ Namespaced CatalogSource
Catalog will only be available in a single namespace

Namespace *

Figure 3-6 Adding an IBM registry as an OperatorHub source

4. The IBM Spectrum Protect Plus Operator is now available from the OperatorHub window by clicking **Administrator** → **Operators** → **OperatorHub**, as shown in Figure 3-7.

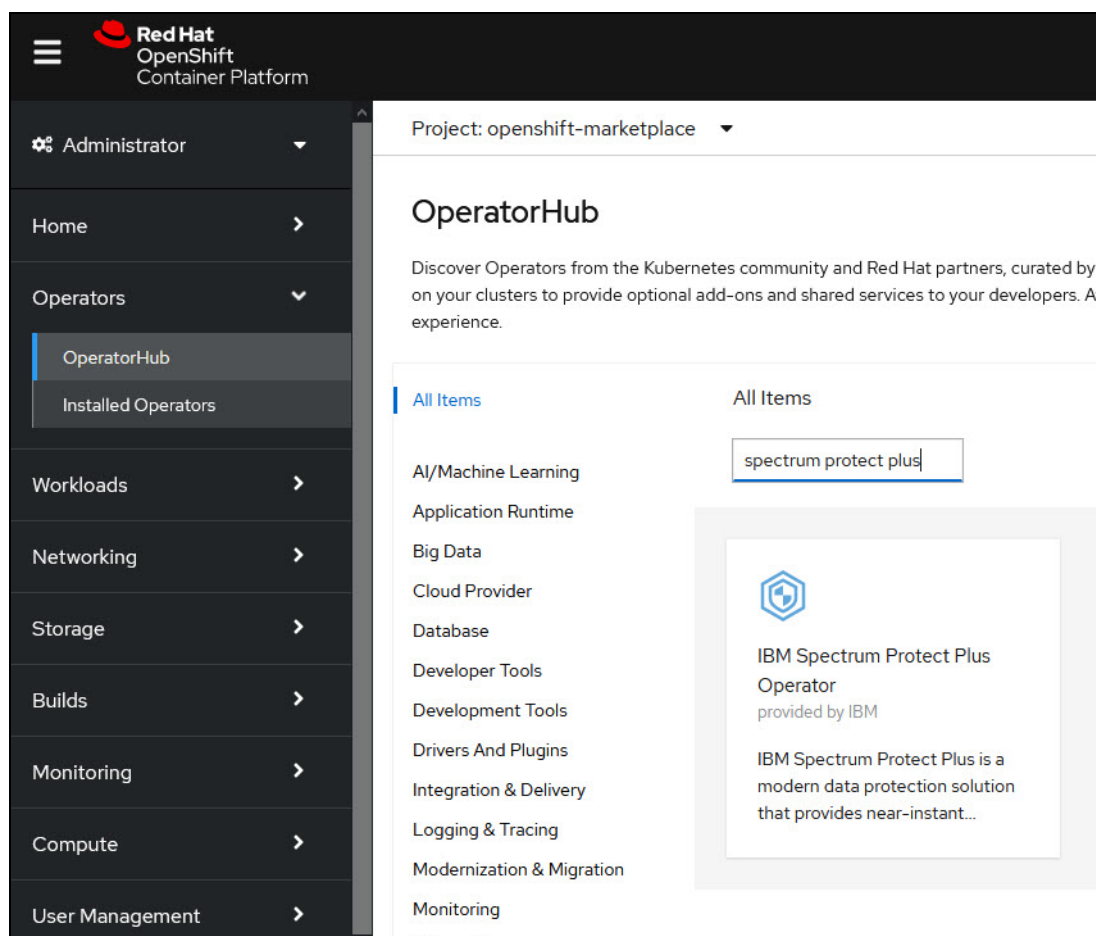


Figure 3-7 IBM Spectrum Protect Plus Operator available in the OperatorHub

3.2.3 Steps 4 and 5: Creating project and image pull secret

Note: A project or namespace must be created and the pull secret added before the IBM Spectrum Protect Plus operator is deployed.

Projects are used to group and isolate related objects. Projects that start with `openshift-` and `kube-` are default projects. These projects host cluster components that run as pods and other infrastructure components.

As an administrator, you can grant developers access to specific projects, allow them to create their own projects, and grant them administrative rights *within* individual projects.

Note: Deploying IBM Spectrum Protect Plus server on its own project or namespace is a good practice.

Creating a project for IBM Spectrum Protect Plus server deployment is a one-step process, as shown in Figure 3-8.

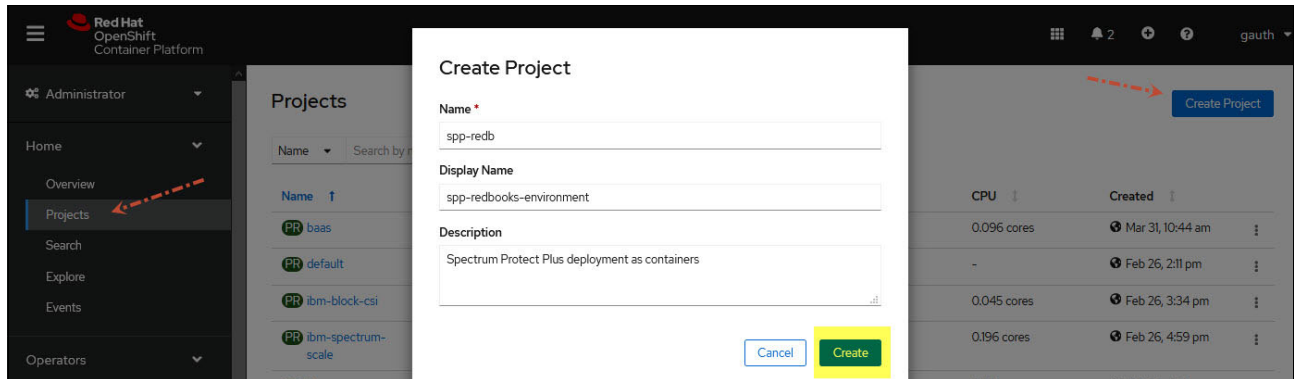


Figure 3-8 Creating the spp-redb project for deploying the IBM Spectrum Protect Plus Operator and server containers

After the project is created, you must create an image pull secret; that is, the credential that is used to access the IBM registry to pull the Docker Image that is used to build the IBM Spectrum Protect Plus operator.

Create that pull secret by clicking **Administrator** → **Workloads** → **Secrets**. On the right side of the window, click **Create** → **Image Pull Secret** (see Figure 3-9).

Be careful to select the correct project space (in our scenario, **spp-redb**) when you are creating this image pull secret.

The screenshot shows the OpenShift console interface. On the left is a dark sidebar with a navigation menu. The 'Workloads' section is expanded, and 'Secrets' is selected. The main content area is titled 'Create Image Pull Secret'. At the top of this area, a dropdown menu shows 'Project: spp-redb', which is circled in red. Below the title, there is a descriptive text: 'Image pull secrets let you authenticate against a private image registry.' The form contains several fields: 'Secret Name' with the value 'ibmspp-image-secret', 'Authentication Type' set to 'Image Registry Credentials', 'Registry Server Address' with the value 'cp.icri.io/cp/sppsver', 'Username' with the value 'cp', 'Password' (masked with dots), and an empty 'Email' field. At the bottom of the form, there is a link '+ Add Credentials' and two buttons: 'Create' and 'Cancel'.

Figure 3-9 Creating an image pull secret in the project

The next step is to install the IBM Spectrum Protect Plus Operator, as described next. After this process is done, you can instantiate an IBM Spectrum Protect Plus server.

Installing IBM Spectrum Protect Plus Operator

After you create the image pull secret, added the IBM image registry, and created the project along with the image pull secret, you are now ready to install the IBM Spectrum Protect Plus operator.

Complete the following steps:

1. Click **Administrator** → **Operators** → **OperatorHub**. Search for IBM Spectrum Protect Plus operator in the GUI by using the keyword spp in the search bar. Then, select the IBM Spectrum Protect Plus Operator from IBM SPP Operator, as shown in Figure 3-10.

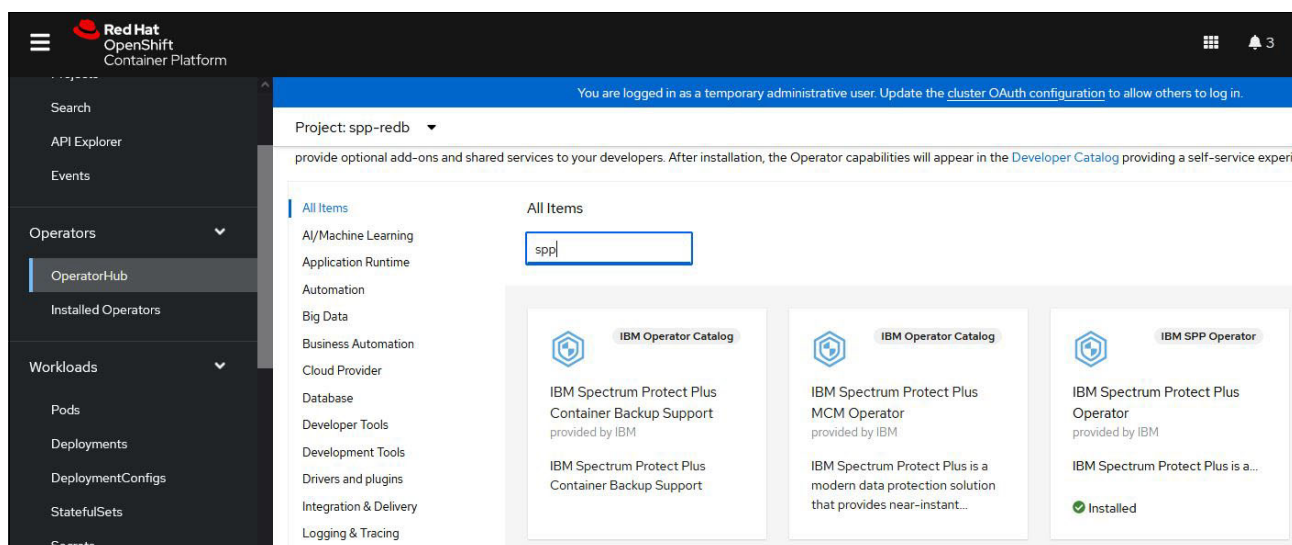


Figure 3-10 Searching the IBM Spectrum Protect Plus Operator from the Operator Hub

2. Click **IBM Spectrum Protect Plus Operator** tile and then, click **Install** (see Figure 3-11).

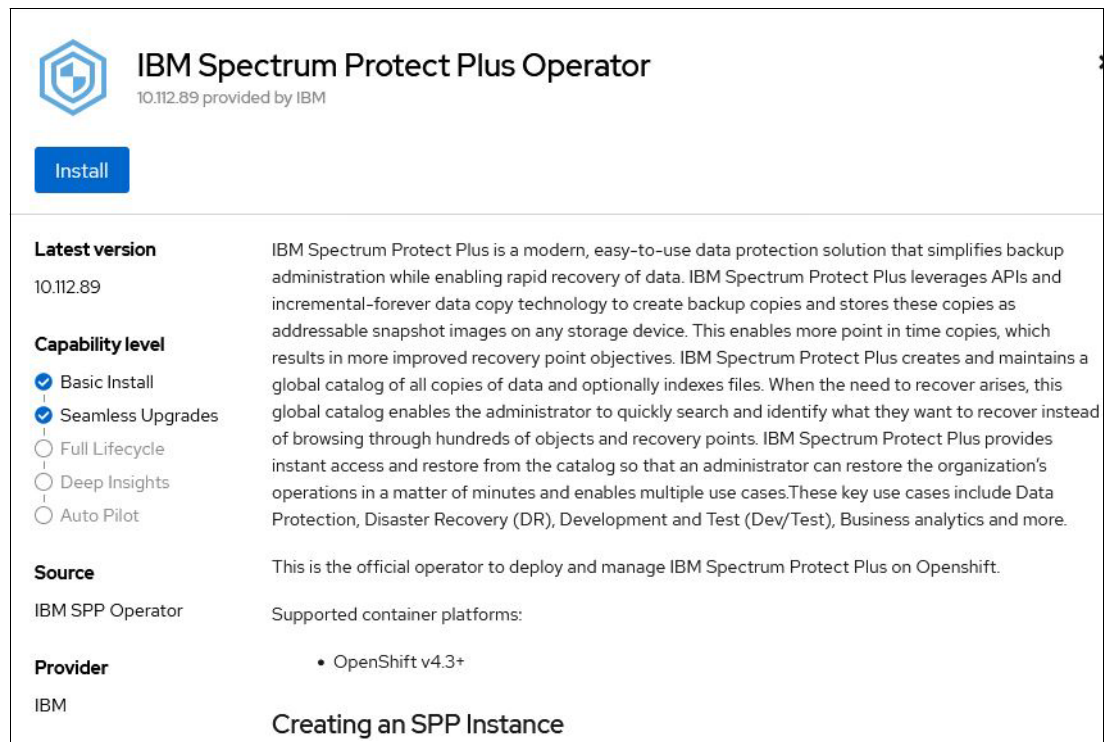


Figure 3-11 Installing the IBM Spectrum Protect Plus Operator from the Operator Hub

3. When the wizard starts, specify the following information:
- The Installation Mode on a specific namespace on the cluster, which is the mode that we created earlier (spp-redb).
 - Update channel: Choose the latest version available (here, it is v10.112), as shown in Figure 3-12.




Figure 3-12 IBM Spectrum Protect Plus operator update channel selection

- Update Approval Strategy: Choose Automatic or Manual. Although your choice depends on your strategy, Manual might be a better choice because you can decide when the upgrade occurs to avoid conflict with backup workload.

- Then, click **Install** to start the installation.

If you selected the Manual update strategy, a message might appear in which you are prompted to approve the chosen installation plan, as shown in Figure 3-13.

IBM Spectrum Protect Plus Operator
10.19.373 provided by IBM

Manual approval required

Review the manual install plan for operators **spp-operator.v10.19.373**. Once approved, the following resources will be created in order to satisfy the requirements for the components specified in the plan. Click the resource name to view the resource in detail.

[Approve](#) [Deny](#) [View installed Operators in Namespace spp-redb](#)

spp-operator.v10.19.373

Name	Kind	Status	API version
CSV spp-operator.v10.19.373	ClusterServiceVersion	Unknown	operators.coreos.com/v1alpha1
CRD ibmspps.ocp.spp.ibm.com	CustomResourceDefinition	Unknown	apiextensions.k8s.io/v1
SA spp-operator	ServiceAccount	Unknown	core/v1
CR spp-operator.v10.19.373-586556d9ff	ClusterRole	Unknown	rbac.authorization.k8s.io/v1
CRB spp-operator.v10.19.373-586556d9ff	ClusterRoleBinding	Unknown	rbac.authorization.k8s.io/v1

Figure 3-13 IBM Spectrum Protect Plus operator Manual Update approval confirmation wizard

As a result, the IBM Spectrum Protect Plus operator is installed as a container that is named `spp-operator-xxx` (where `xxx` is a randomly generated unique string), as shown in Example 3-5.

Example 3-5 IBM Spectrum Protect Plus operator pod running

```
[root@ocp-helper ~]# oc project spp-redb
Now using project "spp-redb" on server "https://api.ocp4.isv.escc.lab:6443".
[root@ocp-helper ~]# oc get pods -n spp-redb
```

NAME	READY	STATUS	RESTARTS	AGE
spp-operator-5f6bf988bc-94wst	2/2	Running	0	12m

The IBM Spectrum Protect Plus server can now be instantiated from the operator. The operator manages the installation for you.

3.3 Deploying IBM Spectrum Protect Plus by using an operator

After you created the image pull secret, added the IBM image registry, created the project along with the image pull secret, and installed the IBM Spectrum Protect Plus operator, you are now ready to create an instance of the IBM Spectrum Protect Plus server by using IBM Spectrum Protect Plus operator.

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 containers, where the different IBM Spectrum Protect Plus Server components are running. For more information, see 3.1.1, “Containers and pods” on page 23.

3.3.1 Creating IBM Spectrum Protect Plus instance

By using the IBM Spectrum Protect Plus operator, the deployment of IBM Spectrum Protect Plus server is straightforward. Complete the following steps:

1. Click **Administrator** → **Operators** → **Installed Operators**. Search for the IBM Spectrum Protect Plus operator tile and then, click **Create Instance**, as shown in Figure 3-14.

(Selecting the **Create IBMSPP** option that is available in the IBM Spectrum Protect Plus section of that same window results in the same action.)

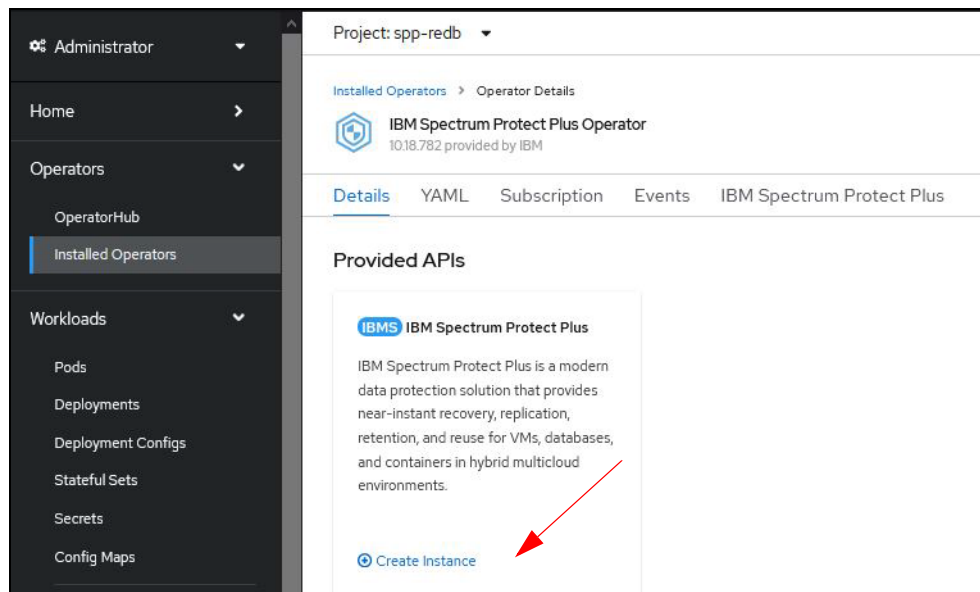


Figure 3-14 Creating an instance from the IBM Spectrum Protect Plus Installed Operator menu

2. The Create Instance wizard prompts you to specify the following information:
 - The IBM Spectrum Protect Plus deployment Name: `ibmspp`.
 - License agreements (click the twistie to display and accept the agreement).
 - Specify the Image registry from where the IBM Spectrum Protect Plus container images are downloaded: `cp.icr.io/cp/sppserver`.
 - The Image Pull Secret (the one you added as part of preparation steps in your IBM Spectrum Protect Plus project environment; in our example, **spp-redb**). The image pull secret in our example is named `ibmspp-image-secret`.

- The hostname, which is used for the route definition, to access that IBM Spectrum Protect Plus deployment from the outside of the Red Hat OpenShift Cluster (the IBM Spectrum Protect Plus GUI). This hostname must not be in use. In our example, we specify: `sppredb.apps.ocp4.isv.escc.lab`.
- The Storage Class, where the persistent storage volumes are created for some of the IBM Spectrum Protect Plus containers, as described in 3.1.2, “Persistent storage” on page 24.
- The IBM Spectrum Protect Plus version that you want to deploy. Choose the latest available from the drop down list.

The first two windows of the installation wizard are shown in Figure 3-15 and Figure 3-16 on page 40.

Figure 3-15 IBM Spectrum Protect Plus Server creation wizard: part 1

Image Registry *
cp.icrio/cp/sppserver
Specify the image registry.

Image Pull Secret *
Select Secret
Select the credentials used to pull the images.

Hostname *
sppredb.apps.ocp4.isv.escc.lab
A public hostname for the SPP instance

Storage Class *
ocs-storagecluster-ceph-rgw
The name of the storage class.

IBM SPP Version *
10.112.124
Select the version to install.

Use same registry for RedHat images
☐ usesameregistryforredhat
Use same registry for RedHat images.

Create **Cancel**

Figure 3-16 IBM Spectrum Protect Plus Server creation wizard: part 2

3. Click **Create**. The Spectrum Protect Plus operator creates the config maps, the secret, and the deployment, which enables creating the pods, persistent storage, and associated PVCs, services, and eventually the routes.

You can monitor pod creation activities by using the CLI, as shown in Example 3-6.

Example 3-6 Using the CLI

```
[[root@ocp-helper ~]# oc get pods -n spp-redb
```

NAME	READY	STATUS	RESTARTS	AGE
spp-awsebs-6b7f7ffc6f-f9z2b	1/1	Running	0	105s
spp-awsec2-9489c5d49-9zp55	1/1	Running	0	108s
spp-ingressproxy-54d7dcc67c-6m2lr	0/1	Running	0	69s
spp-manager-586c57559b-84pd7	0/1	Running	0	115s
spp-operator-664f985744-z9qmh	2/2	Running	0	17h
spp-plugins-mongo-66f478dbd7-6d6vq	1/1	Running	0	113s
spp-plugins-redis-6cbb55b57-w6cgl	1/1	Running	0	110s
spp-proxy-755d6bf5c9-mqlcg	0/1	Running	0	65s
sppdbmongo-5d94fdf7f9-kbdbx	1/1	Running	0	100s
sppdbmongo2-6695ffcf4c-zcdp9	1/1	Running	0	99s
sppdbpostgres-64b9d68989-tprtzt	1/1	Running	0	94s
sppkc-77698479dd-5nfcg	1/1	Running	0	92s
sppnodejs-65fbcd4d55-nk6zl	0/1	Running	0	88s
sppui-58c76fd85b-gthjj	0/1	Running	0	74s
sppvadb-5dd67568b5-t88rz	1/1	Running	0	77s
sppvirgo-55886cc57d-gvc89	0/1	ContainerCreating	0	81s
...				

4. Alternatively, you can use the Red Hat OpenShift web console. Click **Operators** → **Installed Operators**, and select your project. Then, click the IBM Spectrum Protect Plus Operator.
5. Select the **IBM Spectrum Protect Plus** tab and then, click the name for the instance (in our example, ibmspp), as shown in Figure 3-17 on page 41.

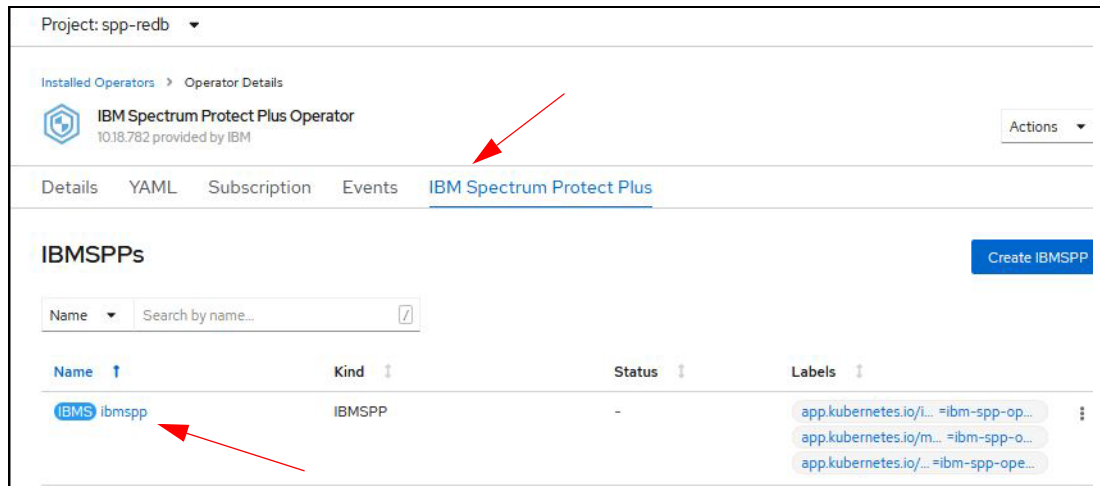


Figure 3-17 Browsing the menu to find the IBM Spectrum Protect Plus custom resource

6. For verification, you can select the **Resources** tab, as shown in Figure 3-18.

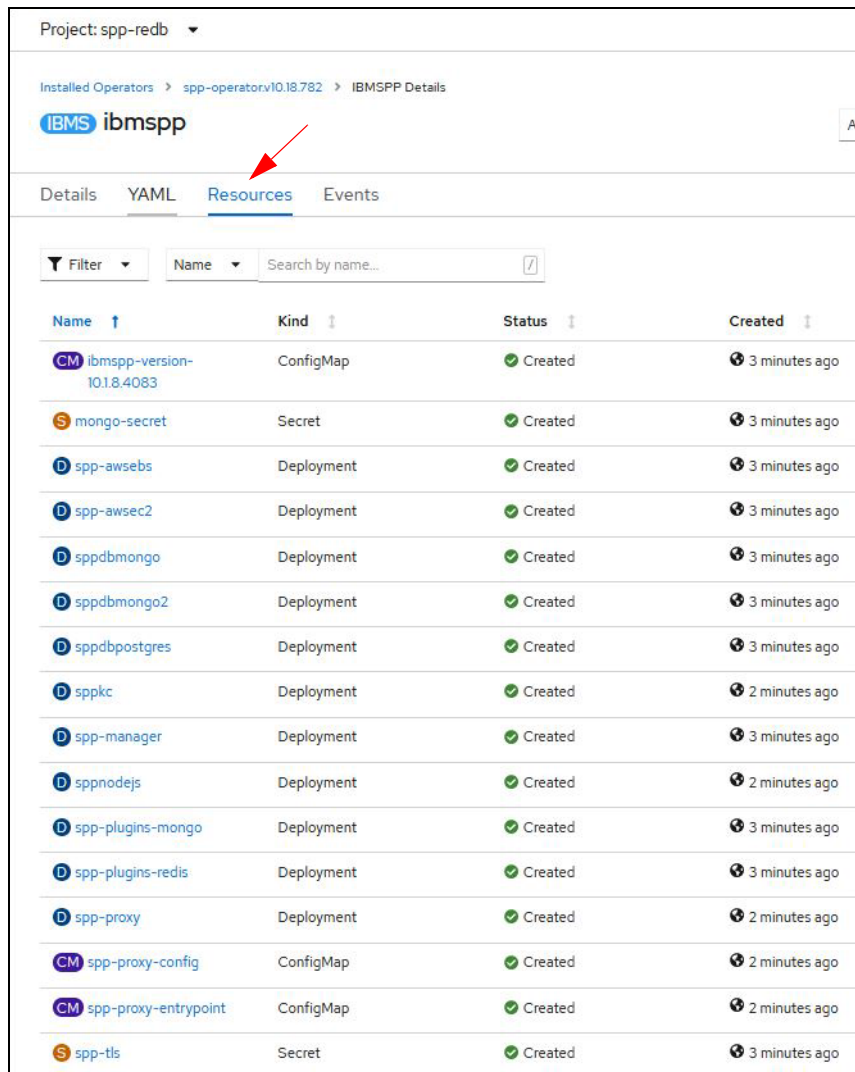


Figure 3-18 IBM Spectrum Protect Plus Resource deployment after instance creation

7. After the operator creates the route as part of the setup process, you can access the IBM Spectrum Protect Plus Server User Interface and proceed with steps that are specific to IBM Spectrum Protect Plus configuration, such as:
 - Defining an administrator
 - Configuring object storage or a vSnap server
 - Configuring the IBM Spectrum Protect Plus Server catalog backup
 - Creating service-level agreements (SLAs)
 - All IBM Spectrum Protect Plus administrative-related tasks
8. Access the IBM Spectrum Protect Plus web interface. The connection information is available from by selecting **Red Hat OpenShift Container Platform user interface** → **Networking** → **Routes** menu or the use of the CLI, as shown in Example 3-7.

Example 3-7 Getting the route information to access IBM Spectrum Protect Plus web interface

```
[root@ocp-helper ~]# oc get routes -n spp-redb
```

NAME	HOST/PORT	PATH	SERVICES	PORT
TERMINATION	WILDCARD			
spp-p9lwx	sppredb.apps.ocp4.isv.escc.lab	/	sppingressproxy	<all>
edge/Redirect	None			

3.3.2 Reviewing the IBM Spectrum Protect Plus deployment

As a result of the IBM Spectrum Protect Plus server deployment through the IBM Spectrum Protect Plus Operator, the following components are created:

- ▶ Sixteen pods
- ▶ Three services that are dedicated to the IBM Spectrum Protect Plus server
- ▶ One route that is used to access the application

If you want to configure the backup of your different projects, namespaces, and PVCs, you must deploy the backup as a service (BaaS) components. For more information, see Chapter 4, “Container Backup Support” on page 61.

The following commands that are run from the Red Hat OpenShift Container platform console can help you to check whether the resources are running and available:

- ▶ `oc project spp-redb`
Log on to the project that you created for your IBM Spectrum Protect Plus server deployment.
- ▶ `oc get pods`
List the pods where IBM Spectrum Protect Plus server components are running, and their status.
- ▶ `oc get services | grep spp`
List services that are associated with IBM Spectrum Protect Plus Server.
- ▶ `oc get routes`
List the route that was created to access IBM Spectrum Protect Plus server User Interface from outside the Red Hat OpenShift cluster.

3.4 Configuring an IBM Spectrum Protect Plus Server catalog backup

After you successfully installed the IBM Spectrum Protect Plus server as a containerized application, its catalog backup must be configured. The catalog contains all metadata information that is required to track the backups activities and recovery points information.

The first time that you connect to the IBM Spectrum Protect Plus server user interface, you are prompted to change the default credentials information by specifying a new administrator name and password of your choice.

Note: Default login credential is: admin/password.

Then, you can log in and start configuring the IBM Spectrum Protect Plus environment.

To backup the IBM Spectrum Protect Plus server catalog, the first step is to configure a vSnap or IBM Cloud Object Storage. Then, you must define an SLA policy that points to that newly defined storage.

In this section, we describe how to configure object storage and define an SLA to backup the catalog.

3.4.1 Defining a cloud storage to be used by IBM Spectrum Protect Plus

First, define cloud storage that points to your S3 object storage. Browse to the **System Configuration** → **Storage** menu and click the **Cloud Storage** tab. From this menu, click **Add cloud storage**. The wizard opens and guides you through the configuration of your S3 object storage provider.

Note: For more information about a list of supported Object Storage providers, see this IBM Documentation [web page](#).

When adding your object storage as a storage target for backup, it can be used for many different backup workloads, including the following examples:

- ▶ Catalog backups
- ▶ Container backups that can direct backup to S3
- ▶ Copies and archive processes that use S3 as a copy destination

Different object storage buckets must be specified to separate the backup and copies processes from a logical standpoint on the object storage.

As part of the wizard steps, after the endpoint URL is specified, you can retrieve the list of buckets that are defined on your object storage. On the object storage side, be sure that different buckets are prepared in advance and have meaningful names so that you can identify them when you configure the IBM Spectrum Protect Plus server side.

Note: The catalog backup to object storage uses the backup storage bucket, which is defined as a parameter in the Object Storage definition.

Figure 3-19 shows the Cloud storage window, which is accessible by using the **System Configuration** → **Storage** menu. From here, you can start the wizard that helps you to add your cloud storage.

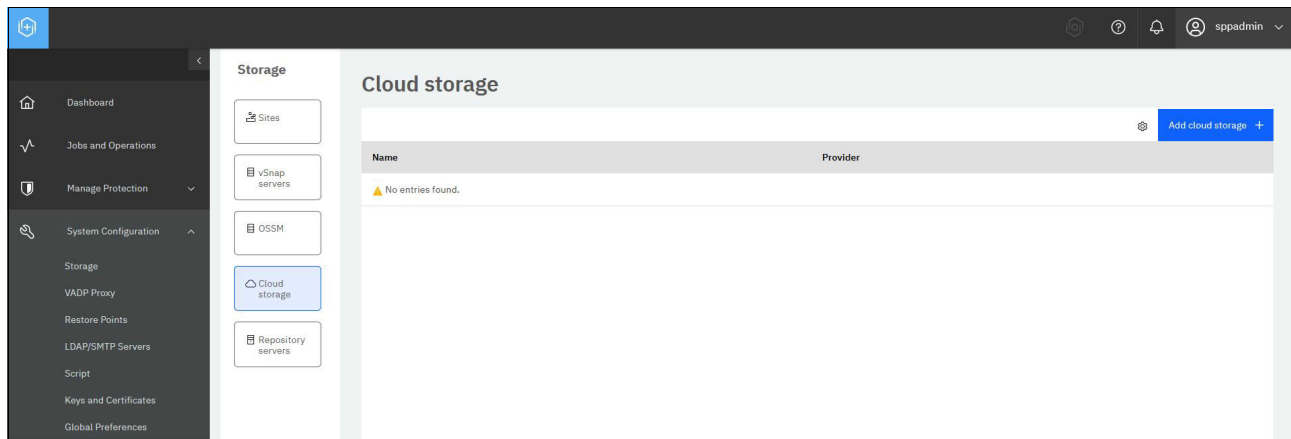


Figure 3-19 Cloud Storage window in IBM Spectrum Protect Plus

By using the wizard, you specify your Object Storage type (in our case, S3 compatible Storage tile). Click **Next**.

Then, in the Cloud Details window, you specify the name of your choice and the required credentials to access that storage endpoint (Access key and Secret key). Eventually, you also specify the certificate because the S3 connection is going to be encrypted. Click **Next**.

The next window is where you specify your object storage endpoint URL, which then enables you to connect and retrieve the list of defined buckets, as shown in Figure 3-20.

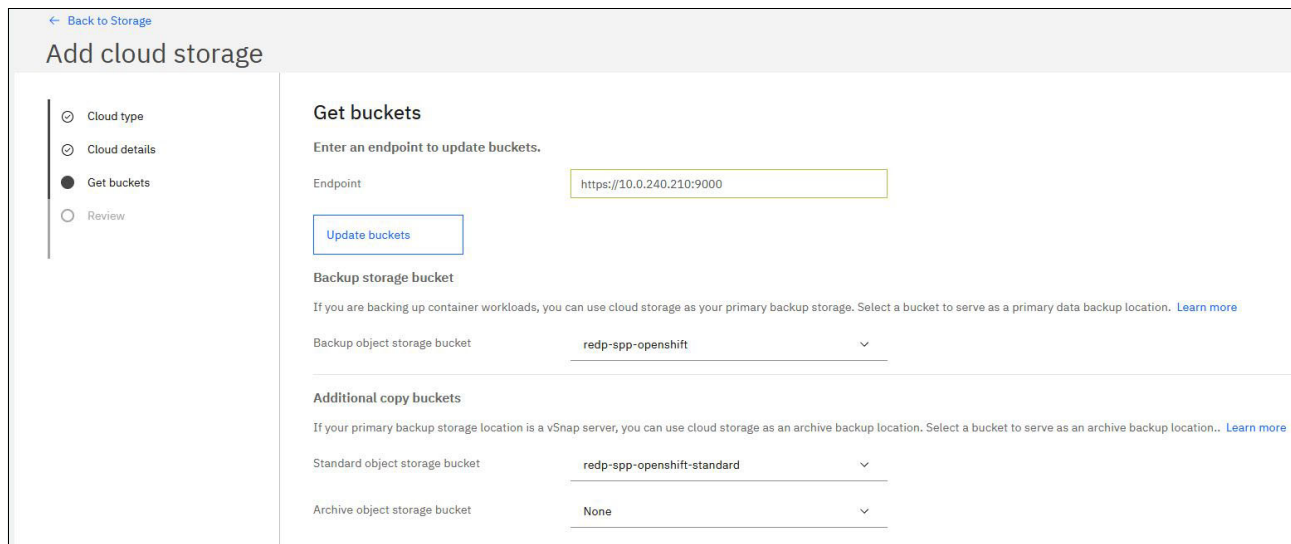


Figure 3-20 Listing buckets that are defined in the specified S3 object storage endpoints

The following buckets are specified in the window that is shown in Figure 3-20 on page 44:

- ▶ **Backup storage bucket**

This bucket is used to store any type of workload that is eligible for direct backup to object storage, whenever cloud storage is selected in an SLA. As of IBM Spectrum Protect Plus 10.1.12, these workloads can be container backups or IBM Spectrum Protect Plus server catalog backups.

- ▶ **Additional copy buckets → Standard object storage bucket**

This bucket is used to store *incremental* copies. The incremental copy feature is valid for any supported workload in IBM Spectrum Protect Plus, except the Application logs (for example, databases logs).

- ▶ **Additional copy buckets → Archive object storage bucket**

This bucket is used to store *full* copies. The archive copy feature is valid for any supported workload in IBM Spectrum Protect Plus, except the Application logs (for example, databases logs).

After the cloud storage is created, you can use that storage target within an SLA policy as described next.

3.4.2 Defining an SLA policy to protect the IBM Spectrum Protect Plus catalog

Starting with IBM Spectrum Protect Plus 10.1.12, you can configure the internal catalog to be backed up directly on Object Storage. After an Object Storage endpoint is configured on the IBM Spectrum Protect Plus server (as described in 3.4.1, “Defining a cloud storage to be used by IBM Spectrum Protect Plus” on page 43), you must create an SLA with the specific IBM Spectrum Protect Plus catalog category, as shown in Figure 3-21.

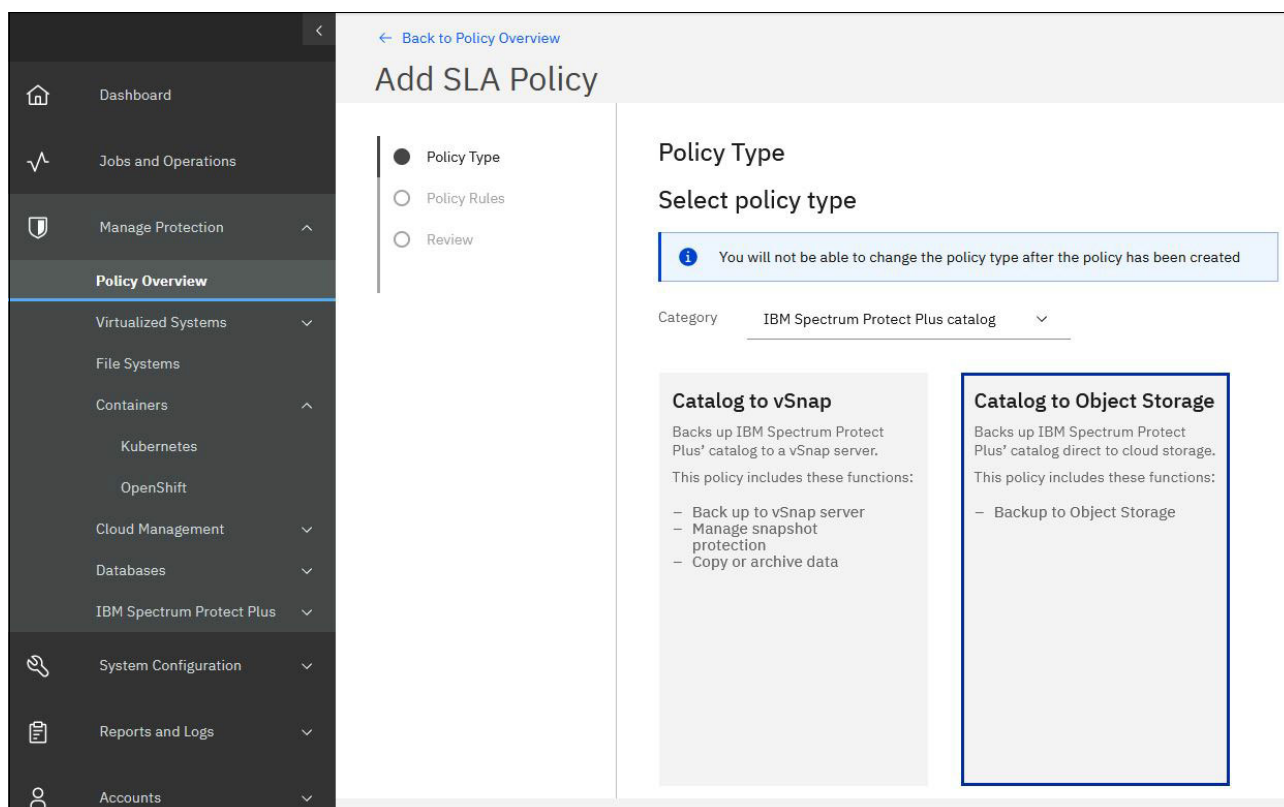


Figure 3-21 Adding an SLA to back up an IBM Spectrum Protect Plus catalog that uses its SLA-specific category

This SLA is created by clicking **Add SLA Policy** in the user interface and then, **Manage Protection** → **Policy Overview** window.

The following settings are configured by using the wizard:

- ▶ Name: The name that you want to use for this SLA.
- ▶ Retention: The number of days that you want to keep this catalog backup.
- ▶ Schedule information: Frequency and start time.
- ▶ Destination: Object Storage.
- ▶ Target Object Storage: The Object storage that you configured as described in 3.4.1, “Defining a cloud storage to be used by IBM Spectrum Protect Plus” on page 43.

An example is shown in Figure 3-22.

← Back to Policy Overview

Add SLA Policy

☑ Policy Type

● Policy Rules

○ Review

Policy Rules

Current Policy Type: Catalog to Object Storage ⓘ

Name
SPPCatalogtoS3

Retention 5 - + Days ▾

☐ Disable Schedule

Repeats Daily ▾ Every: 1 - + day(s)

Start Time 10/24/2022 01:00 Europe/Paris

Destination Object Storage ▾

Target Object Storage minio ▾

Figure 3-22 IBM Spectrum Protect Plus SLA for catalog backup details

Now that the object storage is configured and the SLA is defined, one last step must be completed: assign the specific SLA that was created to the catalog backup. This process is done by clicking menu **Manage Protection** → **IBM Spectrum Protect Plus** → **Backup**.

In this window, select the SLA policy that was created and click **Save**, as shown in Figure 3-23.

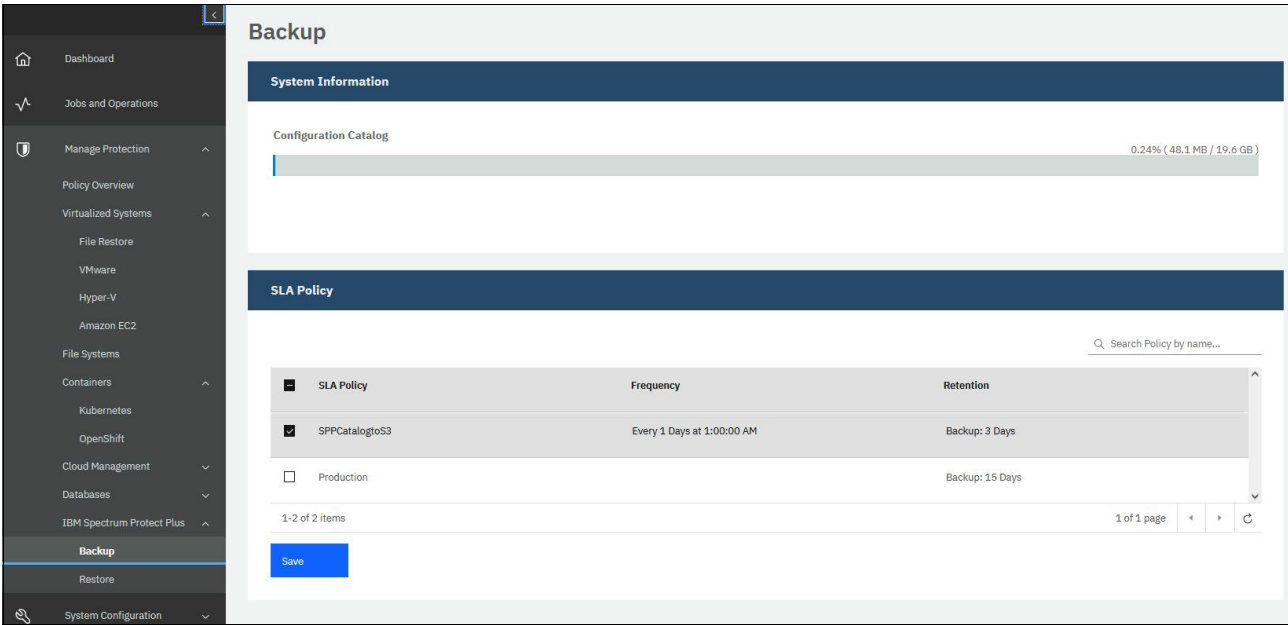


Figure 3-23 Assigning a catalog backup SLA to enable an IBM Spectrum Protect Plus catalog backup

After the SLA is assigned, you can immediately trigger a catalog backup by clicking **Action** at the bottom of the window (see Figure 3-24).

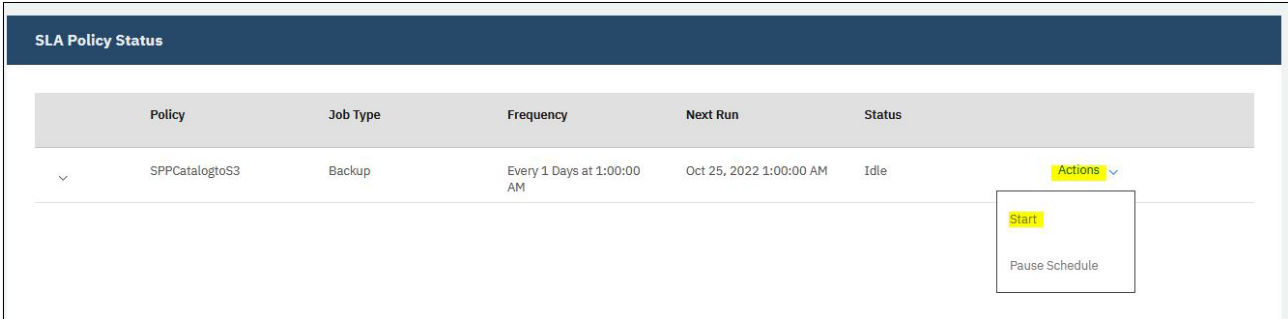


Figure 3-24 Manually starting a catalog backup by selecting the SLA action

Note: IBM Spectrum Protect Plus catalog backup can be directly stored to Object Storage or copied to Object Storage. Depending on how you configured it, the procedure and requirements to recover it from Object Storage are different.

If the catalog was stored directly to object storage, it can be restored directly from that Object Storage by clicking **Manage Protection** → **IBM Spectrum Protect Plus** → **Restore** and then, selecting the **From cloud storage** tab.

If the catalog was stored to vSnap and then copied to an object storage, a vSnap is required (for staging the content back from Object Storage). To access this catalog recovery point, click **Manage Protection** → **IBM Spectrum Protect Plus** → **Restore** and then, select the **From vSnap** tab.

3.5 Upgrading an IBM Spectrum Protect Plus containerized server

IBM Spectrum Protect Plus server can require product updates, such as security fixes or new features. They also can include minor fixes, which are part of the same update channel, or they might require a subscription to a new update channel, as in when a new product version is released.

Note: IBM Spectrum Protect Plus server upgrade is a disruptive process. Be sure that no jobs are running before starting the upgrade process. For more information about general recommendations about how to prepare for an upgrade, see this IBM Documentation [web page](#).

At the time of this writing, the latest IBM Spectrum Protect Plus version is 10.1.12 with the associated update channel v10.112 for the IBM Spectrum Protect Plus Operator.

Update channels also are available for IBM Spectrum Protect Plus version 10.1.11 (Operator update channel v10.111) and IBM Spectrum Protect Plus version 10.1.8 (Operator update channel v10.18).

Thanks to the advanced automation capabilities brought by the IBM Spectrum Protect Plus Operator, code updates are greatly simplified.

When the updates are within the same code stream (for example, 10.1.12_, they are automatically made available for installation. Depending on the configuration you selected for the Update Approval as part of the IBM Spectrum Protect Plus Operator installation, the updates are automatically or manually installed.

The available updates are shown on the Operator window (see Figure 3-25). Here, you can see the message Upgrade available because we chose the Manual Update approval strategy. If you opt in for the Automatic approval, the operator is automatically updated.

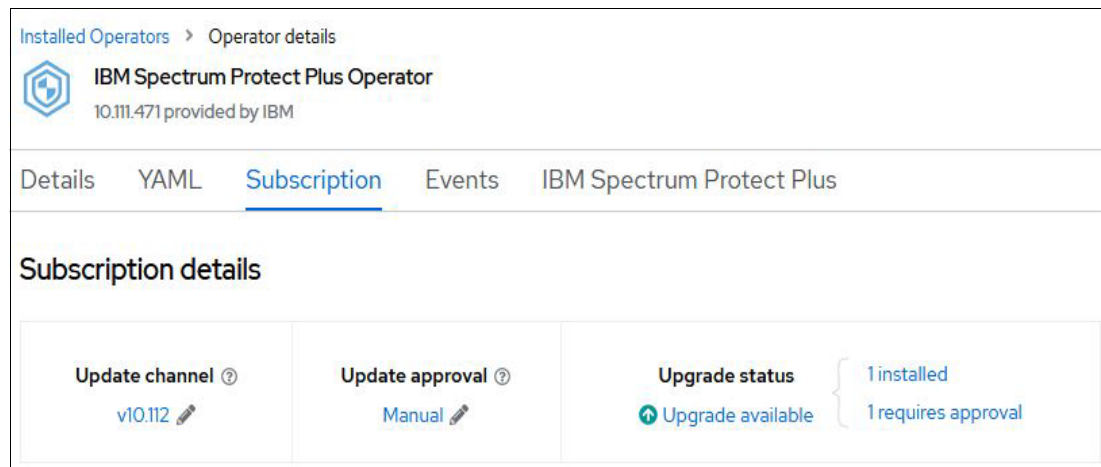


Figure 3-25 Operator Subscription version and upgrade status

If you need to switch from one IBM Spectrum Protect Plus version to the next (for example, you want to upgrade from version 10.1.11 to version 10.1.12), you must update your Operator subscription channel. This process first upgrades the IBM Spectrum Protect Plus Operator, which then enables you to trigger the upgrade of IBM Spectrum Protect Plus server components.

Figure 3-26 shows how to change the IBM Spectrum Protect Plus Operator subscription channel from Version 10.1.11 to Version 10.1.12. Complete the following steps:

1. From the **Operators** → **Installed Operators** menu, select the installed operator of the specific project environment. Select the **Subscription** tab and then, click the pen icon to edit the Subscription channel.

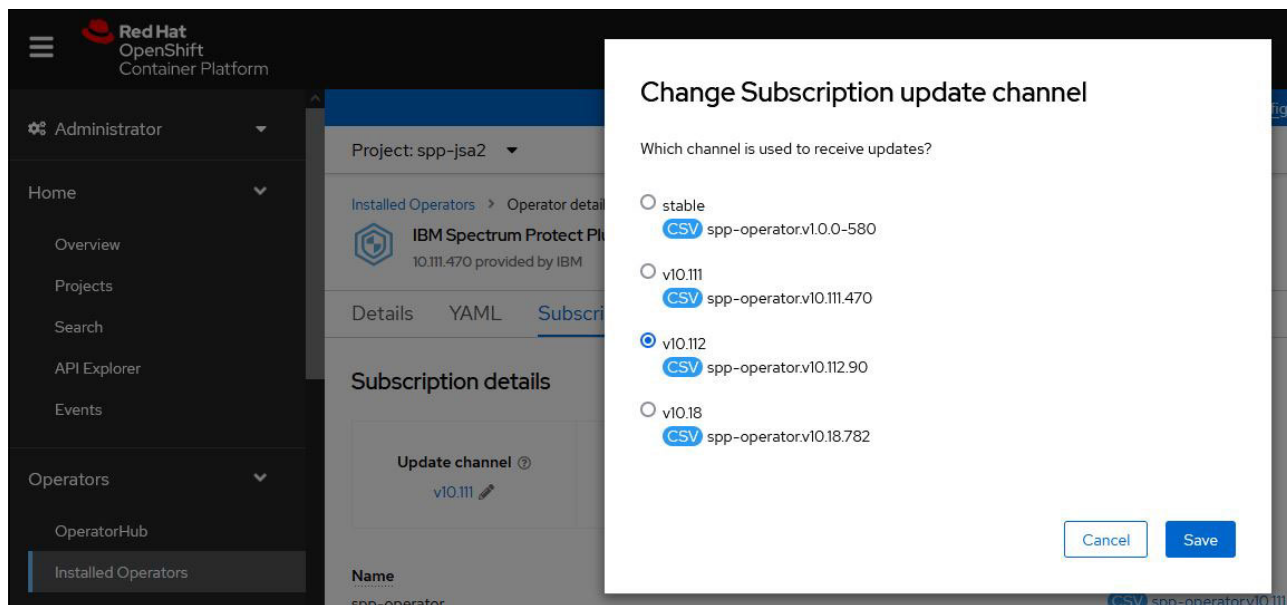


Figure 3-26 IBM Spectrum Protect Plus Installed Operator update channel modification to Version 10.112

2. After you change the Operator update channel, the Operator pods are updated to the corresponding version. Again, depending on the update approval strategy that you used, you might need to confirm the installation by approving the installation plan, as shown in Figure 3-27.

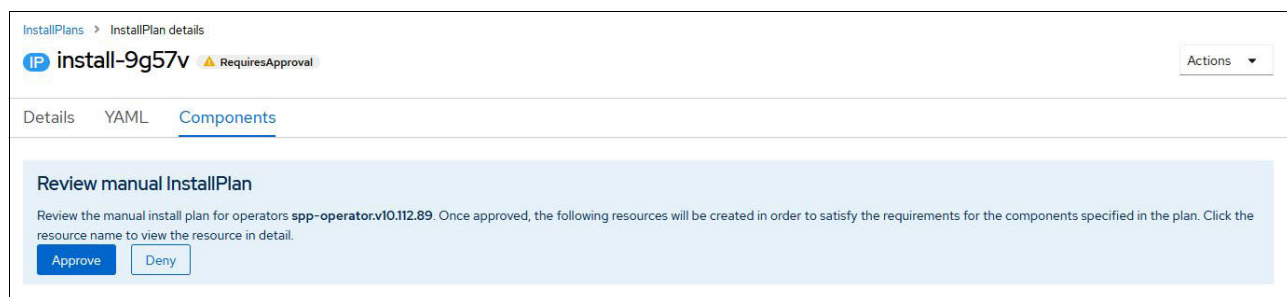


Figure 3-27 Reviewing and approving the spp-operator pods upgrade plan following an upgrade channel update

3. You can follow the installation from the user interface by clicking **Operators** → **Installed Operator**. Then, select your operator and click the **Events** tab.

4. Follow the operator upgrade process by using the CLI. Checking the operator pods, you see that the spp-operator pod is changing and the new pod is reporting an updated version, as shown in Example 3-8.

Example 3-8 Checking the IBM Spectrum Protect Plus operator pods version after an upgrade

```
[root@ocp-helper ~]# oc describe pods spp-operator-6d45bbc7f5-p5twm | grep
OPERATOR_CONDITION_NAME
      OPERATOR_CONDITION_NAME: spp-operator.v10.111.471
      OPERATOR_CONDITION_NAME: spp-operator.v10.111.471

....

[root@ocp-helper ~]# oc get pods | grep operator
spp-operator-64dbdc86c6-xsmnx          2/2      Running      0          22s
spp-operator-6d45bbc7f5-p5twm         2/2      Terminating 0          24m

....

[root@ocp-helper ~]# oc describe pods spp-operator-64dbdc86c6-xsmnx | grep
OPERATOR_CONDITION_NAME
      OPERATOR_CONDITION_NAME: spp-operator.v10.112.89
      OPERATOR_CONDITION_NAME: spp-operator.v10.112.89
```

5. Following the spp-operator pods upgrade, the IBM Spectrum Protect Server also can be upgraded to the corresponding version. The following options are available:
 - The use of the IBM Spectrum Protect Plus User Interface wizard, as shown in Figure 3-28.

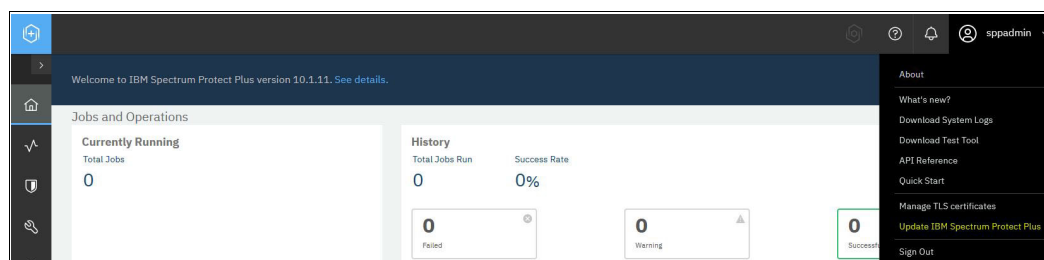


Figure 3-28 IBM Spectrum Protect Plus UI upgrade wizard

- Modify the IBMSPP instance YAML file, as shown in Figure 3-29.

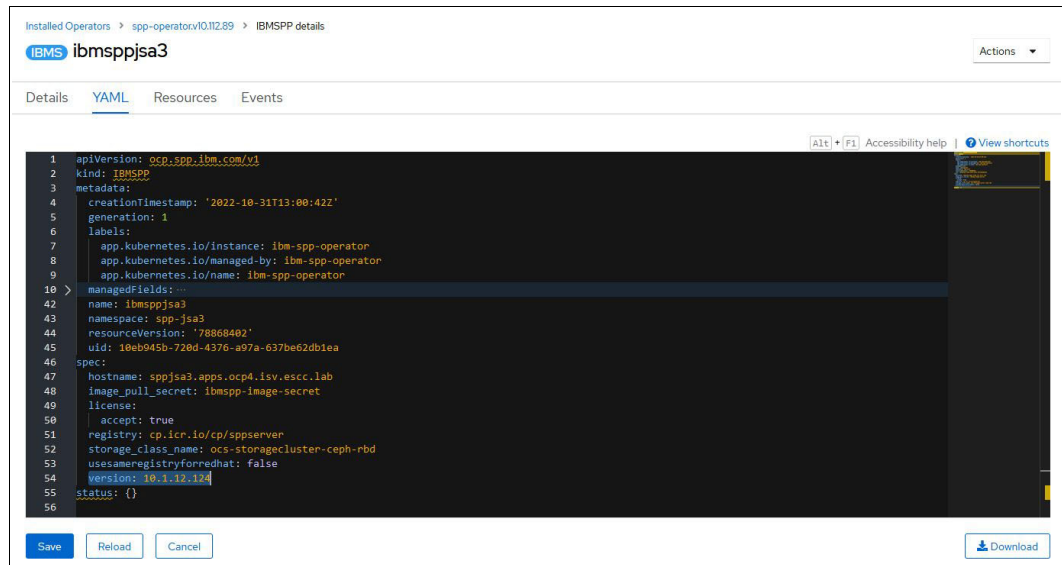


Figure 3-29 Editing the IBM Spectrum Protect Plus server instance to change the version

This YAML file is accessible by completing the following steps:

- Click **Operators** → **Installed Operators** → **IBM Spectrum Protect Plus Operator**.
- Select the **IBM Spectrum Protect Plus** tab.
- Select your IBMSPP instance, and then select the **YAML** tab.
- Search for the keyword `version:` and enter the suitable version information (in our example, 10.1.12.124).

A few seconds after you save the modification that is done in the YAML, the operator triggers the upgrade process of all IBM Spectrum Protect Server pods.

Note: To find the specific IBM Spectrum Protect Plus version to specify here (in our example, 10.1.12.124), click **Workloads** → **ConfigMaps** and find the latest ConfigMaps that was created for the `ibmspp` operator (which is named `ibmspp-version-****`, where `****` is the version for which you are looking).

Alternatively, you can use the commands that are shown in Example 3-9. The details of the ConfigMaps give you the specific version that is going to be deployed for each of the IBM Spectrum Protect Plus server components (it is specified in the tag information under the Data section).

These commands must be used in the context of the project where your IBM Spectrum Protect Plus Server is deployed.

Example 3-9 Determining the correct IBM Spectrum Protect Plus server version

```
[root@ocp-helper ~]# oc get configmaps | grep ibmspp
ibmspp-version-10.1.11.3010    20    49m
ibmspp-version-10.1.12.124    20    27m

[root@ocp-helper ~]# oc describe configmaps ibmspp-version-10.1.12.124
Name:      ibmspp-version-10.1.12.124
Namespace: spp-jsa3
```

```
Labels:      app=ibmspp
            version=10.1.12.124
Annotations: cloudpakId: c77058d714fc4bb89bc29781e0161826
            cloudpakName: IBM Cloud Pak for Spectrum Protect Plus
            cloudpakVersion: 1.0.0
            productChargedContainers: All
            productCloudpakRatio: 1:1
            productID: c77058d714fc4bb89bc29781e0161826
            productMetric: VIRTUAL_PROCESSOR_CORE
            productName: IBM Spectrum Protect Plus
            productVersion: 10.17.5
...
```

Whatever path you opt for, the operator checks for the latest version that is available in the select upgrade channel. Then, it re-creates the different pods that are composing the IBM Spectrum Protect Plus Server to include all of the updates.

The upgrade operation can take several minutes, depending on the size of the environment. One possibility to monitor the IBM Spectrum Protect Server restart after the upgrade is to check the virgo logs by using the `oc logs sppvirgo-xxxx -f` command, which displays the virgo logs in real time. When the virgo logs entries stop, the IBM Spectrum Protect Plus Server is accessible by using its graphical interface.

3.6 IBM Spectrum Protect Plus deployment troubleshooting

The IBM Spectrum Protect Plus Server deployment relies on the IBM Spectrum Protect Plus operator. Therefore, if a problem occurs in the deployment, investigation can be done by checking the operator logs.

3.6.1 Pods information

Understand which pods are running and from which image they built upon, as shown in Example 3-10.

You can see that some of the containers that are deployed are coming from the `cp.icr.io`, which is the IBM registry that you added as part of the set up process. Those containers are the core components of the IBM Spectrum Protect Plus server. Others are coming from `docker.io` Docker registry or `registry.access.redhat.com` and `registry.redhat.io` Red Hat Registry.

Example 3-10 Listing pods image details (including their source repository)

```
[root@console ~]# oc get pods -n spp-redb -o jsonpath="{..image}" | tr -s
'[:,space:]' '\n' | sort -u
cp.icr.io/cp/sppserver/spp-awsebs@sha256:9e862d55a710bcd3bf512f1e2becb8d50e7dcc356
f490703e5e8cd6c51d672af
cp.icr.io/cp/sppserver/spp-awsec2@sha256:c74bca7db1259a110327d4fb95c29fbea8ab7ec38
0b0752b636eaf9594f653a7
cp.icr.io/cp/sppserver/spp-kc@sha256:626c2fa00a65092708e0fc40e53506c95438575143e46
9b7219e0d9defb02aa8
cp.icr.io/cp/sppserver/spp-manager@sha256:51d9d2755110531681155ef52d62c087038ff21b
06d42870fd5a576e62033e13
```

```

cp.icr.io/cp/sppserver/spp-node@sha256:637c23e231adb69a38b9e38a5869e3aea92a951a4ee
e048880d91e2d3ab18c28
cp.icr.io/cp/sppserver/spp@sha256:98029ebc79216aad3431d897c99eeec63b126d57898b161a
e448b88db3144c0f
cp.icr.io/cp/sppserver/spp-uinginx@sha256:5bebed76016499fbc8039714b754b5d36951ad4a
472b32c830edbec1ed7ccf17
cp.icr.io/cp/sppserver/spp-vadp@sha256:92897eb1afb06e399578d47b00a37f83c5902502103
7178a6363aec3ff10c5f9
docker.io/ibmcom/spp-operator@sha256:79789506cd01bf426c6e078ea47b83572ed6405651f0f
d247b74bfab156671a8
registry.access.redhat.com/rhsc1/mongodb-36-rhel7:latest
registry.redhat.io/rhel8/postgresql-96:latest
registry.redhat.io/rhsc1/nginx-116-rhel7:latest
registry.redhat.io/rhsc1/redis-5-rhel7:latest

```

The IBM Spectrum Protect Plus operator runs as one pod, two containers, which can be found by using the command that is shown in Example 3-11, where spp-redb is the project under which IBM Spectrum Protect Plus is deployed. First, find the operator pod name and then, query its details.spec.containers to get the types of container that the pod is running.

Example 3-11 IBM Spectrum Protect Plus operator pod and containers

```

[root@console ~]# oc project spp-redb
[root@console ~]# oc get pods | grep -E "NAME|operator"
NAME                                READY    STATUS    RESTARTS   AGE
spp-operator-6ddc8c7897-bsnmh      2/2      Running   0           10d

[root@console ~]# oc get pods spp-operator-6ddc8c7897-bsnmh -o
jsonpath='{.spec.containers[*].name}*' ; echo
ansible operator*
[root@console ~]#

```

Apparently the IBM Spectrum Protect Plus operators rely on an Ansible container and an operator container. This information is important to know where you need to query the logs for troubleshooting.

3.6.2 Accessing the pods and container logs

The pods logs can be accessed from the Red Hat OpenShift Container Platform CLI by using the **oc logs <pod name>** command. If the pod includes multiple containers, you must specify the **-c** option to indicate which container you want to query, as shown in Example 3-12. This example shows the pod spp-operator, which has two containers: operator and Ansible.

Example 3-12 Getting logs from the operator pod and Ansible container

```

[root@console ~]# oc logs spp-operator-6ddc8c7897-bsnmh
error: a container name must be specified for pod spp-operator-6ddc8c7897-bsnmh,
choose one of: [ansible operator]
[root@console ~]# oc logs spp-operator-6ddc8c7897-bsnmh -c ansible |tail -10
36751 1621520083.40838: getting the next task for host localhost
36751 1621520083.40900: done getting next task for host localhost
36751 1621520083.40974: ^ task is: None
36751 1621520083.41041: ^ state is: HOST STATE: block=5, task=0, rescue=0,
always=0, run_state=ITERATING_COMPLETE, fail_state=FAILED_NONE,

```

```
pending_setup=False, tasks child state? (None), rescue child state? (None), always
child state? (None), did rescue? False, did start at task? False
36751 1621520083.41102: running handlers
```

```
PLAY RECAP *****
localhost      : ok=50   changed=0    unreachable=0    failed=0
skipped=5      rescued=0    ignored=0
```

```
36751 1621520083.41476: RUNNING CLEANUP
```

Checking the log of the operator container (spp-operator pod), you can find more information about all of the tasks that are managed by the operator. Details about Red Hat Ansible tasks that are run by the Ansible container are logged and can be viewed.

Example 3-13 shows the task summary and status of the vadp pod deployment, which occurs when you instantiate an IBM Spectrum Protect Plus server.

Example 3-13 Operator container log showing the vadp related build activity

```
[root@console ~]# oc logs spp-operator-6ddc8c7897-bsnmh -c ansible |tail -10
----- Ansible Task StdOut -----

TASK [ibmspp : Create vadp persistent volume claim] *****
task path: /opt/ansible/roles/ibmspp/tasks/spp.yaml:127

-----
{"level":"info","ts":1621520064.509277,"logger":"logging_event_handler","msg":"[p
laybook task]","name":"ibmspp2","namespace":"spp-redb","gvk":"ocp.spp.ibm.com/v1,
Kind=IBMSPP","event_type":"playbook_on_task_start","job":"2374361282119289216","Ev
entData.Name":"ibmspp : Create vadp persistent volume claim"}

----- Ansible Task StdOut -----

TASK [ibmspp : Create vadp deployment] *****
task path: /opt/ansible/roles/ibmspp/tasks/spp.yaml:134

-----
{"level":"info","ts":1621520066.373031,"logger":"logging_event_handler","msg":"[p
laybook task]","name":"ibmspp2","namespace":"spp-redb","gvk":"ocp.spp.ibm.com/v1,
Kind=IBMSPP","event_type":"playbook_on_task_start","job":"2374361282119289216","Ev
entData.Name":"ibmspp : Create vadp deployment"}

----- Ansible Task StdOut -----

TASK [ibmspp : Create vadp service] *****
task path: /opt/ansible/roles/ibmspp/tasks/spp.yaml:140

-----
{"level":"info","ts":1621520068.3327763,"logger":"logging_event_handler","msg":"[p
laybook task]","name":"ibmspp2","namespace":"spp-redb","gvk":"ocp.spp.ibm.com/v1,
Kind=IBMSPP","event_type":"playbook_on_task_start","job":"2374361282119289216","Ev
entData.Name":"ibmspp : Create vadp service"}
-----
```

The **oc logs** command is valid for all other IBM Spectrum Protect Plus Server pods, such as the `sppvirgo` pod, which can help you track the start of the IBM Spectrum Protect Plus server interface task. A typical situation is when you are waiting for the UI to be ready, you might want to follow the `virgo` activity to understand when that interface is ready. Often, it is ready when the `virgo` log stops intensive message logging (see Example 3-14).

Example 3-14 Checking the `virgo` logs

```
[root@console ~]# oc logs sppvirgo-6664cdbcd6-6dvvc |tail-f
[2021-05-20T14:11:43.134Z] start-signalling-38      <DE0005I> Started bundle
'com.catalogic.ecx.ibmproviderfcm' version '3.8.0'.
[2021-05-20T14:11:45.346Z] fs-watcher              <DE0004I> Starting bundle
'com.catalogic.ecx.serviceprovider.catalog.ibmfcml' version '3.8.0'.
[2021-05-20T14:11:45.458Z] start-signalling-38      <DE0005I> Started bundle
'com.catalogic.ecx.ibmproviderfcm.impl' version '3.8.0'.
[2021-05-20T14:11:46.063Z] start-signalling-37      <WE0001I> Started web
bundle 'com.syncsort.dp.xsb.api.endeavour.update' version '3.8.0' with context
path '/api/endeavour/update'.
[2021-05-20T14:11:46.063Z] start-signalling-37      <DE0005I> Started bundle
'com.syncsort.dp.xsb.api.endeavour.update' version '3.8.0'.
[2021-05-20T14:11:48.290Z] start-signalling-37      <DE0005I> Started bundle
'com.catalogic.ecx.serviceprovider.catalog.ibmfcml' version '3.8.0'.
[2021-05-20T14:11:48.291Z] start-signalling-37      <DE0005I> Started plan
'xsb-ibmproviderfcm' version '2.0.0'.
[2021-05-20T14:12:46.251Z] start-signalling-37      <DE0005I> Started bundle
'com.syncsort.dp.xsb.autoupdate' version '3.8.0'.
[2021-05-20T14:12:46.252Z] start-signalling-37      <DE0005I> Started plan
'xsb-update' version '2.0.0'.
[2021-05-20T14:12:46.258Z] start-signalling-37      <DE0005I> Started plan
'xsb-app' version '2.0.0'.
```

The pods logs also are accessible from the Red Hat OpenShift Container Platform web interface by clicking **Administrator** → **Workloads** → **Pods**, selecting your pod, and then, clicking the **Logs** tab, as shown in Figure 3-30.

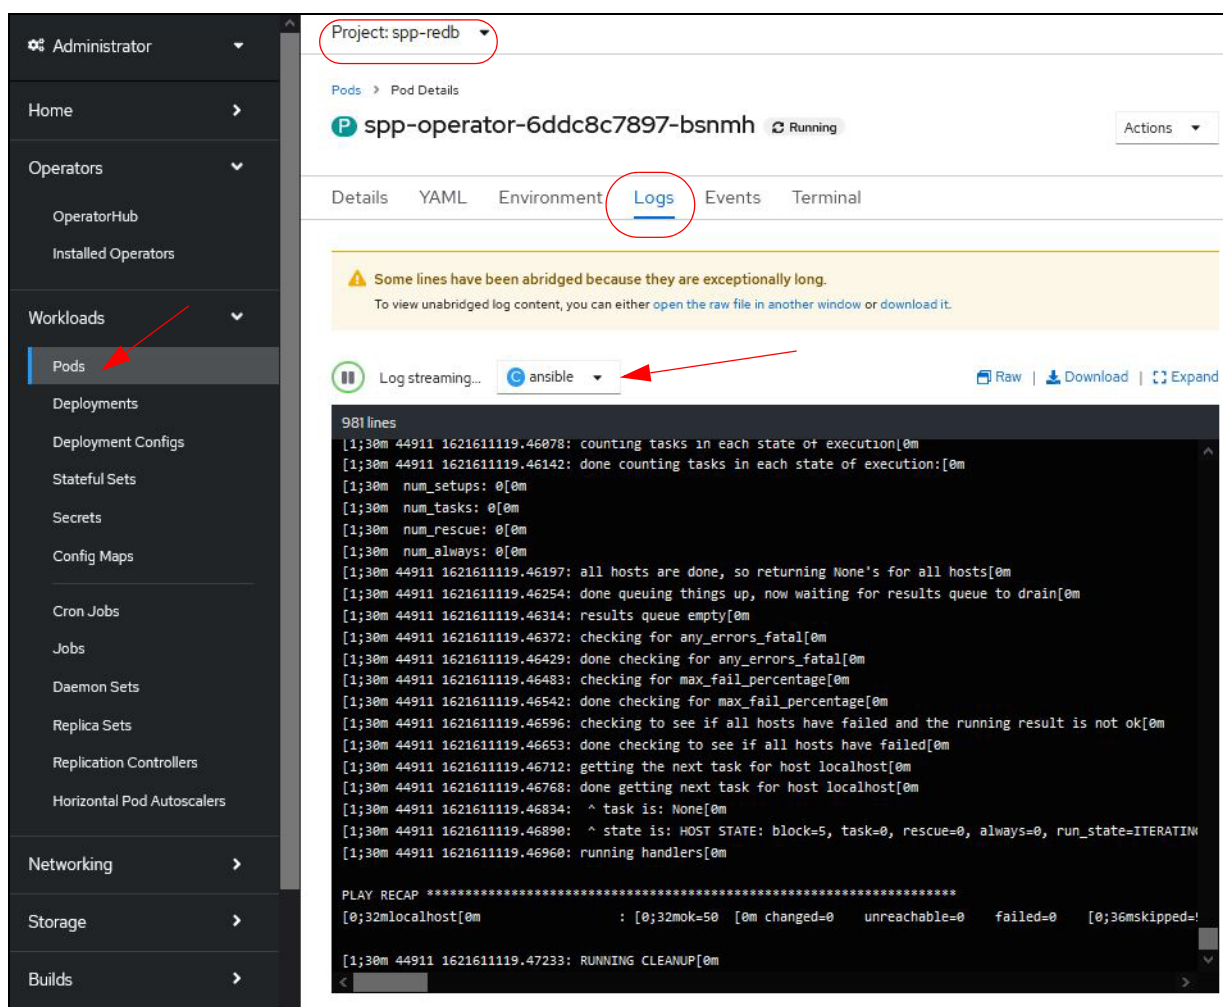


Figure 3-30 Accessing pods logs from the Red Hat OpenShift Container Platform web interface

One of the benefits of the use of the web interface is that it offers log streaming, with which you can watch the logs in real time. This feature is helpful to follow modifications that are applied to a pod.

Note: If you review the operator pods logs and select the operator container, you see summarized output of the Ansible tasks that are performed by the operator, which is helpful to track what was done without too much information about the action.

However, if you select the Ansible container, the information is too verbose about the various Ansible operations that are done to perform all the deployment and updates activities.

3.6.3 Monitoring the IBM Spectrum Protect Plus Server resource usage

Many performance indicators are automatically monitored when you deploy pods within a Red Hat OpenShift Container Platform environment.

After you deployed the IBM Spectrum Protect Plus server by using the IBM provided operator, you can access performance indicators by clicking **Administrator** → **Observe** → **Dashboards**. From the drop-down menu at the upper right, select the dashboard type that you want; for example, **Compute Resources/Namespace (pods)**. Then, select your namespace (spp-redb is our sample project under which you deployed the IBM Spectrum Protect Plus server), as shown in Figure 3-31.

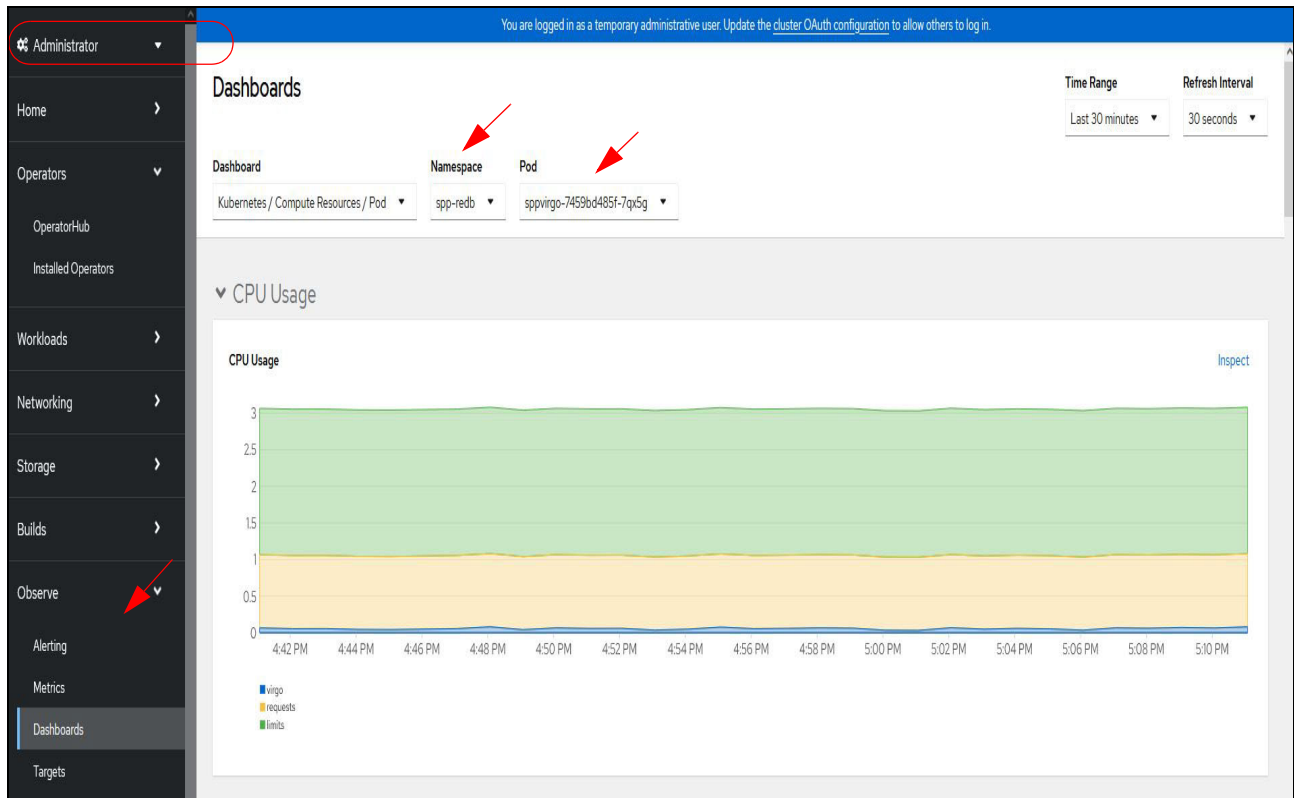


Figure 3-31 Monitoring feature of Red Hat OpenShift Container Platform

This dashboard provides more information about the Compute resources that are used by the IBM Spectrum Protect Plus pods. It also includes information about network traffic (Receive and Transmit bandwidth), and all of that information is available at the pod level.

This information is useful when you must troubleshoot performance-related problems.

3.6.4 Changing the hostname of IBM Spectrum Protect Plus server

A situation might exist in which you need to change the hostname of your IBM Spectrum Protect Plus server.

The Red Hat OpenShift operator makes this task easy because it manages all of the required pods communication updates in the background.

Complete the following steps to change the IBM Spectrum Protect Plus hostname:

1. Click **Operators** → **Installed Operators**.
2. Select the IBM Spectrum Protect Plus Operator that is installed in your project (in our example, spp-redb).
3. In the **IBM Spectrum Protect Plus** tab, click your IBMSPP instance (in our example, **ibmspp**). After the instance is selected, edit the YAML view to access the details of your specific IBMSPP deployment, as shown in Figure 3-32.

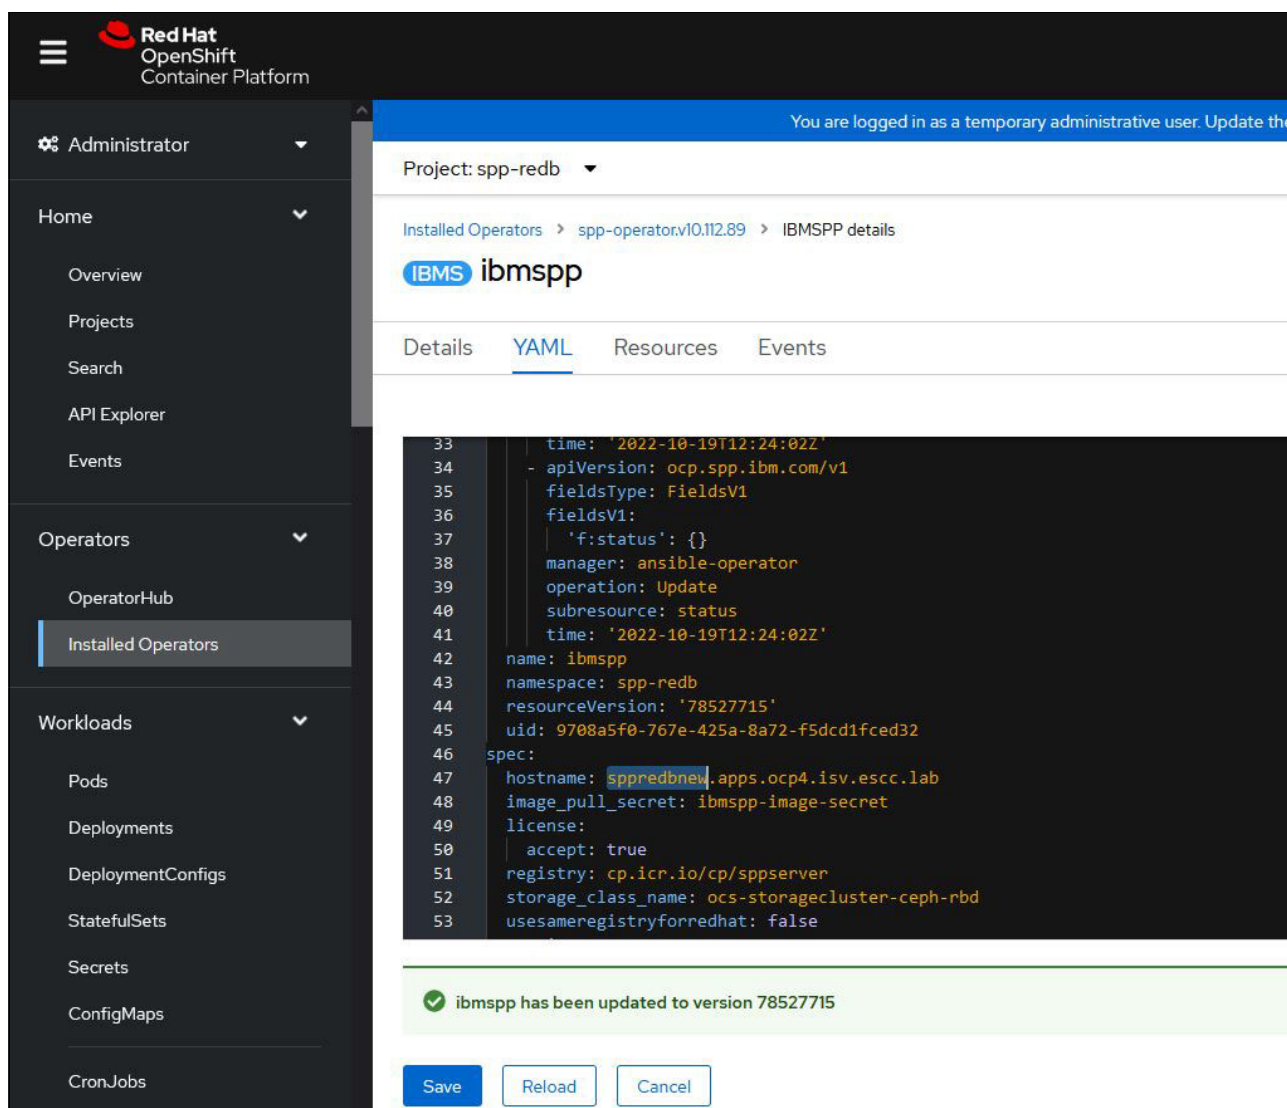


Figure 3-32 Editing the `ibmspp` operator YAML to change the IBM Spectrum Protect Plus server hostname

4. Scroll down to the YAML and edit the hostname value to specify the new hostname. Click **Save**.

After a few seconds, you see that the vadb, virgo, and nodejs pods are re-created automatically, as shown in Example 3-15.

Example 3-15 Routes and pods information

```
[root@ocp-helper ~]# oc get routes
```

NAME	HOST/PORT	PATH	SERVICES	PORT
TERMINATION	WILDCARD			
spp-tbj18	sppredb.apps.ocp4.isv.escc.lab	/	sppingressproxy	<all>
edge/Redirect	None			


```
[root@ocp-helper ~]# oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
spp-awsebs-6b7f7ffc6f-jw4n9	1/1	Running	0	11d
spp-awsec2-9489c5d49-rdsqh	1/1	Running	0	11d
spp-ingressproxy-54d7dcc67c-w2b22	1/1	Running	0	11d
spp-manager-586c57559b-xczz4	1/1	Running	0	11d
spp-operator-f78cd7949-rjjb9	2/2	Running	0	56m
spp-plugins-mongo-66f478dbd7-zvlrb	1/1	Running	0	11d
spp-plugins-redis-6cbb55b57-5pgsg	1/1	Running	0	11d
spp-proxy-755d6bf5c9-khxfs	1/1	Running	0	11d
sppdbmongo-5d94fdf7f9-h6h44	1/1	Running	0	11d
sppdbmongo2-6695ffc4c-sxf6s	1/1	Running	0	11d
sppdbpostgres-64b9d68989-lpzqk	1/1	Running	0	11d
sppkc-77698479dd-mjbtj	1/1	Running	0	11d
sppnodejs-69bbd6b7f6-wswfw	1/1	Terminating	12 (22m ago)	11d
sppui-58c76fd85b-v4jkc	1/1	Running	0	11d
sppvadb-699f4954b6-6dbz1	1/1	Terminating	0	11d
sppvirgo-7459bd485f-xt81k	1/1	Terminating	0	31m

Also, a route is created that corresponds to the new IBM Spectrum Protect Plus hostname that you specified in the operator YAML file, as shown in Example 3-16.

Example 3-16 Routes and pods information after the operator finishes the hostname change

```
[root@ocp-helper ~]# oc get routes
```

NAME	HOST/PORT	PATH	SERVICES	PORT
TERMINATION	WILDCARD			
spp-6mb4c	sppredbnew.apps.ocp4.isv.escc.lab	/	sppingressproxy	<all>
edge/Redirect	None			


```
[root@ocp-helper ~]# oc get pods
```

NAME	READY	STATUS	RESTARTS	AGE
spp-awsebs-6b7f7ffc6f-jw4n9	1/1	Running	0	11d
spp-awsec2-9489c5d49-rdsqh	1/1	Running	0	11d
spp-ingressproxy-54d7dcc67c-w2b22	1/1	Running	0	11d
spp-manager-586c57559b-xczz4	1/1	Running	0	11d
spp-operator-f78cd7949-rjjb9	2/2	Running	0	59m
spp-plugins-mongo-66f478dbd7-zvlrb	1/1	Running	0	11d
spp-plugins-redis-6cbb55b57-5pgsg	1/1	Running	0	11d
spp-proxy-755d6bf5c9-khxfs	1/1	Running	0	11d
sppdbmongo-5d94fdf7f9-h6h44	1/1	Running	0	11d
sppdbmongo2-6695ffc4c-sxf6s	1/1	Running	0	11d
sppdbpostgres-64b9d68989-lpzqk	1/1	Running	0	11d
sppkc-77698479dd-mjbtj	1/1	Running	0	11d
sppnodejs-74fb964bb4-n85bk	0/1	Running	1 (27s ago)	2m22s
sppui-58c76fd85b-v4jkc	1/1	Running	0	11d

sppvadp-54d4df5bc4-kp2hr	1/1	Running	0	2m11s
sppvirgo-7b7b86d84f-9w8w5	0/1	Running	0	2m15s



Container Backup Support

This chapter provides information about the Container Backup Support feature of IBM Spectrum Protect Plus. It covers IBM Spectrum Protect Plus components and prerequisites for Container Backup Support in Red Hat OpenShift.

This chapter includes the following topics:

- ▶ 4.1, “Overview” on page 62
- ▶ 4.2, “Components” on page 64
- ▶ 4.3, “Red Hat OpenShift prerequisites and supported environments” on page 68
- ▶ 4.4, “IBM Spectrum Protect Plus and storage prerequisites” on page 72
- ▶ 4.5, “Container backup and restore types” on page 77
- ▶ 4.6, “Service-level agreement policies for container backup” on page 80
- ▶ 4.7, “Container Backup Support security features” on page 82
- ▶ 4.8, “IBM Spectrum Fusion and IBM Cloud Pak for Data” on page 83

4.1 Overview

Container Backup Support that is provided with IBM Spectrum Protect Plus protects the data of persistent volumes, namespace-scoped resources, and cluster-scoped resources that are associated with containers in a Kubernetes or Red Hat OpenShift Container Platform environment. You can run snapshot backup operations to create locally stored backup snapshots in the cluster, or you can store backup copies on the IBM Spectrum Protect Plus vSnap server or object store for longer-term retention.

4.1.1 Version history

Container Backup Support for Red Hat OpenShift was initially introduced with IBM Spectrum Protect Plus V10.1.7. With more recent versions of IBM Spectrum Protect Plus, more features that are related to container backup were added, as listed in Table 4-1.

Table 4-1 IBM Spectrum Protect Plus Container Backup Support history of versions and features

IBM Spectrum Protect version	Specific features of Red Hat OpenShift container environments
10.1.7	<ul style="list-style-type: none">▶ Initial Red Hat OpenShift container platform support on AMD64.▶ Metadata backups by using Red Hat OpenShift API for Data Protection and Velero.▶ Deploy IBM Spectrum Protect Plus as a container.▶ Restore to alternative namespaces or clusters.▶ IBM block storage Container Storage Interface (CSI) driver support.
10.1.8	<ul style="list-style-type: none">▶ Protect virtual machines (VMs) that are running in Red Hat OpenShift Container Platform containers (10.1.8.1).▶ Back up and restore application-consistent data by using Velero pre- and post-backup hooks (10.1.8.1).▶ Ceph File System (CephFS) and Scale CSI driver support.
10.1.9	<ul style="list-style-type: none">▶ Backup Red Hat OpenShift Container Platform data directly to cloud storage.▶ Install Container Backup Support by using an operator.▶ Enhance security by switching from SSH to RestAPI.
10.1.10	<ul style="list-style-type: none">▶ Reduce storage footprint for containerized IBM Spectrum Protect Plus server (v10.1.10.2).▶ Hitachi NAS (HNAS) and Net App Trident CSI driver support.
10.1.11	<ul style="list-style-type: none">▶ Create a namespace during the restore of persistent volume claims (PVCs).▶ Vertical Pod Autoscaling for Container Backup Support pods.▶ Backup and restore Red Hat OpenShift platform on IBM zSystems.
10.1.12	<ul style="list-style-type: none">▶ Back up the IBM Spectrum Protect Plus catalog directly to cloud storage.▶ Enhanced options for PVC restore jobs.

4.1.2 User roles

Depending on their role, developers of containerized applications and backup administrators interact with different user interfaces to protect persistent data in containers.

Application developer

The enterprise application developer completes the following tasks:

- ▶ Defines application protection requirements (for example, RTO and RPO).
- ▶ Defines the application components that need to be protected.

- ▶ Communicates the above to the backup administrator.
- ▶ Uses the Red Hat OpenShift GUI or command-line interface (CLI) (oc) to complete the following tasks:
 - Creates scripts that allow to manipulate applications in preparation for a backup or after a restore.
 - Assigns annotations to pods to use these scripts as backup or restore pre- and post-hooks.
 - Assigns labels to Red Hat OpenShift Container Platform components that can be bound to backup policies later.
- ▶ Completes the following tasks in the IBM Spectrum Protect Plus user interface (can be done by the backup administrator as well):
 - Starts self-service backup and restore requests.
 - Selects service-level agreement (SLA) policies 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 Red Hat 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 Red Hat OpenShift cluster or updates the cluster properties.
 - Manually runs an inventory to detect cluster resources.
 - Creates SLA policies according to the requirements as provided by the application administrator.
 - 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 Red Hat OpenShift environment and viewing trace log files for Container Backup Support.

4.2 Components

In addition to the IBM Spectrum Protect Plus server and vSnap servers, the container backup support relies on several other components (See Figure 4-1) that must be installed into the backup as a service (BaaS) namespace on a Red Hat OpenShift cluster. These components are described next.

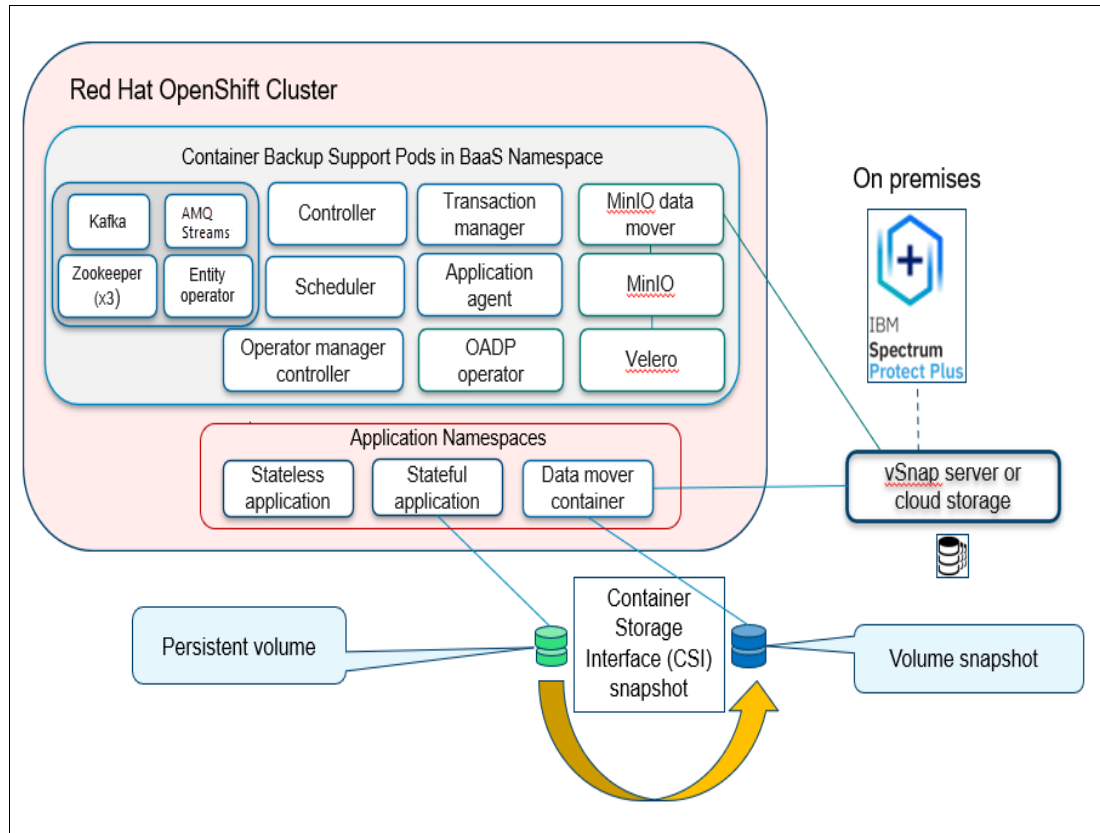


Figure 4-1 Red Hat OpenShift deployment

4.2.1 Container Backup Support (BaaS) namespace

To protect persistent volumes (PVs) that are used by containers and cluster-scoped and namespace-scoped Red Hat OpenShift resources, you must install and configure IBM Spectrum Protect Plus Container Backup Support in a Red Hat OpenShift Container Platform environment.

On Red Hat OpenShift and Kubernetes clusters, you can install Container Backup Support by running an installation script at the CLI. However, as an alternative on Red Hat OpenShift clusters, you also can install Container Backup Support from the Red Hat OpenShift web console.

As a prerequisite, an instance of the IBM Spectrum Protect Plus server must be installed and configured. This instance can be deployed as a containerized application into Red Hat OpenShift or as a virtual appliance on a VMware or Microsoft Hyper-V hypervisor. Container Backup Support protects only persistent storage that was allocated by a storage plug-in that supports the CSI.

Velero and Red Hat OpenShift API for Data Protection

Velero is an open-source tool that is used in container environments to safely back up and restore, perform disaster recovery tasks, and migrate Kubernetes cluster resources and persistent volumes.

Velero and the Red Hat OpenShift API for Data Protection operators are deployed together with the IBM Spectrum Protect Plus Container Backup Support installation. They do not need to be installed separately in the Red Hat OpenShift cluster.

Kafka and AMQ streams

Apache Kafka is an event-streaming platform that collects, stores, manipulates, and distributes messages between processes.

Red Hat AMQ Streams is a Kafka distribution for the Red Hat OpenShift container platform. It provides an event-streaming backbone with which the application agent and data movers exchange data with high throughput and low latency.

Figure 4-2 shows the pods that are part of the BaaS project installation.

Project: baas ▾							
Pods							
Filter ▾	Name ▾	<input type="text" value="Search by name."/>					
Name ↑	Status ▾	Ready ▾	Restarts ▾	Owner ▾	Memory ▾	CPU ▾	
amq-streams-cluster-operator-v2.2.0-l-547978bf59-2muhf	Running	1/1	0	amq-streams-cluster-operator-v2.2.0-l-547978bf59	318,2 MiB	0,019 cores	
baas-entity-operator-674d845745-jt25g	Running	3/3	0	baas-entity-operator-674d845745	495,7 MiB	0,012 cores	
baas-kafka-0	Running	1/1	0	baas-kafka	715,5 MiB	0,013 cores	
baas-minio-0	Running	1/1	0	baas-minio	168,3 MiB	0,004 cores	
baas-scheduler-bddc8bb6f-2pwn9	Running	1/1	0	baas-scheduler-bddc8bb6f	16,2 MiB	0,001 cores	
baas-spp-agent-5c8b6c5ddb-vz4dg	Running	1/1	0	baas-spp-agent-5c8b6c5ddb	416,4 MiB	0,001 cores	
baas-transaction-manager-655d589bc8-kc4cp	Running	3/3	0	baas-transaction-manager-655d589bc8	412,6 MiB	0,008 cores	
baas-transaction-manager-655d589bc8-tcwn7	Running	3/3	2	baas-transaction-manager-655d589bc8	404,6 MiB	0,008 cores	
baas-transaction-manager-655d589bc8-vcwdw	Running	3/3	0	baas-transaction-manager-655d589bc8	416,0 MiB	0,007 cores	
baas-zookeeper-0	Running	1/1	0	baas-zookeeper	636,7 MiB	0,009 cores	
baas-zookeeper-1	Running	1/1	0	baas-zookeeper	617,9 MiB	0,007 cores	
baas-zookeeper-2	Running	1/1	0	baas-zookeeper	385,8 MiB	0,008 cores	
bmsspc-operator-controller-manager-9c655855c-nvbcf	Running	2/2	0	bmsspc-operator-controller-manager-9c655855c	1,408,2 MiB	1,160 cores	
openshift-adp-controller-manager-59dd87cf96-pbf77	Running	1/1	0	openshift-adp-controller-manager-59dd87cf96	40,4 MiB	0,004 cores	
velero-b87f65b7-wrrvj	Running	1/1	0	velero-b87f65b7	86,1 MiB	0,006 cores	

Figure 4-2 Pods that are installed as part of the BaaS project

4.2.2 Data mover pods

Data mover containers are controlled by the IBM Spectrum Protect Plus server and are used to move copies of local snapshot backup data to external storage; for example, a vSnap server or IBM Cloud Object Storage. The following types of data movers are available and are deployed with IBM Spectrum Protect Plus Container Backup Support:

- ▶ The first type of data mover is used to create external copies of PVC data by mounting a clone of the latest CSI snapshot. It is instantiated on-demand (when a backup job runs) and deployed in the application namespaces because access to snapshots and PVCs is possible from within a namespace only. After a PVC backup is completed, the data mover pods are removed again to save resources.
- ▶ The second type of data mover is used to create external copies of data that were created by the Velero tool (“metadata” backups). Velero stores its backups to a MinIO pod that runs inside the BaaS namespace and uses a small PVC to persist these backups. Thus, this second data mover type mounts a clone of the “MinIO-PVC” and copies the contents to external storage (vSnap or IBM Cloud Object Storage). As with the first data mover type, it is instantiated on-demand (when a backup job runs) and is removed again after the metadata backup completes.

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 policy. The data mover copies the data and IBM Spectrum Protect Plus creates a record in its recovery database. Snapshots that are created by on-demand backups also are recorded in IBM Spectrum Protect Plus.

4.2.3 Container Storage Interface

Stateful data must be managed across any containerized environment. For this need, many drivers with different specifications were available. The industry recognized the lack of standardization and developed a specification that is called the CSI, which is now adopted by all major container orchestrator systems.

For more information about how CSI integrates into a Red Hat OpenShift cluster, see [IBM Documentation](#).

The CSI enables Red Hat OpenShift Container Platform to use storage from any storage backend that implements the CSI interface as persistent storage. Multiple storage backends can be attached to the same Red Hat OpenShift cluster in parallel.

By using standardized API calls, applications can perform complex storage actions, such as dynamic provisioning of volumes, creating snapshots, or volume cloning without the need to understand how these activities are technically implemented on a specific storage back end. The CSI driver hides this complexity from the calling application.

IBM Spectrum Protect Plus Container Backup Support mainly uses the following CSI features:

- ▶ Provisioning of volumes: New volumes might be created as part of a restore procedure; for example, when the backup of a PVC is restored to another namespace or cluster. The type of volume is being specified by a StorageClass parameter of the CSI driver.
- ▶ Volume snapshots: Creating volume snapshots is a part of the PVC backup process. Users can decide whether creating regular, local snapshots is sufficient or if external copies of volume data must be created. To enable the creation of volume snapshots, the CSI driver that provided the volume must also implement a VolumeSnapshotClass.

- **Create volume from snapshot:** Whenever disaster protection is required and backup data must be copied to an external storage device (for example, vSnap or IBM Cloud Object Storage), volumes must be attached to data mover pods. Because snapshots cannot be mounted to a pod directly, temporary volumes are created from the latest snapshot and mapped to the data movers.

Examples

Theoretically, IBM Spectrum Protect Plus Container Backup Support works with any type of storage vendor that includes a CSI driver that implements the features as listed in 4.2.3, “Container Storage Interface” on page 66. However, IBM supports only CSI drivers and versions that were tested.

For more information about validated storage providers, see 4.3.3, “Supported storage types” on page 70 or this IBM Support [web page](#).

Types of supported container storage types include the following examples:

- **IBM Block CSI:** This storage plug-in enables a Red Hat OpenShift cluster to dynamically provision volumes on IBM block storage systems that are attached to the worker nodes by way of ISCSI or Fibre Channel, such as the IBM DS8000® family, IBM FlashSystem® A9000/A9000R, and systems from the IBM Spectrum Virtualize family.

Each persistent volume in Red Hat OpenShift features a 1:1 relation to a LUN in the storage system (for example, a VDisk in an IBM FlashSystem). The creation of a volume snapshot in Red Hat OpenShift results in a hardware snapshot (IBM FlashCopy®) on the storage system.

- **IBM Spectrum Scale CSI:** This storage plug-in enables a Red Hat OpenShift cluster to dynamically provision file volumes that are in independent file sets of an IBM Spectrum Scale file system. Each persistent volume in Red Hat OpenShift features a 1:1 relation to an independent file set in IBM Spectrum Scale. The creation of a volume snapshot in Red Hat OpenShift results in a software snapshot of the corresponding file set.

Important: The latest versions of IBM Spectrum Scale CSI support so-called “consistency groups” (CGs). In such a setup, the PVCs are represented by dependent file sets that belong to an independent file set (which represents the CG). If a CSI snapshot is performed for only one of the PVCs, the result is a software snapshot of the IBM Spectrum Scale independent file set, which also affects or snapshots all of the other PVCs in the CG. For more information, see this IBM Documentation [web page](#).

Figure 4-3 shows the relation between persistent volumes and file sets in IBM Spectrum Scale consistency groups.

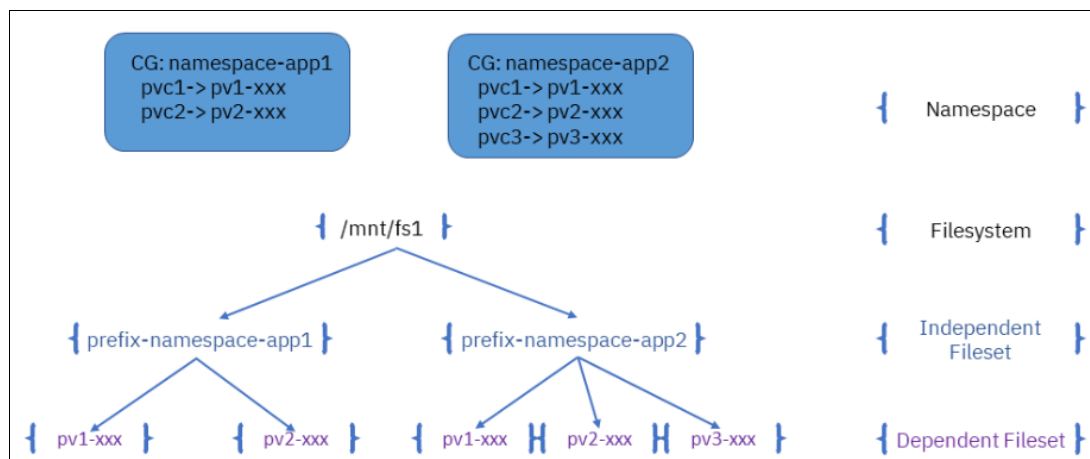


Figure 4-3 IBM Spectrum Scale consistency groups

4.3 Red Hat OpenShift prerequisites and supported environments

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 more information about the current support matrix, see this IBM Support [web page](#).

4.3.1 Installation types

The following types of installations are available:

- Online

Install the IBM Spectrum Protect Plus operator from an online registry, such as the IBM Entitled Registry. By creating an image pull secret, the IBM Spectrum Protect Plus image can be installed from the IBM Cloud Container Registry (cp.icr.io/cp).

For more information, see Chapter 5, “Implementing Container Backup Support” on page 85 or this IBM Documentation [web page](#).

- Air-gapped installation

Some Red Hat OpenShift clusters cannot be attached to the internet and must have a private registry to install the IBM Spectrum Protect Plus operator. This installation is a bit more tricky because more components must be in place. The installation can be done by using a Container Application Software for Enterprises (CASE) package.

For more information, see this IBM Documentation [web page](#).

4.3.2 Supported Red Hat OpenShift architecture

IBM Spectrum Protect Plus data protection for containers can be implemented on Red Hat OpenShift that is running on AMD64 and IBM zSystems platform.

Different versions of cloud environments also are supported for the AMD64 platform. The IBM zSystems platform is supported only in private cloud deployments (on-premised). For more information about the coverage support matrix, see this IBM Support [web page](#).

Notes: Red Hat OpenShift for IBM zSystems is *not* supported in an air-gapped environment. In addition, the data that is protected on IBM zSystems environments cannot be recovered on amd64-based Red Hat OpenShift Container Platform clusters.

The storage type support varies between the AMD64 and IBM zSystems platform. For more information about storage support, see Chapter 5, “Implementing Container Backup Support” on page 85.

At the time of writing, the Red Hat OpenShift environments that are listed in Table 4-2 - Table 4-5 are supported by IBM Spectrum Protect Plus.

Table 4-2 Supported Red Hat OpenShift Container Platform versions (AMD64)

Red Hat OpenShift Container Platform	IBM Spectrum Protect Plus
<ul style="list-style-type: none">▶ Version 4.8 and later updates▶ Version 4.9 and later updates▶ Version 4.10 and later updates▶ Version 4.11 and later updates	<ul style="list-style-type: none">▶ Beginning with Version 10.1.8 iFix2▶ Beginning with Version 10.1.9 iFix3▶ Beginning with Version 10.1.10.2▶ Beginning with Version 10.1.11

Table 4-3 Supported ODF versions (AMD64)

Red Hat OpenShift Data Foundation	IBM Spectrum Protect Plus
<ul style="list-style-type: none">▶ Version 4.8 and later updates▶ Version 4.9 and later updates▶ Version 4.10 and later updates	<ul style="list-style-type: none">▶ Beginning with Version 10.1.8 iFix2▶ Beginning with Version 10.1.9 iFix3▶ Beginning with Version 10.1.11

Table 4-4 Supported Red Hat OpenShift Container Platform versions (IBM zSystems)

Red Hat OpenShift Container Platform	IBM Spectrum Protect Plus
<ul style="list-style-type: none">▶ Version 4.9 and later updates▶ Version 4.10 and later updates▶ Version 4.11 and later updates	<ul style="list-style-type: none">▶ Beginning with Version 10.1.11▶ Beginning with Version 10.1.12▶ Beginning with Version 10.1.12

Table 4-5 Supported ODF versions (IBM zSystems)

Red Hat OpenShift Data Foundation	IBM Spectrum Protect Plus
<ul style="list-style-type: none">▶ Version 4.9 and later updates▶ Version 4.10 and later updates	<ul style="list-style-type: none">▶ Beginning with Version 10.1.11▶ Beginning with Version 10.1.12

4.3.3 Supported storage types

Container Backup Support protects volume data that was allocated and provisioned by a storage plug-in that supports the CSI specification.

The CSI plug-in must provide the following features:

- Dynamic provisioning of volumes.
- Snapshot capabilities that are used for local backup operations.
- Volume-cloning capabilities (“volume from snapshot”) that are used to mount a volume clone to the IBM Spectrum Protect Plus data mover component to create external backup copies; for example, to an IBM Spectrum Protect Plus vSnap server or to Object Storage.

Notes: Back-up operations for raw block device volumes (volumeMode 'Block') are not supported. Only formatted volumes can be mounted to the data mover for copy operations.

You cannot restore a snapshot backup to a different cluster or namespace because a snapshot belongs to the same namespace as the source volume.

File-based, incremental copy backup and restore operations are performed by using the open-source tool restic, which is included in the data mover pods. During incremental backups, only new and changed data is copied to the IBM Spectrum Protect Plus vSnap server. It does not matter if block or file-based storage types are in use.

For IBM Spectrum Scale backups, the following limitations apply:

- Snapshots can be created from independent fileset-based PVCs only. 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.
- Backup and restore operations of IBM Spectrum Scale specific access control lists (ACLs) are *not* supported. During a restore operation, only standard and extended POSIX attributes are restored.

At the time of this writing, the storage types that are listed in Table 4-6 and Table 4-7 on page 71 are supported by IBM Spectrum Protect Plus.

Table 4-6 Supported storage types (AMD64)

Storage	CSI driver	CSI and IBM Spectrum Protect Plus version
External CephFS	Ceph CSI driver with CephFS storage	<ul style="list-style-type: none">► Installed with OCS, or Version 3.2.2 or later (Beginning with IBM Spectrum Protect Plus 10.1.10)► Version 3.6.2 or later (Beginning with IBM Spectrum Protect Plus 10.1.12)

Storage	CSI driver	CSI and IBM Spectrum Protect Plus version
Ceph Rados Block Device (RBD)	Ceph CSI driver with RBD	<ul style="list-style-type: none"> ▶ Installed with OCS, or Version 3.2.2 or later (Beginning with IBM Spectrum Protect Plus 10.1.10) ▶ Version 3.6.2 or later (Beginning with IBM Spectrum Protect Plus 10.1.12)
IBM block storage	IBM block storage CSI for virtualized storage	<ul style="list-style-type: none"> ▶ Version 1.6 or later (Beginning with IBM Spectrum Protect Plus 10.1.8 ifix2) ▶ Version 1.7 or later (Beginning with IBM Spectrum Protect Plus 10.1.9) ▶ Version 1.8 or later (Beginning with IBM Spectrum Protect Plus 10.1.10) ▶ Version 1.9 or later (Beginning with IBM Spectrum Protect Plus 10.1.11)
IBM Spectrum Scale (5.1.1 or later)	IBM Spectrum Scale CSI driver	Version 2.2.0 or later updates (Beginning with IBM Spectrum Protect Plus 10.1.8)
HNAS	HNAS CSI Driver for Kubernetes	Version 1.1.1 or later (Beginning with IBM Spectrum Protect Plus 10.1.9)
NetApp storage	CSI Trident for Kubernetes	Version 21 or later (Beginning with IBM Spectrum Protect Plus 10.1.9)

Table 4-7 Supported storage types (IBM zSystems)

Storage	CSI driver	CSI and IBM Spectrum Protect Plus version
External CephFS	Ceph CSI driver with CephFS storage	Version 3.2.2 or later (Beginning with IBM Spectrum Protect Plus 10.1.12)
Ceph RBD	Ceph CSI driver with RBD	Version 3.2.2 or later (Beginning with IBM Spectrum Protect Plus 10.1.12)
IBM block storage	IBM block storage CSI for virtualized storage	Version 1.8 or later (Beginning with IBM Spectrum Protect Plus 10.1.11)

For more information about supported storage types, see this IBM Support [web page](#).

4.3.4 Red Hat OpenShift cluster prerequisites

Specific requirements must be met to ensure the proper operation of Container Backup Support on your cluster. These requirements consist of the CLI, collecting metrics performance improvements, CSI external-snapshotter, an accessible registry, a storage class, and defining user permissions and settings.

Consider the following points:

- ▶ In a Red Hat OpenShift environment, the CLI (`oc`) must be accessible on the installation host and in the local path.
- ▶ Ensure that the Red Hat OpenShift Cluster Monitoring Operator is installed and running in the environment for metrics and performance improvement.
- ▶ Ensure that the cluster operator `csi-snapshot-controller` is in the Available: True state.
- ▶ A storage class and volume snapshot class must be defined for the persistent volumes that are being protected.
- ▶ The target image registry must be accessible from the Red Hat 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 use a `kubeconfig` file with cluster-admin privileges, `KUBECONFIG`.
- ▶ To create 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.

4.4 IBM Spectrum Protect Plus and storage prerequisites

In this section, we discuss IBM Spectrum Protect Plus and storage prerequisites.

4.4.1 General 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. Consider the following points:

- ▶ An IBM Spectrum Protect Plus instance must be deployed in a container environment or as a VMware or Microsoft Hyper-V virtual appliance.
- ▶ Network connectivity must exist to and from the target Red Hat OpenShift cluster.
- ▶ The Red Hat OpenShift cluster must be registered to IBM Spectrum Protect Plus by using a Fully Qualified Domain Name (FQDN) name or IP address.
- ▶ The IBM Spectrum Protect Plus IP address and port number must be specified in the Container Backup Support configuration. Only https (port 443) can be specified for use with all IBM Spectrum Protect Plus instances.

An example of a BaaS configuration file is shown in Figure 4-4 on page 73.

```

1  apiVersion: sppc.ibm.com/v1
2  kind: IBMSPPC
3  metadata:
4    creationTimestamp: '2022-10-14T12:06:58Z'
5    generation: 1
6    managedFields: ...
61  name: ibmsppc-entitled-registry
62  namespace: baas
63  resourceVersion: '19301631'
64  uid: fc9840c6-cb8a-4cfd-8b83-525154345969
65  spec:
66    networkPolicy:
67      cluster_api_serverips:
68        - 10.0.240.232
69        - 10.0.240.233
70        - 10.0.240.234
71      cluster_api_serverport: 6443
72      cluster_cidr: 10.254.0.0/16
73      is_server_installed_on_another_cluster: false
74      other_cluster_cidr_block: '24'
75      name_override: baas
76      spp_ips: 10.0.240.222
77    license:
78      accept: true
79      spp_fqdn: 10.0.240.222
80      is_icp: false
81      velero_namespace: ''
82      image_registry: cp.icr.io/cp
83      minio_password: ''
84      product_locale: en_US
85      minio_storage_class: ocs-storagecluster-ceph-rbd
86      spp_port: 443
87      image_registry_secret: baas-registry-secret
88      spp_agent_service_node_port: 30001
89      product_loglevel: DEBUG
90      cluster_name: dpr
91      state: present
92      fullname_override: baas
93      is_scale_node_selector: false
94      max_parallel_snapshots: 20
95      version_override: ''
96      arch: amd64
97      managed_by: ansible

```

Figure 4-4 Sample YAML file with IBM Spectrum Protect Plus IP address and port number

- An administrative user account for Container Backup Support must be configured on IBM Spectrum Protect Plus for access to all external components that interact with Container Backup Support. This administrative account can be a local user or configured as a global LDAP account. Figure 4-5 shows the user management dialog of IBM Spectrum Protect Plus.

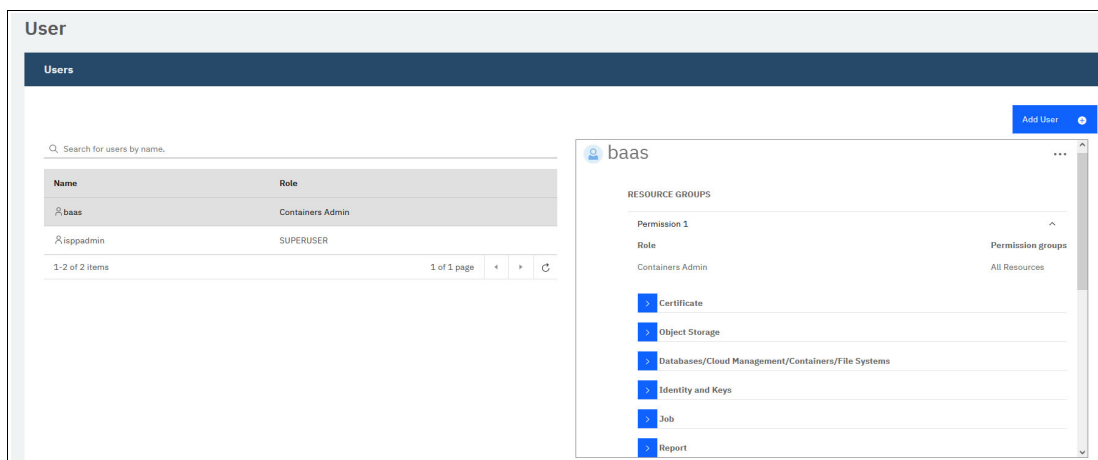


Figure 4-5 BaaS user with the Container Admin role

4.4.2 Managing connection performance in Red Hat OpenShift

A default of two Ingress controller pods is needed to handle routes. Larger Red Hat OpenShift clusters require more Ingress controller pods to handle routes. When the number of Ingress controller pods is insufficient, you might experience connections being dropped from the IBM Spectrum Protect Plus to the agent. The maximum number of Ingress controller pods that you can scale is equal to the number of worker nodes.

Example 4-1 shows how to determine the current number of ingress controllers, update the number to a higher value, and check that the update was successful.

Example 4-1 Modifying the number of Red Hat OpenShift Ingress controllers

```
#> oc get -n openshift-ingress-operator ingresscontrollers/default -o
jsonpath='{$.status.availableReplicas}'
2

#> oc patch -n openshift-ingress-operator ingresscontroller/default --patch
'{"spec":{"replicas": 3}}' --type=merge
ingresscontroller.operator.openshift.io/default patched

#> oc get -n openshift-ingress-operator ingresscontrollers/default -o
jsonpath='{$.status.availableReplicas}'
3
```

Consider the following points:

- ▶ Network connectivity must exist to and from the target Red Hat OpenShift cluster and the IBM Spectrum Protect Plus vSnap instance or S3 object storage.
- ▶ If backups are encrypted at rest, make sure to size suitable resources (memory and CPU) to the vSnap server. Encryption uses 5 - 10% extra CPU resources, depending on the environment.
- ▶ Cluster performance is affected when you back up many PVCs concurrently because each PVC uses multiple threads and uses bandwidth when data is copied. You can specify the maximum concurrent PVCs. Figure 4-6 shows the default concurrent PVCs of 10.

The screenshot displays the 'Manage Clusters' interface. Under the 'Edit Application Properties' section, the 'Use existing user' checkbox is checked, and the 'Select user' dropdown is set to 'dpr'. In the 'Certificate' section, 'Use existing certificate' is selected, and the 'Select a certificate' dropdown is set to 'dpr-cert'. The 'Options' section, highlighted with a red box, contains a slider for 'Maximum concurrent PVCs' which is currently set to 10, with minus and plus buttons for adjustment.

Figure 4-6 Specification of maximum concurrent PVCs

4.4.3 Vertical Pod Autoscaler

The Red Hat OpenShift Container Platform Vertical Pod Autoscaler Operator (VPA) automatically reviews the historic and current CPU and memory resources for containers in pods. It also can update the resource limits and requests based on the usage values it learns over time.

The use of the VPA is optional. If VPA should be used to adapt resource assignments for the BaaS pods, the following steps must be performed before IBM Spectrum Protect Plus Container Backup Support is installed:

1. Install the VPA to the Red Hat OpenShift cluster by using a Red Hat operator. For more information, see this [web page](#).
2. Adapt the BaaS configuration file to use VPA, as described at this [web page](#).
3. Deploy BaaS, as described in Chapter 5, “Implementing Container Backup Support” on page 85.

The VPA parameters that are added to the BaaS configuration enable choosing between two “update modes”:

- ▶ Initial: The VPA adapts the resource assignments to BaaS pods whenever a pod is restarted or evicted.
- ▶ Off: Recommendations can be queried from VPA and applied to the BaaS configuration manually.

4.4.4 Cloud storage for direct backup operations

Object storage can be used in the following ways to store container backup data with IBM Spectrum Protect Plus:

- ▶ Use disk-based vSnap servers as primary target storage for container backup data and create copies from vSnap to an object storage system regularly.
- ▶ Back up data from a Red Hat OpenShift cluster directly to S3 object storage.

Notes: In addition to container backup data, the metadata catalogs of the IBM Spectrum Protect Plus server can be stored to Object Storage by using these approaches.

When direct backup to object storage is used, no other copies (for example, to a vSnap or another object storage) are possible.

The following Object Storage systems are supported for direct backup from Red Hat OpenShift:

- ▶ Amazon Simple Storage Service (Amazon S3)
- ▶ IBM Cloud Object Storage (including IBM Cloud Object Storage Systems)
- ▶ Microsoft Azure Blob storage
- ▶ S3 compatible storage

Note: For IBM Cloud Object Storage, retention-enabled vaults are not supported.

For S3 compatible storage, generic S3 support is based on an external certification process. For more information about supported S3 compatible providers, see this [IBM Support web page](#).

Other object storage (such as MinIO or other open-source implementations) might work, but were not tested and are not supported by IBM.

Figure 4-7 shows an example of an IBM Cloud Object Storage System being registered as cloud storage to IBM Spectrum Protect Plus.

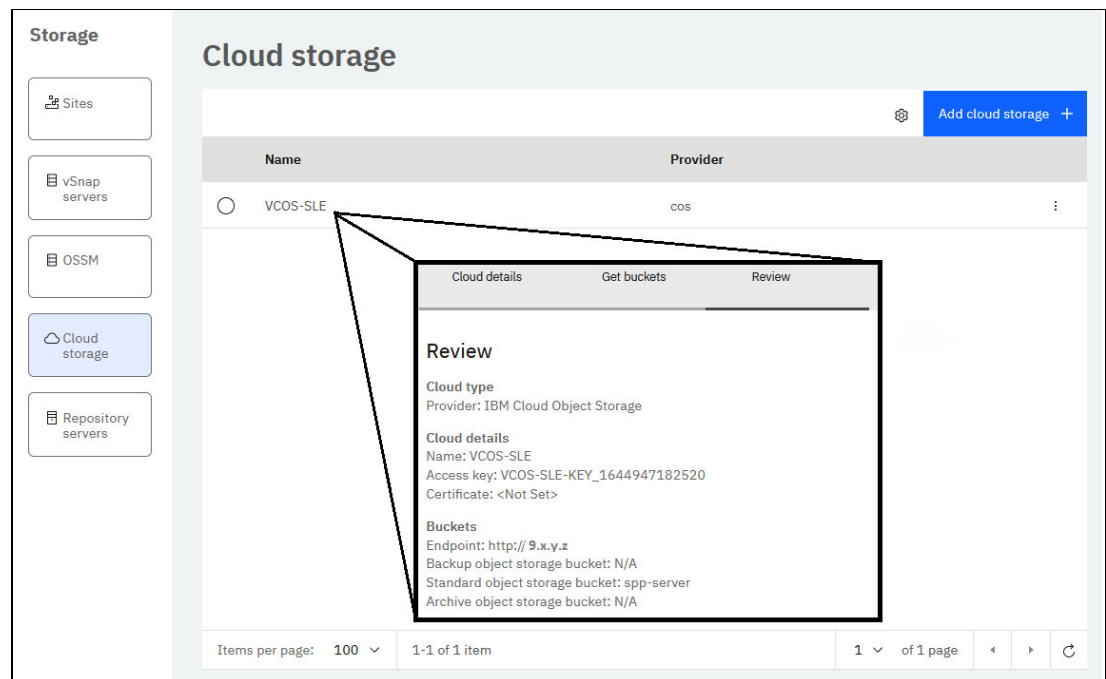


Figure 4-7 Example of defining an IBM Cloud Object Storage System as a backup target

4.4.5 Communication ports

Table 4-8 lists the ports that are used by IBM Spectrum Protect Plus agents.

Table 4-8 Communication ports that are used for Container Backup Support

Port	Protocol	Initiator	Target	Description
443	TCP	IBM Spectrum Protect Plus server	Container Backup Support agent	Used by IBM Spectrum Protect Plus to connect to the data mover container to run agents. Also used for REST API connections to the container backup support agent.
111	TCP and UDP	Container Backup Support agent	vSnap server	Used for NFS data transfer to and from file systems that are mounted from vSnap servers during backup and restore operations.

Port	Protocol	Initiator	Target	Description
443	TCP	Container Backup Support agent	IBM Spectrum Protect Plus server	Used for IBM Spectrum Protect Plus-issued commands to run backup, restore, inventory, and other operations.
2049	TCP and UDP	Container Backup Support agent	vSnap server	Used for NFS data transfer to and from file systems that are mounted from vSnap servers during backup and restore operations.
20048	TCP and UDP	Container Backup Support agent	vSnap server	Used for NFS data transfer to and from file systems that are mounted from vSnap servers during backup and restore operations.

4.5 Container backup and restore types

Container Backup Support provides multiple types of backup and restore functions for your PVCs and other cluster resources. You can use the IBM Spectrum Protect Plus user interface to assign policies and regular backup schedules to cluster components or start ad hoc backup and restore operations.

4.5.1 Backup types

In this section, we describe the available backup operations.

For more information about backup and restore types, see this IBM Documentation [web page](#).

Figure 4-8 shows the data flows for snapshot and copy backups.

Note: Figure 4-8 shows a vSnap server as target for external data copies. Instead of a vSnap server, data also can be backed up to IBM Cloud Object Storage.

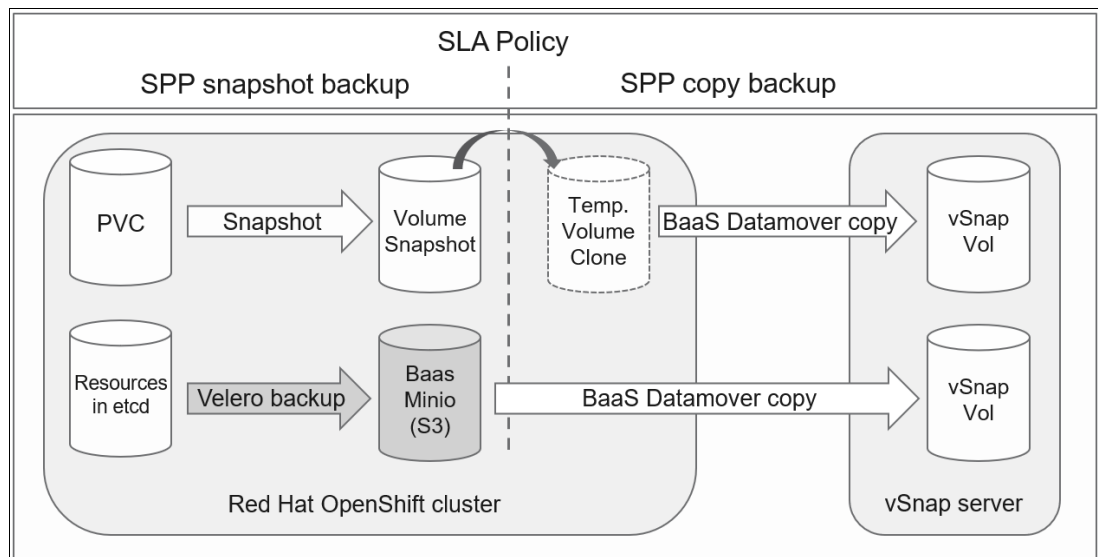


Figure 4-8 Snapshot versus copy backup

The following sections describe the backup options.

Snapshot backup of PVCs

A snapshot backup creates a backup of the persistent volume by using 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 backed up.

The snapshot class must be compatible with the storage class of the persistent volume. That is, the snapshot class and storage class are defined and provided by the same CSI storage plug-in.

Snapshot backup: namespace-scoped and cluster-scoped resources

A snapshot backup creates a backup of the cluster-scoped or namespace-scoped resources, which are backed up to the local MinIO in the 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 or on-demand backup requests. During scheduled backups, snapshot backups are created at intervals that are defined by the assigned SLA policy.

Copy backup: PVCs

This type of backup copies the persistent volume contents to an IBM Spectrum Protect Plus vSnap server or directly to IBM Cloud Object Storage. IBM Spectrum Protect Plus policy definitions allow dissimilar retentions between copy and snapshot backups.

Copy backup: namespace-scoped and cluster-scoped resources

A copy backup copies the namespace-scoped or cluster-scoped resources from BaaS MinIO to an IBM Spectrum Protect Plus vSnap server or directly to IBM Cloud Object Storage. IBM Spectrum Protect Plus policy definitions allow dissimilar retentions between copy and snapshot backups.

Snapshot-only protection

Container Backup Support can be deployed as a snapshot-only solution. In this configuration, the installation of an IBM Spectrum Protect Plus vSnap server or a connection to IBM Cloud Object Storage is not required.

IBM Spectrum Protect Plus policy definitions may define snapshot backups only, without enabling regular copy backups. In this scenario, snapshots are created and stored locally, but data is not copied from the cluster to the external vSnap server or cloud storage. With a snapshot-only deployment, data cannot be restored to another namespace or cluster.

4.5.2 Restore types

In this section, we describe the available types of restore operations.

For more information about restoring container data with IBM Spectrum Protect Plus, see this IBM Documentation [web page](#).

Figure 4-9 on page 79 shows the data flows for snapshot and copy restores.

Note: Figure 4-9 on page 79 shows a vSnap server as source of external data copies. Instead of a vSnap server, data also can be restored from IBM Cloud Object Storage.

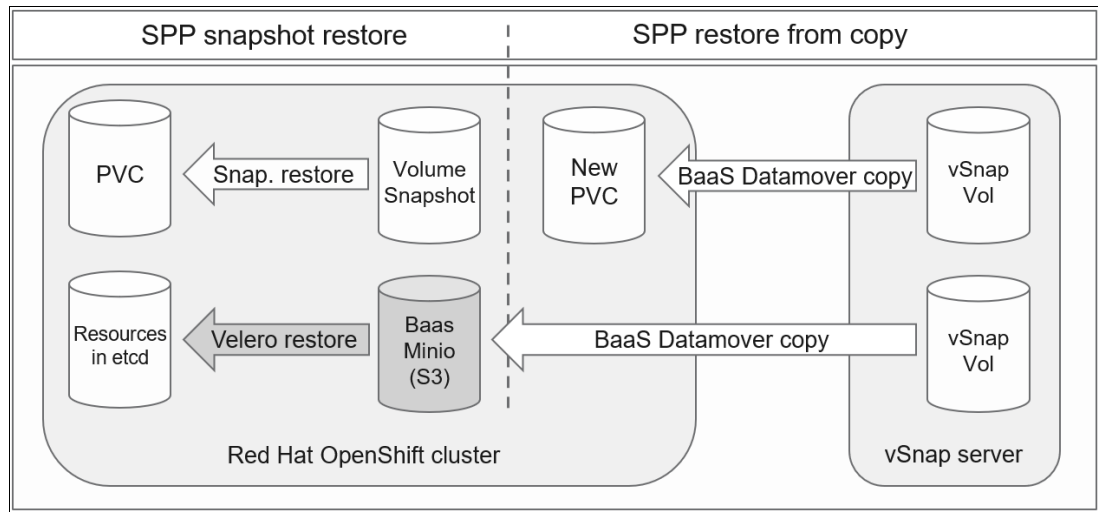


Figure 4-9 Snapshot versus copy restores

The following sections describe the restore options.

Snapshot restore: PVCs

This operation restores a snapshot to a new persistent volume. This type of operation is suitable for rapidly restoring recent snapshot backups.

Snapshot restore: namespace-scoped and cluster-scoped resources

This type of snapshot restore operation restores a snapshot copy from the BaaS MinIO to the same namespace or cluster. Snapshot copy backups cannot be restored to another namespace or cluster; they can be restored to the same cluster or the same namespace only. This type of operation is suitable for rapidly restoring recent snapshot backups.

Copy backup restore: PVCs

This operation restores a copy backup to a persistent volume with the same name of the original volume, or to a new volume. You can select to restore to an alternative 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: namespace-scoped and cluster-scoped resources

This operation restores a backup from the vSnap server or cloud storage 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 is being specified 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.

Restore capability matrix

Table 4-9 lists the four different restore methods in combination with the backup methods.

Table 4-9 Backup and restore methods

Restore to	Restore from Snapshot	Restore from Copy (vSnap or cloud storage)
Original PVC name in same namespace (original PVC must be deleted first)	Y	Y
New PVC name in same namespace (deployments must be updated to include new PVC name)	Y	Y
Other namespace in same cluster	N	Y
Other cluster (both clusters must have same storage type; for example, IBM block CSI) Note: Storage class names can be different	N	Y

4.6 Service-level agreement policies for container backup

SLAs in IBM Spectrum Protect Plus are policies that define the following relevant parameters for a backup:

- ▶ Scheduler: Can be Active or Disabled (ad hoc or external scheduling only).
- ▶ Retention: How long to keep the backup data?
- ▶ Repeats: When and how often is data to be backed up?
- ▶ Target: Where to store the backup data?
- ▶ Replication: Will the data be replicated to another vSnap server?
- ▶ Extra copies: Enables creating copies of backup data that is stored in a vSnap to other storage, such as IBM Spectrum Protect or IBM Cloud Object Storage.

Note: Replication data and extra copies can have their own retentions assigned.

Depending on the type of the backup client (for example, VMs, file systems, databases or applications, Amazon EC2 instances, Kubernetes, or Red Hat OpenShift containers) some SLA options can vary or might not be available at all.

An SLA for Red Hat OpenShift container backup is referred to a *tiered snapshot*. Figure 4-10 on page 81 shows the options of the tiered snapshot SLA type.

Tiered snapshot

This policy includes these functions:

- Snapshot
- Backup to vSnap
- Replication
- Create additional data copies to other locations

These workloads can use this policy:

- Kubernetes
- OpenShift

Figure 4-10 SLA type tiered snapshot for container backup support

The following functions are available:

- ▶ Snapshot Protection (a container backup policy must use CSI snapshots; all other SLA parameters are optional):
 - Scheduler: Active or inactive.
 - Retention for local CSI snapshot backups.
 - Repeats: When and how often will CSI snapshots be created?
 - Snapshot Prefix: A prefix that is assigned to the CSI snapshots that allows identification of snapshots in Red Hat OpenShift Container Platform that were triggered by IBM Spectrum Protect Plus.
- ▶ Backup Policy (optional, snapshots can be “tiered” or copied to external vSnap or cloud storage):
 - Scheduler: Active or inactive.
 - Retention for vSnap or cloud storage backups.
 - Repeats: When and how often will snapshots be tiered to vSnap or cloud?
 - Storage Type: Site (vSnap) or cloud storage.
- ▶ Replication Policy (optional, can be used only if the older SLA section Backup Policy is enabled and configured to copy data to a vSnap. Replication is not possible when the Backup Policy sends snapshot data to directly cloud storage):
 - Scheduler: Active or inactive.
 - Retention for replicated vSnap data.
 - Repeats: When and how often will vSnap data be replicated?
 - Target Site: Specifies which vSnaps will be used as replication targets.

- ▶ Extra Copies (optional, can be used only if older SLA section Backup Policy is enabled and configured to copy data to a vSnap. Extra copies are not possible when the Backup Policy sends snapshot data directly to cloud storage):
 - Scheduler: Active or inactive.
 - Retention for copied vSnap data.
 - Repeats: When and how often will vSnap data be copied?
 - Source: Data can be copied from a backup-target vSnap or replication-target vSnap.
 - Destination: Data can be copied to IBM Spectrum Protect or cloud storage.

The storage administrator can create SLA policies by using the IBM Spectrum Protect Plus user interface:

- ▶ List existing SLAs: 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 and cluster- and namespace-scoped resources. The predefined 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

- ▶ Create new SLA: In the IBM Spectrum Protect Plus user interface, click **Manage Protection** → **Policy Overview** → **Add SLA Policy**. Multiple options for SLA types are displayed.

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. Then, they are deleted by IBM Spectrum Protect Plus maintenance jobs. It is also possible to expire container data backups manually.

For more information, see this IBM Documentation [web page](#).

4.7 Container Backup Support security features

The following advanced security features are provided to help protect containers, secure network connections, encrypt data, and verify installation packages:

- ▶ Security scanning of containers
- ▶ Least privileged containers
- ▶ Authentication of network connections
- ▶ Multitenancy
- ▶ Encryption of data at rest
- ▶ Code signing

For more information about these security features, see this IBM Documentation [web page](#).

4.8 IBM Spectrum Fusion and IBM Cloud Pak for Data

IBM Spectrum Fusion is a data services platform for Red Hat OpenShift. It builds the foundation for container-native applications that are running anywhere on Red Hat OpenShift (on-premises or in a cloud) and provides data-storage and data-protection services.

Key components are based on IBM Spectrum Scale to provide a scalable, flexible container-native storage and IBM Spectrum Protect Plus to provide the container backup capabilities.

IBM Cloud Pak for Data (CP4D) is a cloud-native solution that runs on Red Hat OpenShift and that enables users to get out the most of their data by providing tools and features, such as the following examples:

- ▶ Collect from or connect to various data sources.
- ▶ Organize data by using data classifications.
- ▶ Storing data into various databases or warehouses.
- ▶ Infuse data with AI and machine learning.
- ▶ Analyze and visualize data to get insights and to extract business relevant information.

Depending on which components are implemented and used, protecting a CP4D installation can become challenging. To ensure backup data consistency, some services might need to be stopped during backup or components might need to be backed up in a specific order.

IBM Spectrum Fusion is the only data protection platform that can create a nondisruptive backup of CP4D for DR by orchestrating the backup of multiple data sources by using a prescribed set of steps. This prescription is called a *recipe*, which is provided by the application to be backed up (CP4D) and used by the data protection platform (IBM Spectrum Protect Plus) to perform the backup in the correct way.

At the time of this writing, CP4D is the only Cloud Pak that provides a backup recipe and the containerized version of IBM Spectrum Protect Plus that is packaged with IBM Spectrum Fusion. It also is the only backup solution that can use the recipe to control the backup workflow.



Implementing Container Backup Support

This chapter provides information about how to install, configure, and use the IBM Spectrum Protect Plus and backup as a service (BaaS) component in a Red Hat OpenShift Container environment.

This chapter includes the following topics:

- ▶ 5.1, “Validating the prerequisites” on page 86
- ▶ 5.2, “Installing in an online environment” on page 87
- ▶ 5.3, “Installing in an air-gapped environment” on page 120

5.1 Validating the prerequisites

IBM Spectrum Protect Plus Container Backup Support for Red Hat OpenShift is implemented through the BaaS component. This component acts as an agent for IBM Spectrum Protect Plus and is installed in the Red Hat OpenShift cluster.

Before starting the BaaS installation, the following prerequisites must be met:

- ▶ Your Red Hat OpenShift version is supported by IBM Spectrum Protect Plus.
- ▶ A Storage Provider and Container Storage Interface (CSI) driver is implemented that is supported by IBM Spectrum Protect Plus.
- ▶ A running IBM Spectrum Protect Plus server and optional vSnap servers are available.
- ▶ All necessary software packages are installed in your Red Hat OpenShift Container Platform cluster.

For more information about supported environments, see Chapter 4, “Container Backup Support” on page 61, or this IBM Documentation [web page](#).

5.1.1 Preparing IBM Spectrum Protect Plus backup servers

All backup and restore operations are managed by a central entity: the IBM Spectrum Protect Plus server. The IBM Spectrum Protect Plus vSnap servers act as a storage server for the backup data.

The IBM Spectrum Protect Plus server and the IBM Spectrum Protect Plus vSnap servers must be provisioned and configured by the IBM Spectrum Protect Plus administrator. Consider the following points:

- ▶ An IBM Spectrum Protect Plus server must be deployed in a container environment or as a VMware virtual appliance.

Network connectivity must exist to and from the target cluster and the IBM Spectrum Protect Plus server instance.

- ▶ An administrative account for Container Backup Support must be configured on IBM Spectrum Protect Plus, as shown in Figure 4-5 on page 73. This administrative account can be configured as a local user or as a global LDAP account in the data center. This administrative user account and its password must be available when installing the BaaS component later.

- ▶ Optional: For copy backup and copy restore operations, the IBM Spectrum Protect Plus vSnap servers must be provisioned and configured by the backup 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 Red Hat OpenShift cluster and the IBM Spectrum Protect Plus vSnap instance.

For more information about the IBM Spectrum Protect Plus architecture and components, see the Chapter 2, “IBM Spectrum Protect Plus architecture” on page 11, or *IBM Spectrum Protect Plus Practical Guidance for Deployment, Configuration, and Usage*, REDP-5532.

5.2 Installing in an online environment

The easiest way to install IBM Spectrum Protect Plus Container Backup support is by using IBM and Red Hat online repositories.

To perform a BaaS installation by using the online repositories, the Red Hat OpenShift cluster and the system that is used to pull the required code packages must access the internet. In environments where direct internet access is denied, an http or https proxy can be used.

For more information about how to configure cluster-wide proxy on an Red Hat OpenShift Container Platform cluster, see this Red Hat OpenShift Documentation [web page](#).

The process that is used to install Container Backup Support (BaaS) features the following overall tasks:

- ▶ Obtaining the IBM Container Software Library Entitlement key.
- ▶ Setting up installation variables.
- ▶ Editing the cluster global pull secret.
- ▶ Creating the Container Backup Support namespace.
- ▶ Creating the image pull secret.
- ▶ Creating the general-purpose secret for Container Backup Support.
- ▶ Exchanging certificate between IBM Spectrum Protect Plus and Red Hat OpenShift Container Platform.
- ▶ Adding the online CatalogSource.
- ▶ Installing the Container Backup Support (BaaS) operator.
- ▶ Creating a Container Backup Support (BaaS) instance.
- ▶ Registering the Red Hat OpenShift cluster in IBM Spectrum Protect Plus server.

Note: Unlike previous versions, the latest versions of the Container Backup Support (BaaS) operator deploys the Red Hat OpenShift API for Data Protection and the AMQ Streams operators (as shown in Figure 4-1 on page 64).

Also, Red Hat OpenShift API for Data Protection and AMQ Streams instances are deployed during the Container Backup Support (BaaS) installation.

5.2.1 Obtaining the IBM Container Software Library Entitlement key

During the deployment, container images are pulled from the IBM Container Software Library.

To access this library, you must retrieve your entitlement key from this [web page](#) (log in required).

After you retrieve your entitlement key, complete the following steps:

1. Log in to the IBM Container Software Library by using the IBMid and password that are associated with the entitled software.
2. Click **Get entitlement key**.

3. Click **Copy key** to copy the generated entitlement key (see Figure 5-1).

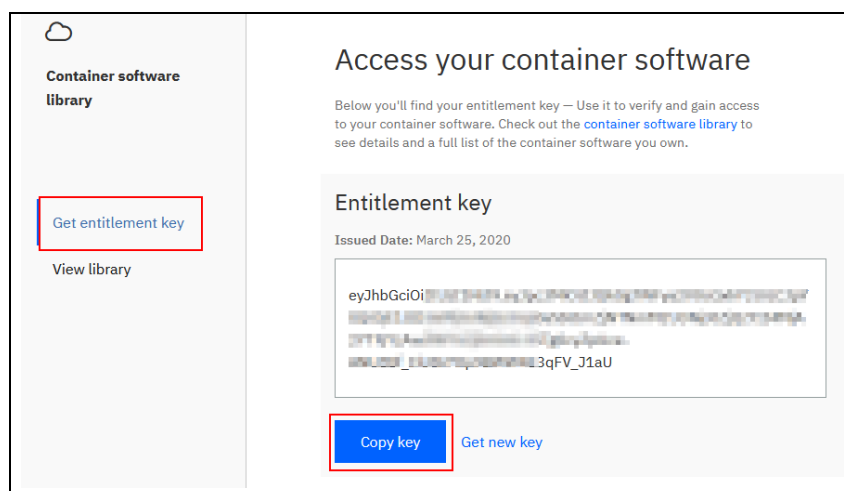


Figure 5-1 Obtaining the entitlement key from the IBM Container Software Library

4. Save the entitlement key securely because it is requested later.

5.2.2 Setting up installation variables

The goal of this process is to prepare the following files required during the installation:

- `baas-options.sh`

This file contains variables that are used to prepare the Container Backup Support installation. It is specially useful for command-line interface (CLI) installation and is used as a reference for GUI installation.

- `baas-values-cr.yaml`

This file describes the CustomResource that is used to deploy the Container Backup Support. It is specially useful for CLI installation and is used as a reference for GUI installation.

During our installation, those files are stored in the `install_vars` directory (see Example 5-1).

Example 5-1 Creating the `install_vars` directory

```
$ mkdir install_vars
$ cd install_vars
```

Preparing the `baas-options.sh` file

The goal of this process is to create the `baas-options.sh` file, which contains variables that are used to prepare the Container Backup Support installation.

Tip: A sample `baas-option.sh` file is available at this IBM Documentation [web page](#).

Although this file is used mostly for the CLI installation, it can be used as a reference for the GUI installation (see Example 5-2).

Example 5-2 Sample baas-option.sh that is used for the lab

```
$ cat baas-options.sh
export DOCKER_REGISTRY_ADDRESS='cp.icr.io/cp'
export DOCKER_REGISTRY_USERNAME='cp'
export DOCKER_REGISTRY_PASSWORD='IBM_CONTAINER_SOFTWARE_ENTITLEMENT_KEY'
export DOCKER_REGISTRY_NAMESPACE='sppc'
export SPP_ADMIN_USERNAME='baas'
export SPP_ADMIN_PASSWORD='Pass40CP!'
export DATAMOVER_USERNAME='dmuser'
export DATAMOVER_PASSWORD='pass2PLAY'
export MINIO_USERNAME='minio'
export MINIO_PASSWORD='pass2PLAY'
```

Use Table 5-1 to determine the suitable environment variables to set in your `baas-options.sh` file and then, update your `baas-options.sh` file as needed.

Table 5-1 The baas-options.sh script environmental variables description

Environment variable	Description
DOCKER_REGISTRY_ADDRESS	For an online installation, images are pulled from the IBM Entitled Registry. You must specify 'cp.icr.io/cp'.
DOCKER_REGISTRY_USERNAME	For an online installation, you must specify 'cp'.
DOCKER_REGISTRY_PASSWORD	For an online installation, the registry password should be the entitlement key that was retrieved as described in 5.2.1, “Obtaining the IBM Container Software Library Entitlement key” on page 87.
DOCKER_REGISTRY_NAMESPACE	For an online installation, you must specify 'sppc'.
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 on All Resources.
SPP_ADMIN_PASSWORD	The IBM Spectrum Protect Plus password for the containers administrator. Optionally, you can specify an environment variable for the password; for example, \${PROTECTPLUS_ADMIN_PW}.
DATAMOVER_USERNAME	<ul style="list-style-type: none"> ► The user ID that is to create for use with the data mover. The value does not have to exist already because it is created during the installation. ► The data mover user ID must adhere to the rules for usernames and passwords for Red Hat Enterprise Linux (RHEL) 7 operating system. The rules for creating a user on RHEL 7 apply; for example, the password and the username cannot be the same.
DATAMOVER_PASSWORD	<p>The user password that is used to create for use with the data mover. The value does not have to exist already because it is created for the installation. The data mover password must adhere to the rules for passwords for RHEL 7. The rules for creating a user on RHEL 7 apply. For example:</p> <ul style="list-style-type: none"> ► The password must be at least 8 characters, and must contain letters and numbers. ► No dictionary words are allowed in the password. ► The password cannot be the same as the username.

Environment variable	Description
MINIO_USERNAME	<ul style="list-style-type: none"> ▶ The username that is used to create for MinIO object storage, which is used to store backups of cluster and namespace resources. ▶ The value does not have to exist already because it is created during the installation. ▶ The username must be at least 3 characters in length.
MINIO_PASSWORD	<ul style="list-style-type: none"> ▶ The password that is used to create for the MinIO user. ▶ The value does not have to exist already because it is created during the installation. ▶ The password must be at least 8 characters.

Preparing the baas-value-cr.yaml file

The goal of this process is to create the baas-value-cr.yaml file, which describes the Container Backup Support custom resource deployment.

Although this file is used mostly for CLI installation, it also is used as a reference for the GUI installation.

Update the example that is shown in Example 5-3 according to your infrastructure, with the help of Table 5-2 on page 91.

Example 5-3 Sample baas-values-cr.yaml that is used for the lab

```

apiVersion: sppc.ibm.com/v1
kind: IBMSPPC
metadata:
  namespace: baas
  name: ibmsppc-entitled-registry
spec:
  state: present
  version: 10.1.12
  license:
    accept: true
  product_loglevel: DEBUG
  cluster_name: dpr
  image_registry_namespace: sppc
  image_registry: cp.icr.io/cp
  image_registry_secret: baas-registry-secret
  spp_port: '443'
  spp_ips: 10.0.240.222
  spp_fqdn: 10.0.240.222
  networkPolicy:
    cluster_api_serverips:
      - 10.0.240.232
      - 10.0.240.233
      - 10.0.240.234
    cluster_api_serverport: 6443
    cluster_cidr: 10.254.0.0/16
    is_server_installed_on_another_cluster: false
    other_cluster_cidr_block: '24'
  minio_storage_class: ocs-storagecluster-ceph-rbd

```

Table 5-2 The `baas-values-cr.yaml` variables description

Parameter	Description	Default Value
<code>license</code>	Set the value to True to indicate that you reviewed and agree to the license agreement.	false
<code>product_loglevel</code>	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
<code>cluster_name</code>	The unique cluster name that is used to register the application host in IBM Spectrum Protect Plus.	None
<code>image_registry_namespace</code>	In an online installation, you must specify <code>sppc</code> .	baas
<code>image_registry</code>	In an online installation, you must specify <code>cp.icr.io/cp</code> .	docker-repo-hostname: 5000
<code>image_registry_secret</code>	The image pull secret for Container Backup Support that provides the credentials for pulling images from the registry where container images are loaded. For an online installation, you must specify <code>baas-registry-secret</code> to pull images from the IBM Entitled Registry.	None
<code>spp_port</code>	The IBM Spectrum Protect Plus server port. The value must be 443.	443
<code>spp_fqdn</code>	<p>The DNS address for the IBM Spectrum Protect Plus server. You can specify an IP address or a fully qualified domain name. Consider the following points:</p> <ul style="list-style-type: none"> ► 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 <code>spp_ips</code> parameter. ► If the IBM Spectrum Protect Plus server and the Container Backup Support are deployed on the same Red Hat OpenShift cluster, <code>spp_fqdn</code> is the full name of the IBM Spectrum Protect Plus server's <code>sppproxy</code> service. It is <code>sppproxy.spp_server_namespace.svc</code>, where <code>spp_server_namespace</code> is the namespace for IBM Spectrum Protect Plus server. ► If the IBM Spectrum Protect Plus server is installed in a Red Hat OpenShift environment and the Container Backup Support is on another Red Hat OpenShift cluster, retrieve the DNS address by issuing the following command: <code>oc get route --namespace spp_server_ns</code> Where <code>spp_server_ns</code> 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

Parameter	Description	Default Value
spp_ips	<ul style="list-style-type: none"> ▶ If the IBM Spectrum Protect Plus server is installed as a virtual appliance, specify its IP address. ▶ If the IBM Spectrum Protect Plus server and the Container Backup Support are deployed on the same Red Hat OpenShift cluster, use the sppproxy IP address: <code>oc get service sppproxy --namespace spp_server_ns</code> ▶ If the IBM Spectrum Protect Plus server and the Container Backup Support are not deployed on the same Red Hat OpenShift cluster, by using the value of the spp_fqdn parameter, use one of the following methods: <ul style="list-style-type: none"> – Run the following command: <code>nslookup my.plus.server.example</code> – 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 DNS debugging resolution. <p>For example:</p> <code>kubectl exec -i -t dnsutils -- nslookup my.plus.server.example</code> 	x.x.x.x
cluster_api_serverips	<p>The master nodes (where the apiserver is running) of the IP addresses. On Red Hat OpenShift, use the following command: <code>oc get endpoints -n default -o yaml kubernetes</code> Use all the IP addresses listed under subset.addresses, or add or remove IP addresses as needed. Specify multiple addresses as follows:</p> <pre>networkPolicy: clusterAPIServerips: - x.x.x.x - y.y.y.y - z.z.z.z</pre>	None
cluster_api_serverport	<p>The port used by the apiserver. On Red Hat OpenShift, use the following command: <code>oc get endpoints -n default -o yaml kubernetes</code> Use the value listed in subsets.ports[name="https"].port in the output.</p>	6443
cluster_cidr	<p>The Classless Inter-Domain Routing (CIDR) value for the cluster. To obtain the CIDR on Red Hat OpenShift, issue the following command: <code>oc get network -o yaml grep -A1 clusterNetwork:</code> Use the displayed IP address as the cluster CIDR address.</p>	192.168.0.0/16
is_server_installed_on_another_cluster	<p>Specifies whether the IBM Spectrum Protect Plus server is installed on another Red Hat OpenShift Cluster. Set the value to <code>false</code> if:</p> <ul style="list-style-type: none"> ▶ You are installing the product on a Kubernetes cluster, or the IBM Spectrum Protect Plus server is installed as a virtual appliance. ▶ You are installing the product on a Red Hat OpenShift cluster and the IBM Spectrum Protect Plus server is installed on the same cluster. <p>Set the value to <code>true</code> if you are installing the product on a Red Hat OpenShift cluster and the IBM Spectrum Protect Plus server is installed on a separate Red Hat OpenShift cluster.</p>	false

Parameter	Description	Default Value
<code>other_cluster_cidr_block</code>	The group of CIDRs and IP addresses of other clusters. This parameter is used to enable IBM Spectrum Protect Plus servers on different clusters to connect to the Container Backup Support container agent in the network policy. If you do not set this value, the network policy acts as a firewall and prevents connections to the cluster on which Container Backup Support is being installed.	None
<code>minio_storage_class</code>	The name of the storage class to use for the MinIO server: <ul style="list-style-type: none"> ► 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. ► To retain snapshot backups of cluster and namespace resources between BaaS uninstalls and reinstalls, use storage class for which reclaim policy is set to Retain or manually set MinIO PV reclaim policy to Retain. For more information, see Reusing MinIO persistent volume on reinstall. 	None
<code>vertical_pod_autoscaler</code>	Optional: For Red Hat OpenShift only. To adjust pod resource levels with the Vertical Pod Autoscaler (VPA) operator, specify the following optional configuration parameters: vertical_pod_autoscaler: namespace: ns_vpa_operator_installed update_mode_policy: Initial Off <ul style="list-style-type: none"> ► vertical_pod_autoscaler is an optional parameter that enables Vertical Pod Autoscaler functions in the Container Backup Support. ► ns_vpa_operator_installed specifies the namespace in which the VPA operator is installed. ► update_mode_policy specifies the policy, which you can set to initial or off: <ul style="list-style-type: none"> – Initial mode applies VPA recommendations when pod is restarted. – Off mode provides only the recommended values so that you can apply them manually. 	None

Note: In Example 5-3 on page 90, we did not use vertical pod autoscaler, and thus did not specify it. For more information, see 4.4.3, “Vertical Pod Autoscaler” on page 75.

5.2.3 Editing the cluster global pull secret

The cluster global pull secret references the different credentials to access the different container registries.

This step adds the credentials for the IBM Container Software Library to this global pull secret, which is in the openshift-config namespace.

Using the CLI

To add the IBM Container Software library to the global pull secret by using CLI, the `jq` tool must be installed. This tool enables processing JSON content from the CLI (see Example 5-4).

Example 5-4 Editing the global pull secret by using the CLI

```
# Source the baas-options.sh file
$ . ./baas-options.sh

# Extract the existing global pull secret configuration to the file
.dockerconfigjson
$ oc extract secret/pull-secret -n openshift-config --confirm > /dev/null

# Create the base64-encoded registry authentication string based on variables
# from baas-options.sh
$ AUTH=$(printf "${DOCKER_REGISTRY_USERNAME}:${DOCKER_REGISTRY_PASSWORD}" |
base64)

# add a docker configuration for cp.icr.io to the existing docker configurations
# and save it to the file new_dockerconfigjson
$ cat .dockerconfigjson | jq --compact-output
".auths[\"${DOCKER_REGISTRY_ADDRESS}\"] |= . + {\"auth\": \"${AUTH}\"}" >
new_dockerconfigjson

# Update the global pull secret with the new docker configurations
$ oc set data secret/pull-secret -n openshift-config \
--from-file=.dockerconfigjson=new_dockerconfigjson

# Remove the temporary files
$ rm .dockerconfigjson new_dockerconfigjson
```

Using the Red Hat OpenShift web console

To edit the cluster global pull secret, complete the following steps:

1. Log in to the Red Hat OpenShift web console.
2. Select **Workloads** → **Secrets**, and look for the `pull-secret` secret from the project `openshift-config`, as shown in Figure 5-2 on page 95. Click the secret to select it.

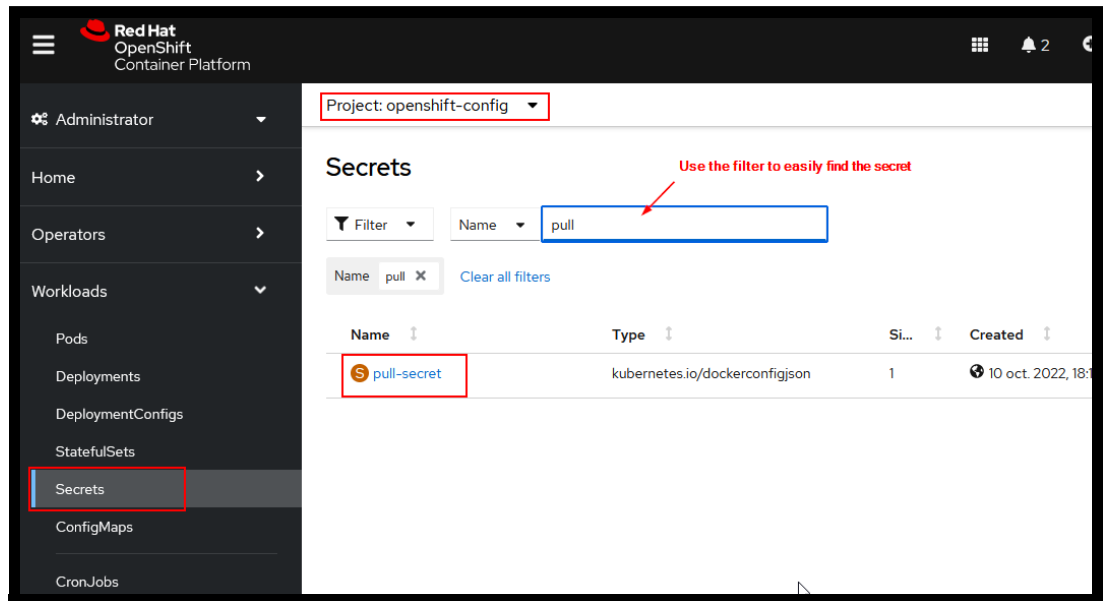


Figure 5-2 Locating the global pull secret

- From the Secret window, select **Edit Secret** from the **Actions** menu, as shown in Figure 5-3.

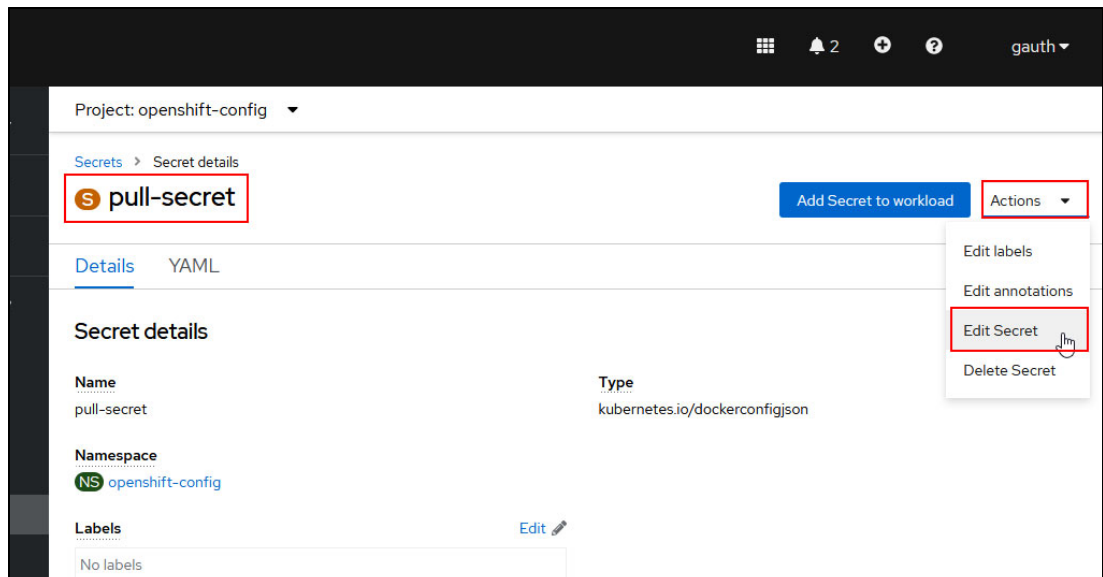


Figure 5-3 Going to the secret edition window

- From the next window, click **Add credentials**, as shown in Figure 5-4, complete the fields that are listed in Table 5-3 and then, click **Save**.

Table 5-3 Fields to be completed

Field	Value
Registry server address	For an online installation, enter cp.icr.io/cp.
Username	For an online installation, enter cp.
Password	For an online installation, enter your entitlement key, which was retrieved as described in 5.2.1, “Obtaining the IBM Container Software Library Entitlement key” on page 87.

The screenshot shows the Red Hat OpenShift Container Platform console. On the left is a sidebar with navigation links: Administrator, Home, Operators, Workloads (Pods, Deployments, DeploymentConfigs, StatefulSets), Secrets (highlighted), ConfigMaps, CronJobs, Jobs, DaemonSets, ReplicaSets, ReplicationControllers, HorizontalPodAutoscalers, Networking, and Storage. The main panel shows the 'Project: openshift-config' configuration page. It has a 'Remove credentials' button and a 'Add credentials' button (highlighted with a red box). Below these are input fields for Username, Password, and Email. A red arrow points to the Password field with the text 'Your IBM Container Library entitlement key'.

Figure 5-4 Editing the global pull secret

5.2.4 Creating the Container Backup Support namespace

In the step, we create the namespace that hosts the Container Backup Support application. The namespace is named baas.

Using the CLI

Example 5-5 shows how to create the Container Backup Support namespace.

Example 5-5 Creating the Container Backup Support namespace by using the CLI

```
$ oc new-project baas
```

Using the Red Hat OpenShift web console

From the Red Hat OpenShift web console, select **Home** → **Projects**. Then, click **Create Project**, as shown in Figure 5-5.

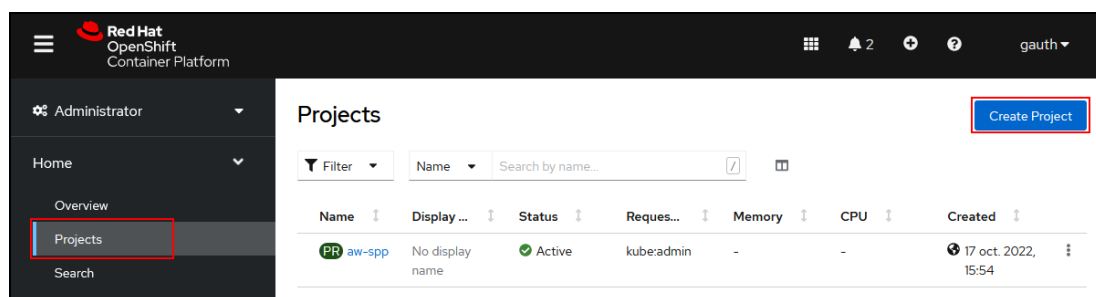


Figure 5-5 Accessing the Project window in Red Hat OpenShift web console

The Create Project window opens. Enter the required values, as shown in Figure 5-6.

The 'Create Project' dialog box is shown. It contains the following fields: 'Name' with the value 'baas', 'Display name' with the value 'Spectrum Protect Plus Baas', and an empty 'Description' text area. At the bottom right, there are 'Cancel' and 'Create' buttons. A mouse cursor is pointing at the 'Create' button.

Figure 5-6 Creating the baas project in the Red Hat OpenShift web console

5.2.5 Creating the image pull secret

In this step, an image pull secret (that is, a secret with registry credentials) is created to allow the Container Backup Support to pull images from the IBM Container Software Library.

This `baas-registry-secret` secret is created in the namespace `baas`.

Using the CLI

Example 5-6 shows how to create the image pull secret by using the CLI.

Example 5-6 Creating the image pull secret

```
# Source the baas-options.sh file to make the variables accessible.
$ ./baas-options.sh
# Create the baas-registry-secret
$ oc create secret docker-registry baas-registry-secret --namespace baas\
    --docker-server=${DOCKER_REGISTRY_ADDRESS}\
    --docker-username=${DOCKER_REGISTRY_USERNAME}\
    --docker-password=${DOCKER_REGISTRY_PASSWORD}
```

Using the Red Hat OpenShift web console

To create the image pull secret by using the Red Hat OpenShift web console, complete the following steps:

1. From the Red Hat OpenShift web console, select **Workloads** → **Secrets**, click **Create**, and then select **Image pull secret**, as shown in Figure 5-7. Be sure to be in the baas project.

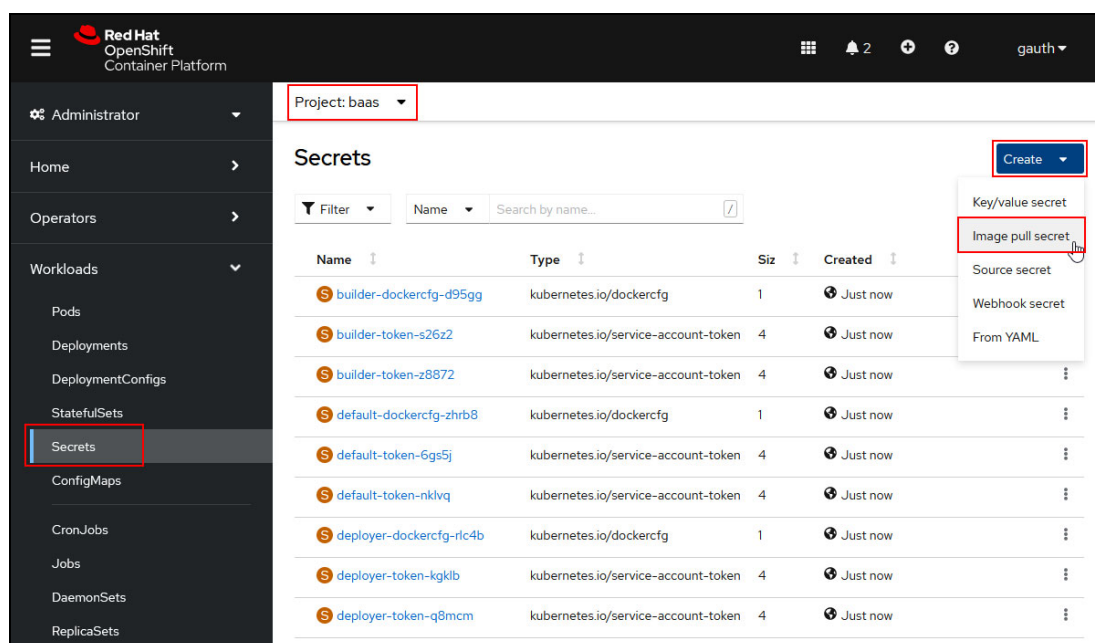


Figure 5-7 Accessing the secrets window in the baas project by using the Red Hat OpenShift web console

2. On the Create Image Pull Secret window (see Figure 5-8 on page 99), complete the fields that are listed in Table 5-4 and then click **Create**.

Table 5-4 Fields to be completed

Field	Value
Secret name	Enter baas-registry-secret.
Authentication type	Select Image registry credentials .
Registry server address	For an online installation, enter cp.icr.io/cp.

Field	Value
Username	For an online installation, enter cp.
Password	For an online installation, enter your entitlement key, which was retrieved in 5.2.1, “Obtaining the IBM Container Software Library Entitlement key” on page 87.

Project: baas

Create image pull secret

Image pull secrets let you authenticate against a private image registry.

Secret name *
baas-registry-secret
Unique name of the new secret.

Authentication type
Image registry credentials

Registry server address *
cp.icrio/cp
For example quay.io or docker.io

Username *
cp

Password *
.....

Email
.....

[Add credentials](#)

[Create](#) [Cancel](#)

Your IBM Container Library entitlement key

Figure 5-8 Creating the baas-registry-secret by using the Red Hat OpenShift web console

5.2.6 Creating the general-purpose secret for Container Backup Support

The Container Backup Support requires credentials for the IBM Spectrum Protect Plus server, data movers, and the MinIO instance. This information is stored in a baas-secret secret in the baas namespace.

Using CLI

Example 5-7 shows how to create the general-purpose secret for Container Backup Support.

Example 5-7 Creating the general-purpose secret for Container Backup Support

```
# Source the baas-options.sh file to make the variables accessible.
$ ./baas-options.sh
# Create the baas-secret
kubectl create secret generic baas-secret --namespace baas \
  --from-literal='baasadmin="'${SPP_ADMIN_USERNAME}'"' \
  --from-literal='baaspassword="'${SPP_ADMIN_PASSWORD}'"' \
  --from-literal='datamoveruser="'${DATAMOVER_USERNAME}'"' \
  --from-literal='datamoverpassword="'${DATAMOVER_PASSWORD}'"' \
```

```
--from-literal='miniouser='${MINIO_USERNAME}' ' \
--from-literal='miniopassword='${MINIO_PASSWORD}' '
```

Using the Red Hat OpenShift web console

To create the general-purpose secret for Container Backup Support by using the Red Hat OpenShift web console, complete the following steps:

1. From the Red Hat OpenShift web console, select **Workloads** → **Secrets**, click **Create**, and then select **Key/value secret**, as shown in Figure 5-9. Be sure to be in the baas project.

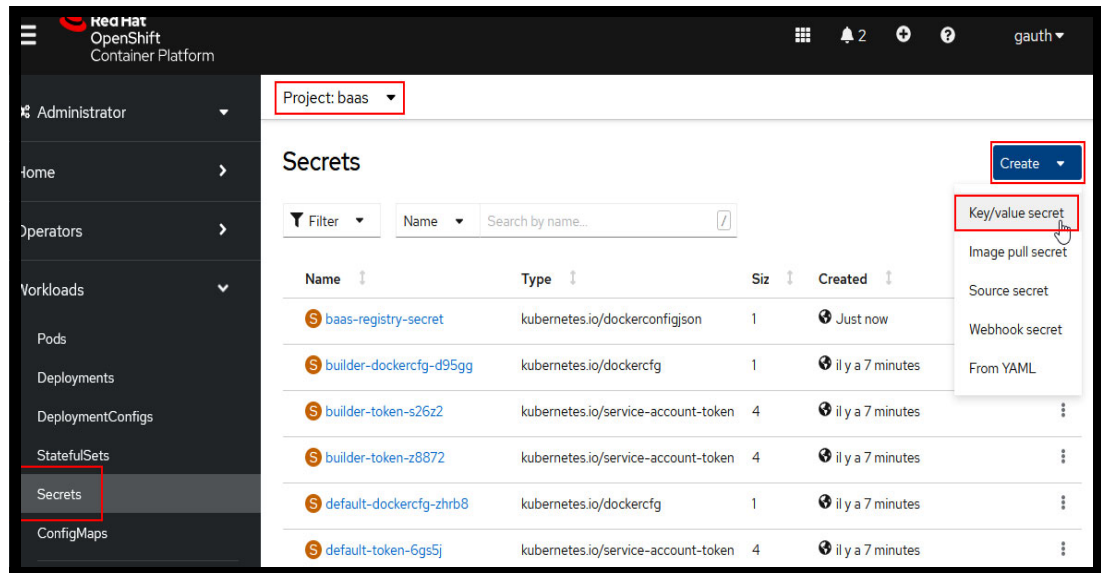


Figure 5-9 Accessing the secrets window in the baas project by using the Red Hat OpenShift web console

2. On the Create key/value Secret window, as shown in Figure 5-10 on page 101, complete the fields that are listed in Table 5-5 and click **Create**.

Tip: Use the baas-options.sh file that was created as described in “Preparing the baas-options.sh file” on page 88 to find the required different values for this secret.

Click the **Add key/value** link to add key/value pairs to the secret, as shown in Figure 5-10 on page 101.

Table 5-5 Fields to be completed

Key	Value
Secret name	baas-secret
baasadmin	The user ID of the IBM Spectrum Protect Plus containers administrator. This value of SPP_ADMIN_USERNAME is from the baas-options.sh file.
baaspassword	The password of the IBM Spectrum Protect Plus containers administrator. This value of SPP_ADMIN_PASSWORD is from the baas-options.sh file.
datamoveruser	The user ID to use with the data movers. This value of DATAMOVER_USERNAME is from the baas-options.sh file.

Key	Value
datamoverpassword	The password to use with the data movers. This value of DATAMOVER_PASSWORD is from the baas-options.sh file.
miniouser	The user ID to use with for the MinIO instance. This value of MINIO_USERNAME is from the baas-options.sh file.
miniopassword	The password to use with for the MinIO instance. This value of MINIO_PASSWORD is from the baas-options.sh file.

Project: baas ▾

Create key/value secret

Key/value secrets let you inject sensitive data into your application as files or environment variables.

Secret name *

Unique name of the new secret.

Key *

Value

Drag and drop file with your value here or browse to upload it.

baas

[+ Add key/value](#)

Figure 5-10 Creating the baas-secret from the Red Hat OpenShift web console

5.2.7 Exchanging a certificate between IBM Spectrum Protect Plus and Red Hat OpenShift Container Platform

In the latest Container Backup Support version, communications between the IBM Spectrum Protect Plus Server and the Container Backup Support use APIs that are secured by TLS certificates.

You do not exchange a certificate in the following environments:

- ▶ IBM Spectrum Protect Plus and Container Backup Support agent are deployed in the same Red Hat OpenShift cluster.
- ▶ IBM Spectrum Protect Plus is configured with the certificates from a known certificate authority.

In our lab, we use self-signed certificates, so we must perform the actions that are described in the following sections.

Obtaining the IBM Spectrum Protect Plus server certificate

If the IBM Spectrum Protect Plus server is installed as a virtual appliance, obtain the required certificates by running the following command:

```
openssl x509 -in <(openssl s_client -connect spp_server_address:443 -prexit  
2>/dev/null) | base64 --wrap=0
```

Where `spp_server_address` is the IP address or Fully Qualified Domain Name (FQDN) of your IBM Spectrum Protect Plus instance (see Example 5-8).

Example 5-8 Obtaining the IBM Spectrum Protect Plus server certificate when the server is installed as a virtual appliance

```
$ openssl x509 -in <(openssl s_client -connect 10.0.240.222:443 -prexit  
2>/dev/null) | base64 --wrap=0  
LS0tLS1CRUdJTiBDRVJUSUZJQ0FURSB0tLS0tCk1JSUZZVENDQTBtZ0F3SUJBZ01VVFMEFZ10Uprb3dYN2  
FwMwdjY1dWcjrZ2Zm5NdORRWUpLb1pJaHJzJTkFRRUUwKQ1FBd1dqRUxNQWtHQTFVRUJoTUNWVnk14Q3pBSk  
JnTlZCQWdNQWslWk1ROHdEUUVlEVlFRSEBkVWkjb2F2Ym1zeApHREFXQmdOVKBjB01EMGxDVFNCRGIzSndiM0  
poZEdsdmJqRVRNqkVHQTFVRUF3d0tjMOJ3TFh0bGNUMmxjakF1CkZ3MH1NakV3TVRFeE1qRTBNRFZlRncw  
ek1qRXdnRGd4TWpFME1EVmFNRM94Q3pBSk1JnTlZCQV1UQWxWVE1Rc3cKQ1FZRFZRUU1EQUpPV1RFUE1BME  
dBMVVFQnd3R1FYSnRiMjYyTVJnd0ZnWURWUWFLREES1FrmGdRMj15Y0c5eQpZWFJwYj10eEV6QVJCZ05W  
KkFNTUNuTndjQzF6W1hKMLpYSXdnZ01pTUEwR0NTcUdTSWIZRFFFQkFRVUFBNE1DCkR3QXdnZ01LQW9JQ0  
FRQzRjZjV5bkZQRHN5cE5lOFJwOG9FaG1GV21yYjFQZEV6b3BD0Vd2NTJmOV1xck1JK1EKZnUybmttS1FB  
VUFhVFcwOWxkRG1XZkZSd2d3bFFocGF1SE1QK1JwZ096U0NUSm9RRXJSSmtWTFg2VUZP0VhYSgppqWT1kZ1  
NhWGPdUd0cvOE11STF1Zm1WbWNPbjV5WWYyRGZITjJWQzRxbnFiWFFQUXM0bEY5TWIOMVE4SW5VMmxvCm1p  
UjRLRVcyc3pvem5qSUVPOGtMRGpmUG9UNEk5NUpTNGtLUEkxWGHjWE9iRVRwNUVNOGJCvZdmSnJiZThvRH  
YKUGQ2Q1N6dE9pL09qR1M4KOMzUE56dzQweFgzRUdHRU1oSDdyNG51a0tVVDcyEEdDU1g1UkpOM01scER2  
OW02WApEY1FWTX1sSVhxM1FId1hkdw9Sc0piZFpqUERPUzthbTlEWnYzY2gwcK5PVURvMFppd1FPQTGTW  
p3Z0xweHC0ClD0VDRHUUFQU1Q3UFcwY19KQnBVNkFYRm1MNF14L0VkeTFBRGswWkwyemdbmJkRmFpVnN3  
NWozVUpjWjYXJFYXkKVTvub1IwV05aUDh3TE9RaUpYVWxSWU0b11KbTJZYW1uQXNzSy9TK0xFMz10W1JKUH  
gzcEdpb21qdEQ4bDVMMApPUF1WQXhCRmY4em1EQi85MOozcXhD0ThWTTfXeThqUG1HaDREb01CUEQ5dGxK  
UWFzQWVmRV1rZFIyYk1uS2FLC1FDVVM2enZsbGtoZGVnQ0pwUHYyWEwyS1Fwa1V1a3o5M1dWM0x0ZDQ0Z1  
pJMTVqWxVIUVd2SWdSNwt4c1p1M1YKb1VtZE56ZVBMdnR4WHJjQVFze1U5ZG1KUkdJQXcwZFd5aGFtZ0JH  
YORYRU10ckEyOHJUT2wzMzRvd01EQVFBQgppveDh3SFRBYk1JnTlZiUkVFRkRBU2dnchPjSEF0YzJWeWRTVn  
1od1FLQVBEZU1BMEdDU3FHU01iMORRRUJDd1VBCKE0SUNBUUJaziTcWwJYQk9jUE1PYX1MZDBoY1kvekor  
cVc1c1BPQnJSSWpMeWnDa25ubkEwMUTvVXR0T3hpUnAKbGRJNTNsR25kQ1FhcDNBaEowSHFYedIvaGFGVU  
xFUUZHVF1FSc2ZmN1FLQTAvRUPOZ3NzaWp1R2VnOXdGVndWeQpQWUwS2NONE1TeGnydWdKMjFidTdIKOZJ  
WHpkZmxsYnJNRVNSNkVCM1ZJVnJiB11TXhLdm1FdEFCN1JaN3NvCmhWc0sxMk85QkF4Wk1xdm1zcys5d2  
pvTzN3eVpIcE54U3orVkyWMS91R1Zrc2xLb31rVVNkbS90dENNdfk0Zy8KTkFrb1pKRHN6aFVsY1NVSkRi  
WmRZR1pNTHBVCVVerNXZmUENTOUZ6MnNpbFpGekJwUG1zVKEvQ29zZVhMTGx4cwpwMFNKY1FoNG85WmFDWG  
F3aX1OMGhwdWc5ZTBSbGYwZGpVNXMvTjdpbV1TWdHhZQVhRbGpGengyQndNdDZ1UTgwCnhTYz1FQkZncG5H  
VzBzUmpZeFBuV0ZBN0FpM2tyMTVQcWdjTXIxSutHOGhCdD12RDVwVzNwekUxZXRVeDZ4Vm0KQmpQUWQxc  
toZ1NOK1g4Wkg3QzJWRDJXSWFpYmtol3dKZ0t1am95WHFaTmNuTGtqNHVCWUQyaVVGWEhzNm1DQworbEdo  
Ump0L3RaK3BDeDM5L09ZazE2NzkrRE93WnJuWHFyVG5yQ1p4WDdSVUZLe1V1WFJXbXh1VzZwOF1kt1M3Cn  
p00DBZMjVSeFMOQzZH21hTZFJZWERvWnR2Ym0zdGgraOdrRkpJTFM0ZGtmbHBQb0k4Q3hnZWht1JWQ09m  
OUOKcnNwSkxQb2JTRU1LK2JWtZn6N0pvYkhxViTjBZRR3Bjb1NaK2xtESzaFRnOXFJW1E9PQotLS0tLU  
VORCDBRVJUSUZJQ0FURSB0tLS0tCg==
```

If the IBM Spectrum Protect Plus server is installed on a Red Hat OpenShift cluster, obtain the required certificates by running the following command (see Example 5-9 on page 103):

```
oc get secret router-ca -n openshift-ingress-operator -o  
jsonpath="{.data['tls.crt']}"
```


Example 5-9 Obtaining the IBM Spectrum Protect Plus server certificate when the server is installed on Red Hat OpenShift

```
$ oc get secret router-ca -n openshift-ingress-operator -o
jsonpath="{.data['tls\.crt']}"
LS0tLS1CRUdJTiBDRVJUSUZJQ0FURS0tLS0tCk1JSURERENDQWZTZ0F3SUJBZ01CQVRBTkNa3Foa2lHOX
cwQkFRc0ZBREftTVNRd0lnWURWUVFEREJ0cGJtZHKWl1h0ekxXOXdaWEpoZEc5eVFERTJ0ak0zTlRjeU56
Y3dIaGNOTWpJd09USXhNVEEwTnpVMl doY05Na1F3T1RJdwpNVEEwTnpVMldqQW1NU1F3SWdZRFZRUUREQn
RwYm1keVpYTnpMVz13W1hKaGRHOX1RREUyTnpNM05UY310emN3CmdnRW1NQTBHQ1NxR1NJYjNEUUVVCQVFV
QUE0SUJEd0F3Z2dFS0FvSUJBURhtFBUEXM2cTROUHVSVQ2VHM1oydFUKR1ZaL1loOHzONWJuc09ieUIrUD
Y3RVFKTKZwZW41eHA0a2JxWFlkTWBbnjJMcH1zVW11OTBNT0F5SndvNjA2bAp6TWhpSnRmdmR6cytUTjRF
QWVpQmEyYVJMOwtjbUtPSVRSTmRqZUQRUW9NK25ibktIOU5nRjJ4MDhVajRkaHJ0c1VFc3BheVYvdVJjc2
44Mkh1RmFSNT1zZURDd1FyTi9uQzdVdUN5ZWVrVHJXex1kUXB3MU9TdTN0ZHF5U1dHSEMkM0p4N3plYm8r
N3M5Sz12L0d4Z0J2bXBDenNxRnR3a0Fsa0szTzcxMGQraS8waE9KVGv2UHZ6djJERkdoSGNWRAPtCnVHT2
hpMG1aK1BU0GhwcmhZNWFpUDF4R11oeFYxWDFHbWVMROxMUHV2R1R4ODhhYndSU285T1BPQ01hM0tYckFn
TUJBQUdqU1RCRE1BNEdBmVvKRHdFQi93UUVBd01DcERBU0JnT1ZiUk1CQWY4RUNEQUdBUUgVQWdFQU1CME
cKQTFVZERNUVdCQ1FSQjdLeTFwTU81ZUZwUHDHdEh2dk9ERm9UU2pBTkNa3Foa2lHOXcwQkFRc0ZBQU9D
QVFFQkQpJa0ZubHdYUzFadC9QVkh6dkhWVUtyYVEONVA2M3VFTUpwZFRVUmxHbXgyOUdIckYzOTg5cGpxV2
RnczdLNE14CjZDOWx6bXc5M11JQ1VF3dWeD1Hbk9WSmp4T1pqWkJEaFgvd1U2a1lQMnBhWD1oNG8yVmNW
V2h1MUDRbD1VTkoKVUZKSvVRU3dJaW5TU21uTvc5SGUzUCtYVjA4NE9KY2x1T1E5MGF5OHZ1c1kxc2xCUE
91WEpxYjJ6SEQzc31jSAoyZGd0QU12S1FBUONEQmJSckRnTUpGR1dLSFNZRFhTbFI4UnhteFZFUTd3RvDd
bXdPck1EeFZLc0hNZm9QbTZTCnZNbnpIQ0VXVHhPT0tydk1WRF1aZzg5UE10dKNVZit3T2JTR0dLdDk2bT
11Q29QNmxJczNaNStxYW9wMSthemIKck5QSEpudmJVRG9EWFRGOVNwS3RoZz09Ci0tLS0tRU5EIEENFU1RJ
Rk1DQVRFLS0tLS0K
```

Importing the IBM Spectrum Protect Plus certificate

In this section, a `baas-spp-server-cert` secret is created by using the IBM Spectrum Protect Plus certificate that was obtained as described in “Obtaining the IBM Spectrum Protect Plus server certificate” on page 102. This `baas-spp-server-cert` secret is created in the Container Backup Support namespace, `baas`.

Using the CLI

Create a file `baas-spp-server-cert.yaml` with the content that is shown in Example 5-10.

Example 5-10 baas-spp-server-cert.yaml content

```
apiVersion: v1
kind: Secret
metadata:
  name: baas-spp-server-cert
  namespace: baas
type: Opaque
data:
  tls.crt: <base64-encoded certificate>
```

Where `<base64-encoded certificate>` is the certificate that was obtained as described in “Obtaining the IBM Spectrum Protect Plus server certificate” on page 102.

Create the secret by running the command that is shown in Example 5-11.

Example 5-11 Creating the baas-spp-server-cert secret by using the CLI

```
$ oc apply -f baas-spp-server-cert.yaml
secret/baas-spp-server-cert created
```

Using the Red Hat OpenShift web console

Complete the following steps:

1. From the Red Hat OpenShift web console, click **Import YAML (+)**.
2. Paste the secret YAML manifest that is shown in Example 5-10 on page 103. Click **Create** (see Figure 5-11).

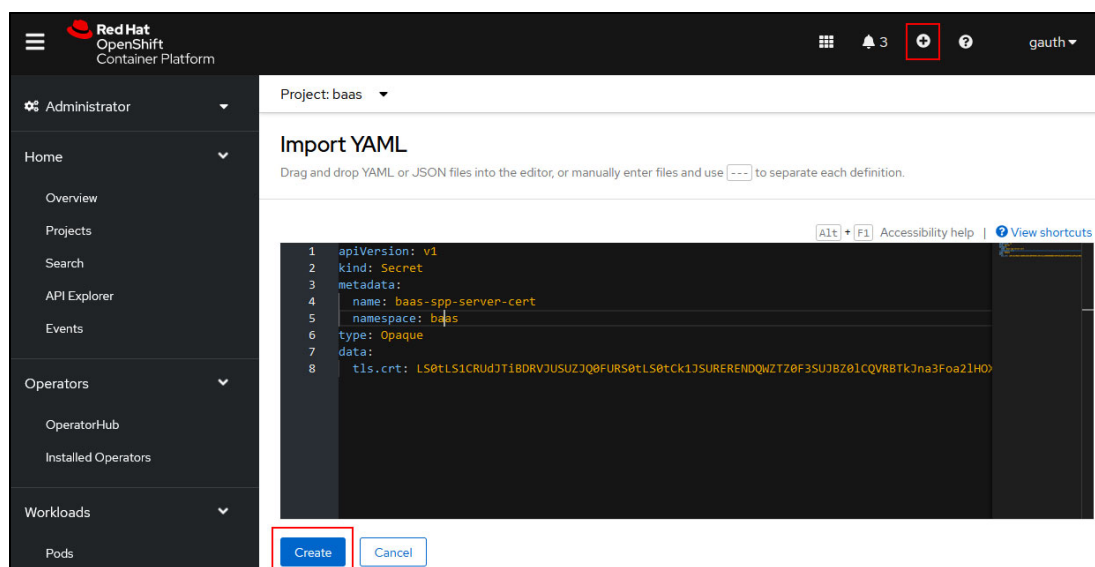


Figure 5-11 Creating the baas-spp-server-cert secret by using the Red Hat OpenShift web console

5.2.8 Adding the online catalog source

A catalog source is a repository of Custom Resource Definitions, Operators, and so on.

In this section, the IBM Container Software Library catalog source is added to the Red Hat OpenShift cluster.

As a result, new operators are available, including the Container Backup Support operator.

Using the CLI

Complete the following steps:

1. Create a file `baas-ocp-catalogsource.yaml` by using the content that is shown in Example 5-12.

Example 5-12 baas-ocp-catalogsource.yaml content

```
apiVersion: operators.coreos.com/v1alpha1
kind: CatalogSource
metadata:
  name: ibm-operator-catalog
  namespace: openshift-marketplace
spec:
```

```
displayName: IBM Operator Catalog
image: 'icr.io/cpopen/ibm-operator-catalog:latest'
publisher: IBM
sourceType: grpc
updateStrategy:
  registryPoll:
    interval: 45m
```

2. Create the catalog source by running the command that is shown in Example 5-13.

Example 5-13 Creating the IBM Operators catalog source by using the CLI

```
$ oc apply -f baas-ocp-catalogsource.yaml
catalogsource/ibm-operator-catalog created
```

3. Check whether the catalog source is ready by running the command that is shown in Example 5-14.

Example 5-14 Checking the catalog source status by using the CLI

```
$ oc describe catalogsource -n openshift-marketplace ibm-operator-catalog | grep
"Last Observed State"
    Last Observed State:      READY
```

If the Last Observed State is Ready, the new operators are available in the Operator Hub.

Using Red Hat OpenShift web console

Complete the following steps:

1. From the Red Hat OpenShift web console, click **Import YAML (+)**.
2. Paste the catalog source YAML manifest that is shown in Example 5-12 on page 104 and click **Create** (see Figure 5-12).

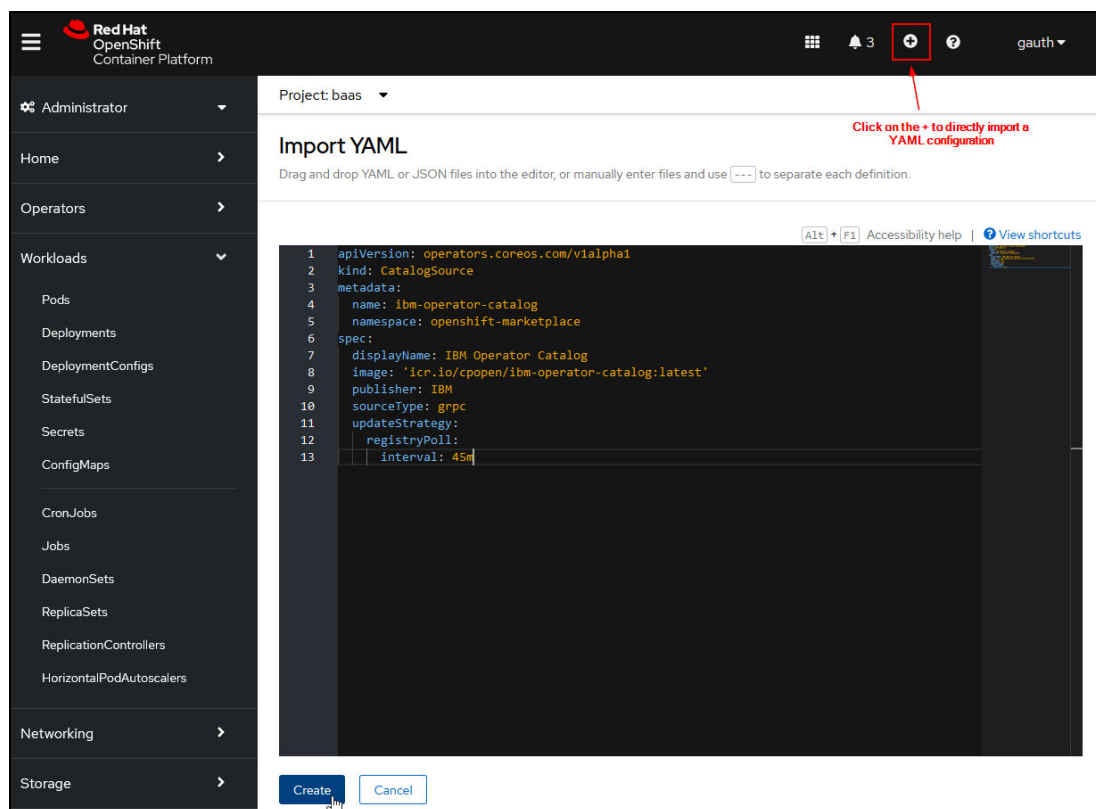


Figure 5-12 Creating the IBM Operators catalog source by using the Red Hat OpenShift web console

You are immediately redirected to the catalog source window, where the status can be monitored, as shown in Figure 5-13.

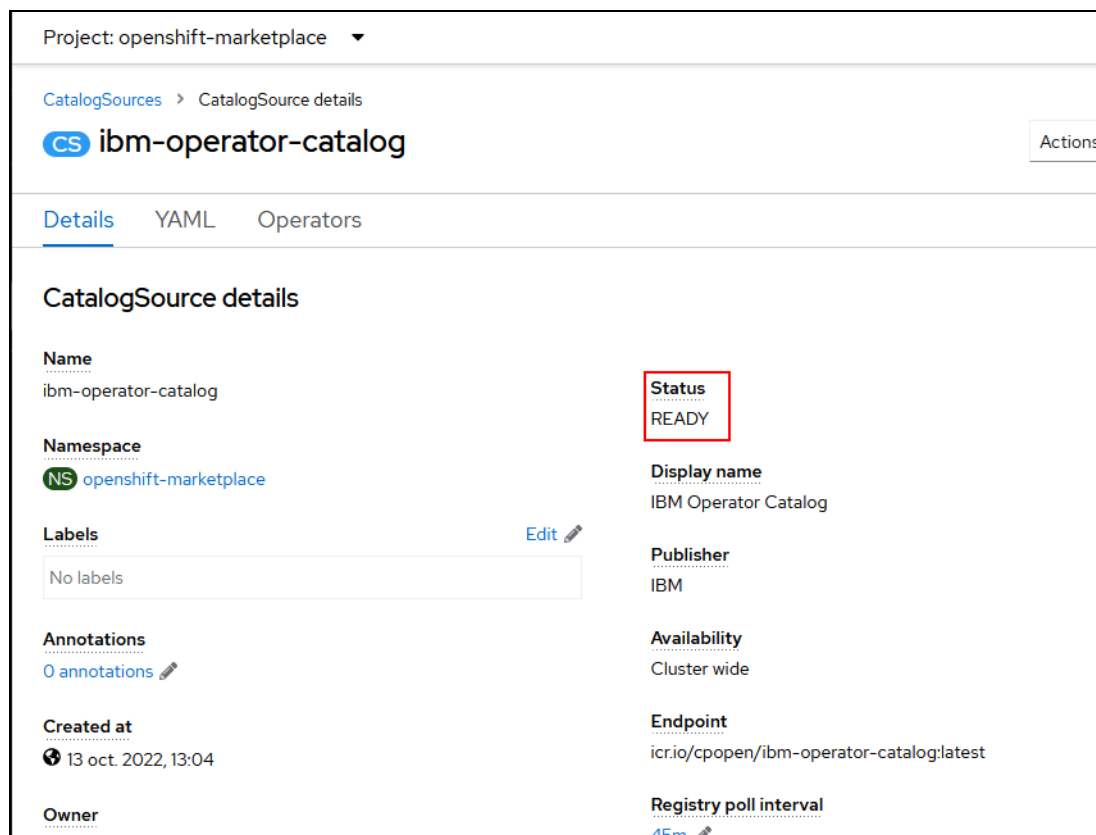


Figure 5-13 Checking the catalog source status by using the Red Hat OpenShift web console

5.2.9 Installing the Container Backup Support (BaaS) operator

After the IBM operators catalog source is created and ready, the Container Backup Support operator is available in the OperatorHub.

The Container Backup Support operator must be installed before you deploy the Container Backup Support instance.

Installing the Container Backup Support also installs the Red Hat OpenShift API for Data Protection operator and the Red Hat AMQ Stream (Strimzi) operator.

Using the CLI

To install the Container Backup Support (BaaS) operator by using the CLI, complete the following steps:

1. Create an Operator Group for Container Backup Support. For more information about Operator Groups, see [Operator groups](#).

2. Create a `baas-ocp-operatorgroup.yaml` with the content that is described in Example 5-15.

Example 5-15 baas-ocp-operator-group.yaml content

```
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: baas-ibmsppc
  namespace: baas
  annotations:
    olm.providedAPIs: IBMSPPC.v1.sppc.ibm.com
spec:
  targetNamespaces:
  - baas
```

3. Apply the YAML file to create the Operator Group, as shown in Example 5-16.

Example 5-16 Applying baas-ocp-operator-group.yaml by using the CLI

```
$ oc apply -f baas-ocp-operator-group.yaml
operatorgroup.operators.coreos.com/baas-ibmsppc created
```

4. Create the Container Backup Support subscription by using the `baas-oc-subscription.yaml` file with the content that is provided in Example 5-17.

Example 5-17 The baas-ocp-subscription.yaml file content

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: baas-operator
  namespace: baas
spec:
  channel: v1.6
  installPlanApproval: Automatic
  name: ibmsppc-operator
  source: ibm-operator-catalog
  sourceNamespace: openshift-marketplace
```

5. Apply the YAML file to create the subscription, as shown in Example 5-18.

Example 5-18 Applying the baas-ocp-subscription.yaml file by using CLI

```
$ oc apply -f baas-ocp-subscription.yaml
subscription.operators.coreos.com/baas-operator created
```

6. After a few minutes, you should see three pods in the `baas` project that correspond to the Red Hat OpenShift API for Data Protection and AMQ Stream (Strimzi) operators, as shown in Example 5-19.

Example 5-19 Checking the operator installations

```
$ oc get pod
```

NAME	READY	STATUS	RESTARTS
amq-streams-cluster-operator-v2.2.0-1-67d944c474-d9dkx	1/1	Running	0
ibmsppc-operator-controller-manager-54cf6bd56-51991	2/2	Running	0
openshift-adp-controller-manager-65745b84c5-c5qd9	1/1	Running	0

Using the Red Hat OpenShift web console

To install the Container Backup Support (BaaS) operator by using the Red Hat OpenShift web console, complete the following steps:

1. From the Red Hat OpenShift web console, select **Operators** → **OperatorHub**, and look for the operator IBM Spectrum Protect Plus Container Backup Support, as shown in Figure 5-14. (Be sure to have the project baas selected.)

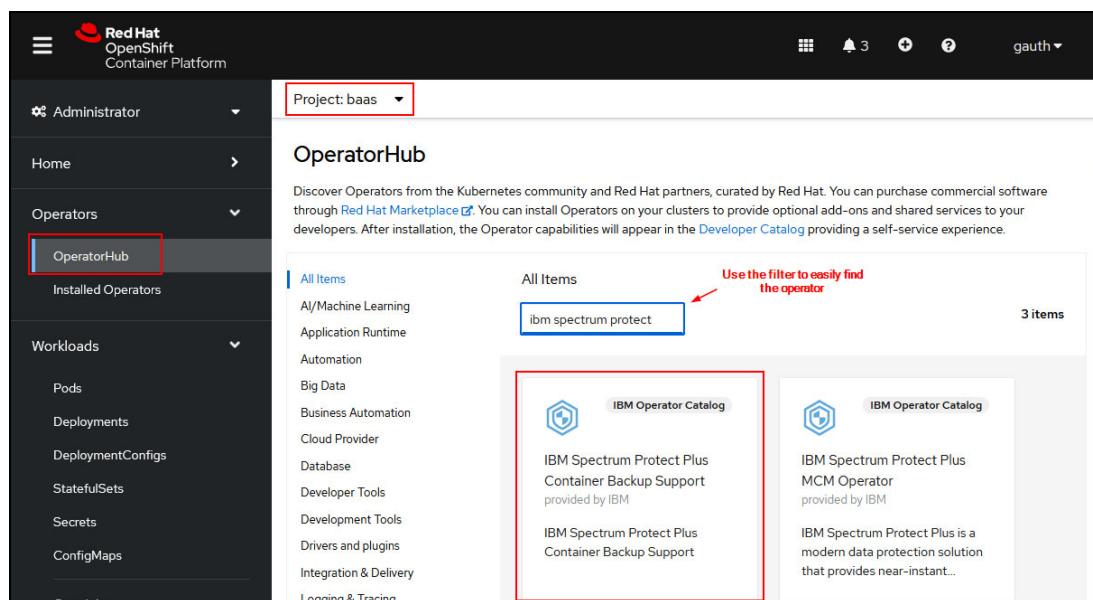


Figure 5-14 IBM Spectrum Protect Plus Container Backup Support in the OperatorHub

2. Click the operator tile, which opens a window. From this window, click **Install**, as shown in Figure 5-15.

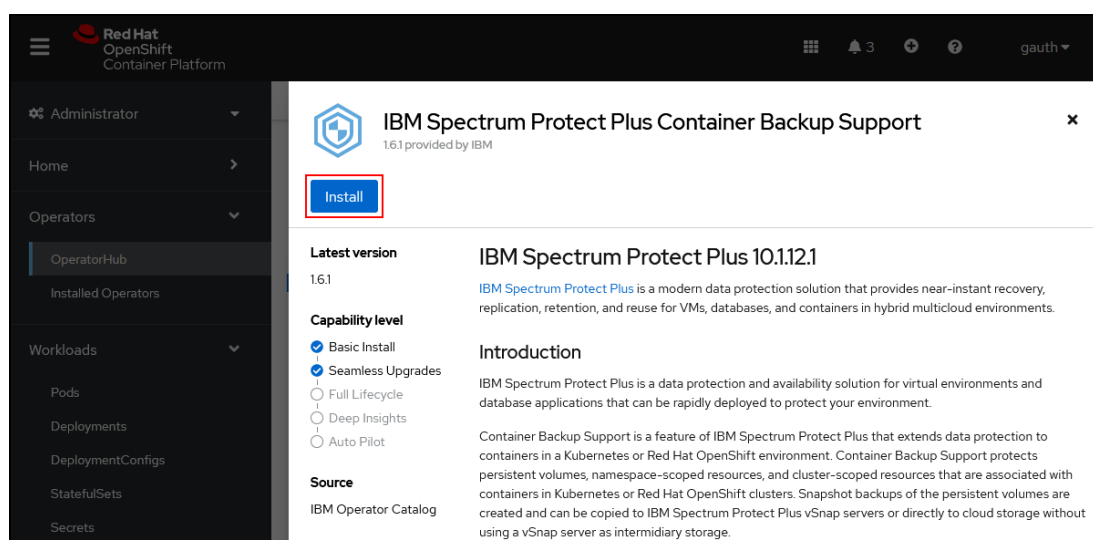


Figure 5-15 Installing the Container Backup Support operator from the Red Hat OpenShift web console part 1

3. In Install Operator window, select **v1.6**, the **baas** namespace as the destination, and the **Automatic** update strategy, as shown in Figure 5-16. Then, click **Install**.

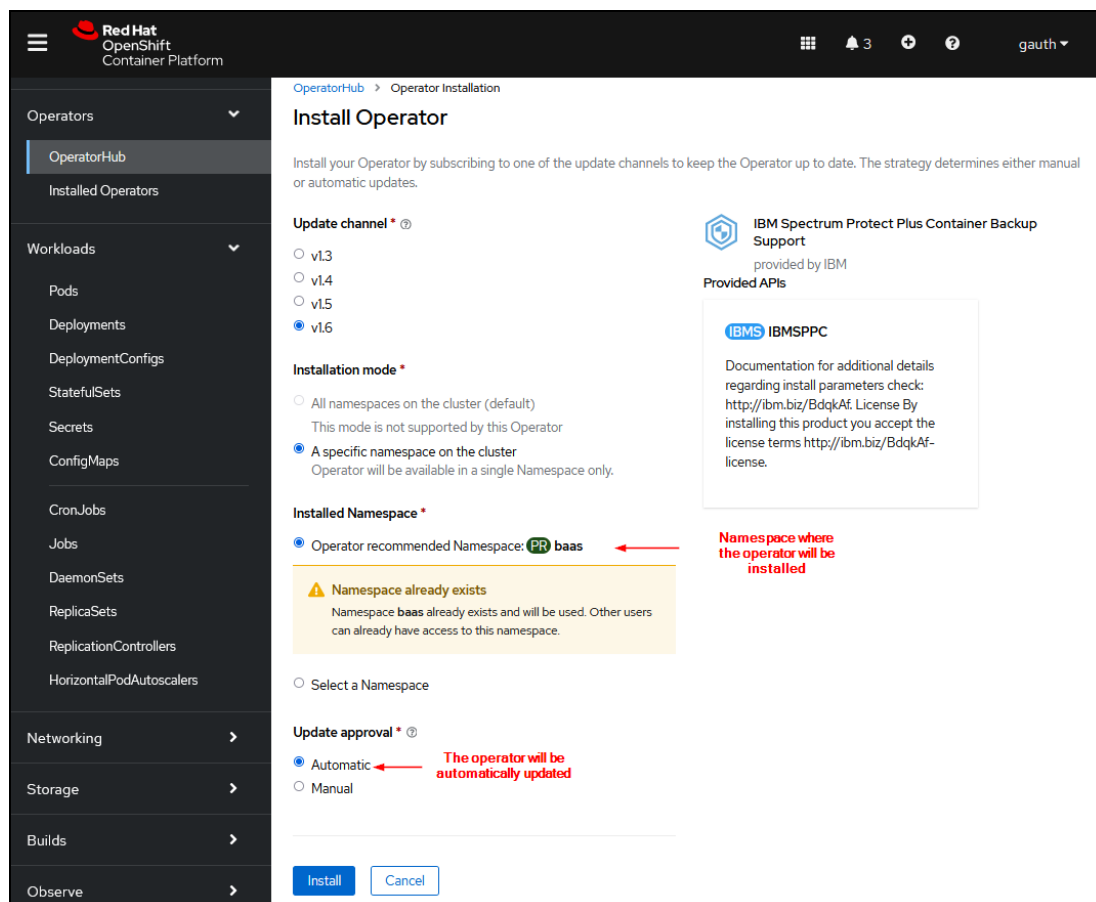


Figure 5-16 Installing the Container Backup Support operator from the Red Hat OpenShift web console part 2

4. The operator installation starts, and it might take a few minutes. After the operation completes, click **View Operator** to go to the operator window, as shown in Figure 5-17.

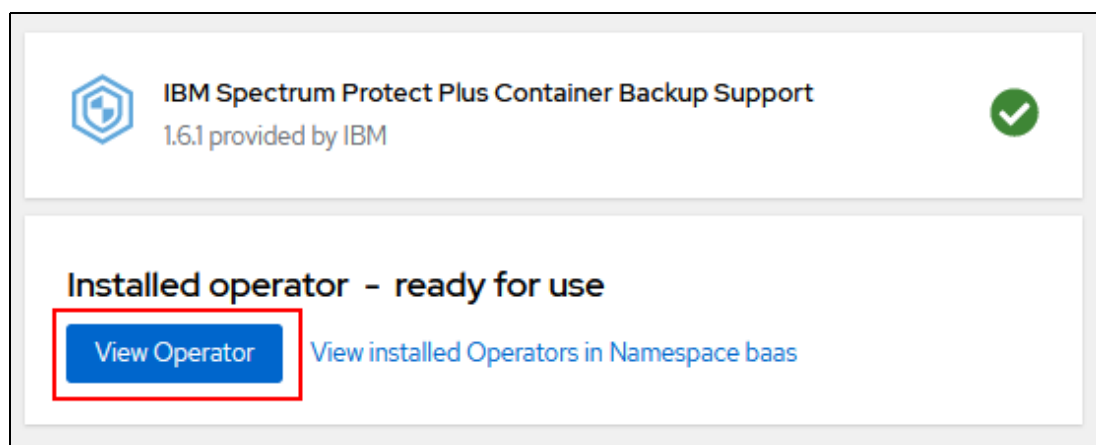


Figure 5-17 Container Backup Support installation

- From the Red Hat OpenShift web console, select **Operators** → **Installed Operator** and select the baas. You see the Container Backup Support operator along with Red Hat Red Hat OpenShift API for Data Protection and AMQ Stream, which were installed by the Container Backup Support operator.

5.2.10 Creating a Container Backup Support (BaaS) instance

After the Container Backup Support operator, along with the Red Hat OpenShift API for Data Protection and AMQ Stream operators, are installed, you can deploy Container Backup Support.

To perform the deployment, use the `baas-value-cr.yaml` that was created in Table 5-1 on page 89, which describes the Container Backup Support instance.

Using the CLI

Because the `baas-value-cr.yaml` file is prepared, apply the file to deploy the Container Backup Support instance, as shown in Example 5-20.

Example 5-20 Deploying the Container Backup Support instance by using the CLI

```
$ oc apply -f ./baas-values-cr.yaml
subscription.operators.coreos.com/baas-operator created
```

This command creates the different resources that are required for Container Backup Support in the baas project.

After a few minutes, if the deployment is successful, all the required pods are running in the baas project, as shown in Example 5-21.

Example 5-21 Checking the Container Backup Support pods' status

```
$ oc get pod -n baas
```

NAME	READY	STATUS	RESTARTS	AGE
amq-streams-cluster-operator-v2.2.0-1-5479-2hf	1/1	Running	0	19d
baas-entity-operator-674d845745-jt25g	3/3	Running	0	19d
baas-kafka-0	1/1	Running	0	19d
baas-minio-0	1/1	Running	0	19d
baas-scheduler-bddc8bb6f-2pvn9	1/1	Running	0	19d
baas-spp-agent-5c8b6c5ddb-vz4dg	1/1	Running	0	19d
baas-transaction-manager-655d589bc8-kc4cp	3/3	Running	0	19d
baas-transaction-manager-655d589bc8-tdvn7	3/3	Running	2 (19d ago)	19d
baas-transaction-manager-655d589bc8-vcwdw	3/3	Running	0	19d
baas-zookeeper-0	1/1	Running	0	19d
baas-zookeeper-1	1/1	Running	0	19d
baas-zookeeper-2	1/1	Running	0	19d
ibmsppc-operator-controller-manager-as5c-nvbcf	2/2	Running	0	20d
openshift-adp-controller-manager-59dd86-pt7	1/1	Running	0	20d
velero-b87f65b7-wwrvj	1/1	Running	0	19d

Using the Red Hat OpenShift web console

To create a Container Backup Support (BaaS) instance by using the Red Hat OpenShift web console, complete the following steps:

1. From the Red Hat OpenShift web console, select **Operators** → **Installed Operators**, and select the baas project. Click the IBM Spectrum Protect Plus Container Backup Support operator, as shown in Figure 5-18.

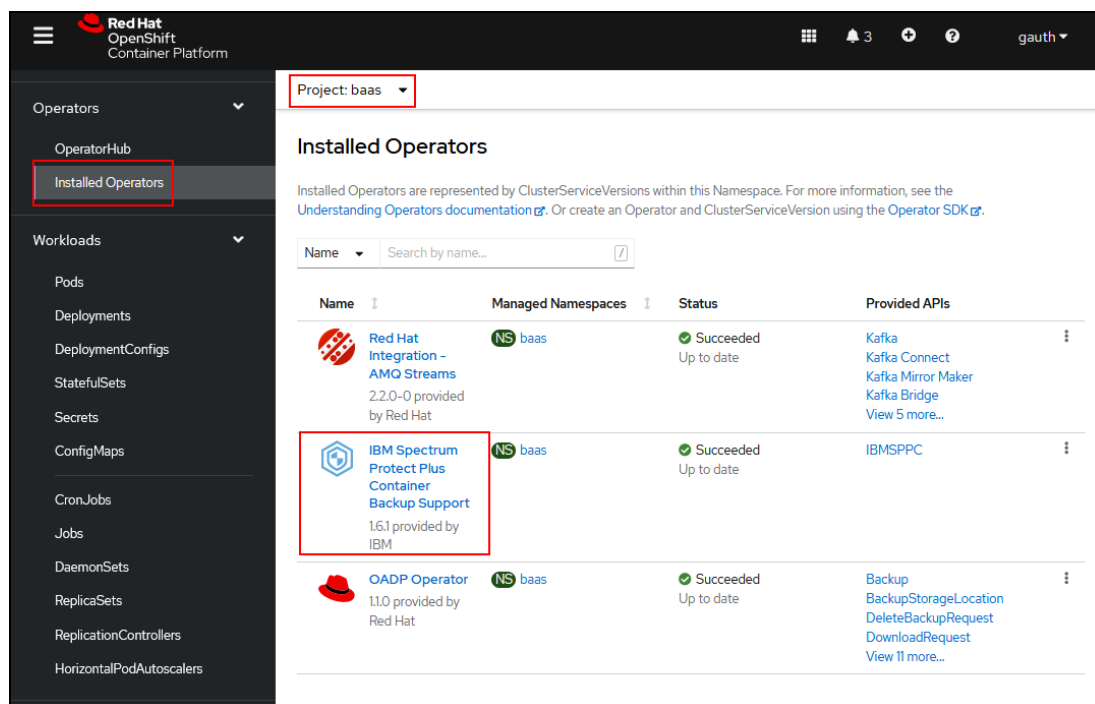


Figure 5-18 Going to the Container Backup Support operator by using the Red Hat OpenShift web console

2. From the IBM Spectrum Protect Plus Container Backup Support operator window, click **Create Instance** in the **IBMSPPC** tile, as shown in Figure 5-19 on page 113.

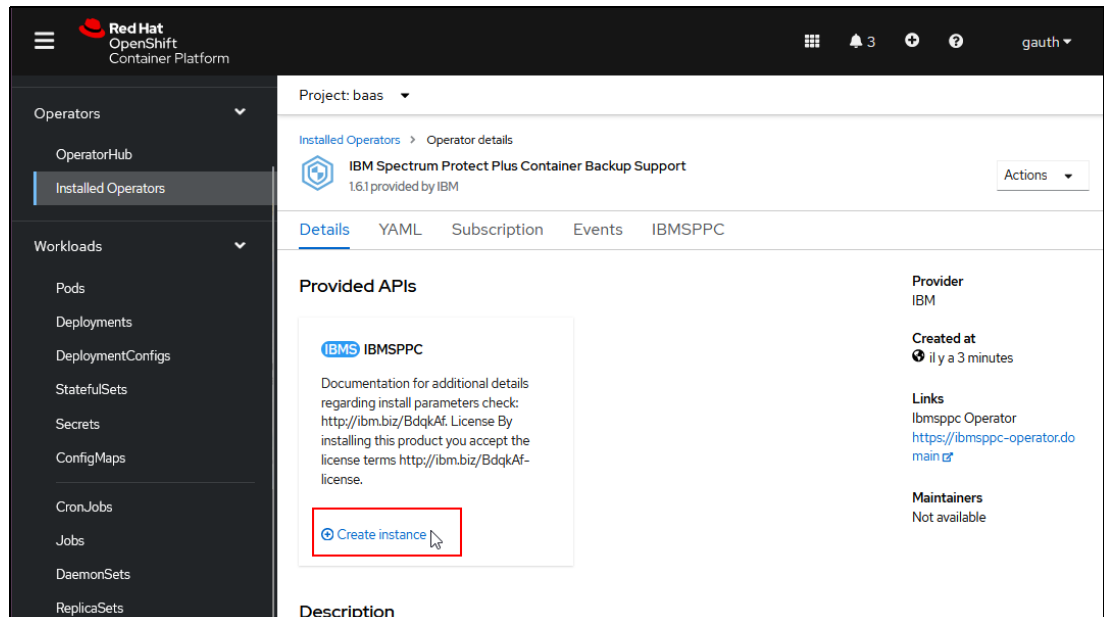


Figure 5-19 Going to the instance creation window by using the Red Hat OpenShift web console

3. From the Create IBMSPPC window, click **YAML view**, and replace the YAML with the content of `baas-values-cr.yaml`, as shown in Figure 5-20.

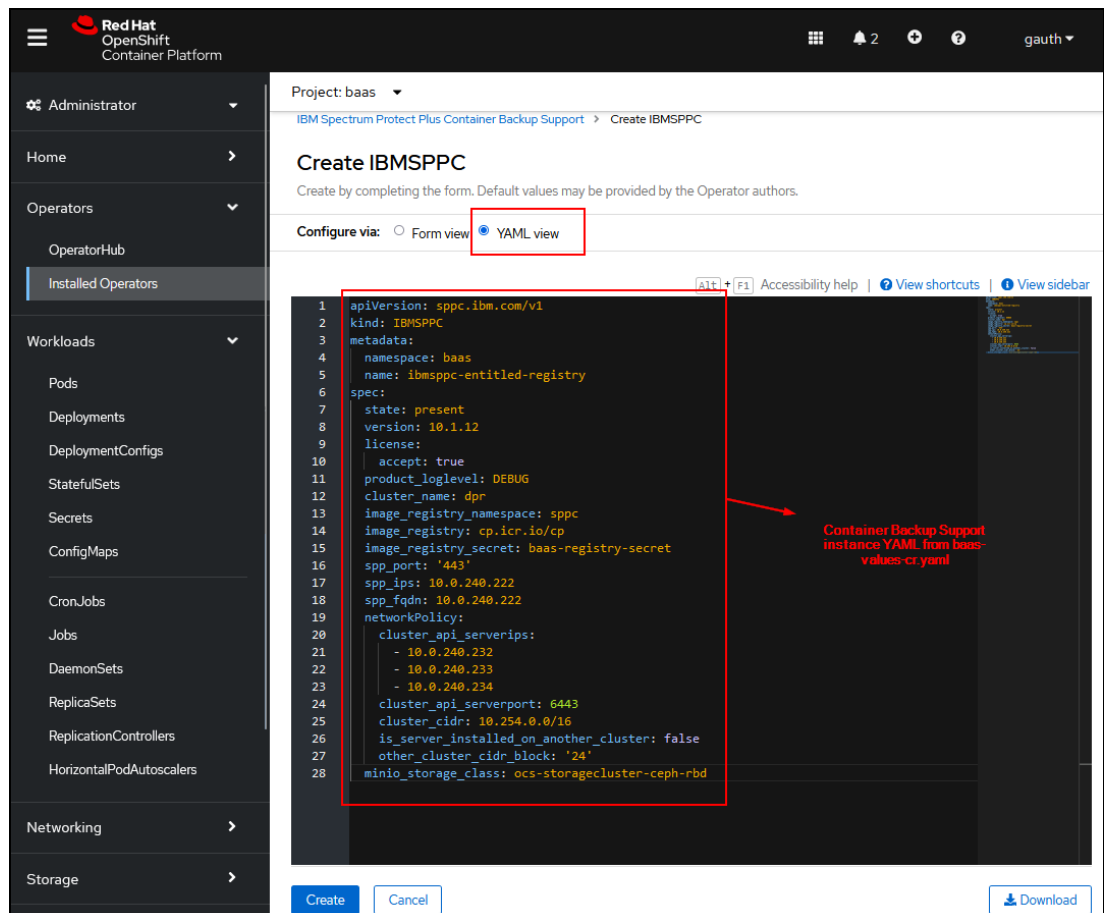


Figure 5-20 Creating the Container Backup Support instance from the Red Hat OpenShift web console

Tip: Make sure that '443' is written with quotation marks, and not as 443 because the system is expecting a string and not a number.

4. Click **Create**.

You are redirected to the **IBMSPPC** tab in the IBM Spectrum Protect Plus Container Backup Support operator window.

When the newly created instance's status changes to Condition Running, as shown in Figure 5-21, the Container Backup Support is correctly installed.

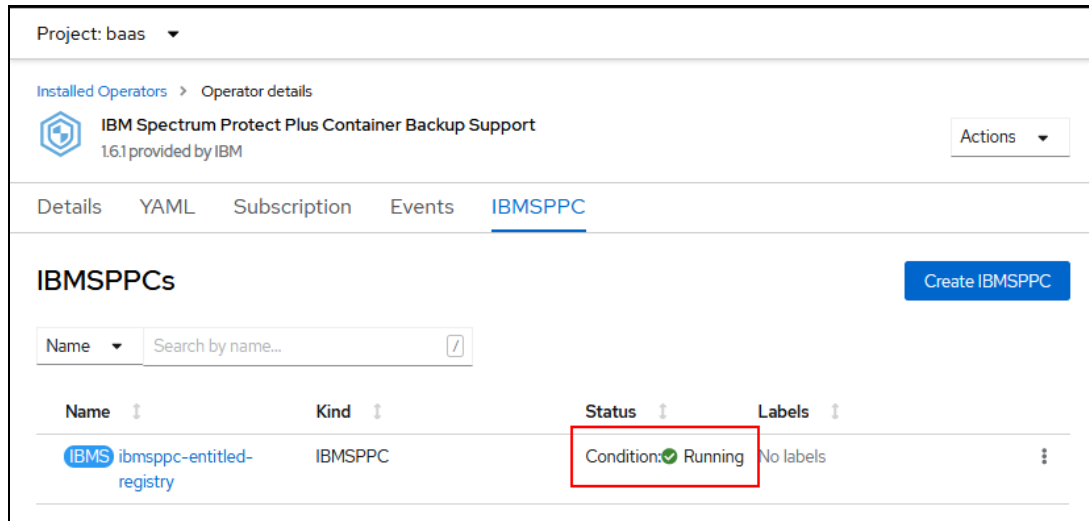


Figure 5-21 Monitoring the Container Backup Support status from the Red Hat OpenShift web console

5. Select **Workloads** → **pods** window and select the baas project. You can verify that all the Container Backup Support pods are running, as shown in Figure 5-22 on page 115.

Project: baas

Pods

Name	Status	Ready	Restarts	Owner	Memory	CPU
amq-streams-cluster-operator-v2.2.0-l-547978bf59-2mxf	Running	1/1	0	amq-streams-cluster-operator-v2.2.0-l-547978bf59	327,2 MiB	0,017 cores
baas-entity-operator-674d845745-jt25g	Running	3/3	0	baas-entity-operator-674d845745	532,9 MiB	0,014 cores
baas-kafka-0	Running	1/1	0	baas-kafka	729,9 MiB	0,015 cores
baas-minio-0	Running	1/1	0	baas-minio	219,8 MiB	0,003 cores
baas-scheduler-bddc8bb6f-2pvn9	Running	1/1	0	baas-scheduler-bddc8bb6f	17,2 MiB	0,001 cores
baas-spp-agent-5c8b6c5ddb-vz4dg	Running	1/1	0	baas-spp-agent-5c8b6c5ddb	495,5 MiB	0,001 cores
baas-transaction-manager-655d589bc8-kt4cp	Running	3/3	0	baas-transaction-manager-655d589bc8	421,8 MiB	0,008 cores
baas-transaction-manager-655d589bc8-tdvn7	Running	3/3	2	baas-transaction-manager-655d589bc8	426,9 MiB	0,008 cores
baas-transaction-manager-655d589bc8-vcwdw	Running	3/3	0	baas-transaction-manager-655d589bc8	427,9 MiB	0,008 cores
baas-zookeeper-0	Running	1/1	0	baas-zookeeper	645,9 MiB	0,008 cores
baas-zookeeper-1	Running	1/1	0	baas-zookeeper	644,5 MiB	0,007 cores
baas-zookeeper-2	Running	1/1	0	baas-zookeeper	640,5 MiB	0,008 cores
ibmsppc-operator-controller-manager-9c655855c-nvbcf	Running	2/2	0	ibmsppc-operator-controller-manager-9c655855c	2,344,2 MiB	1,139 cores
openshift-adp-controller-manager-59dd87cf96-pbft7	Running	1/1	0	openshift-adp-controller-manager-59dd87cf96	38,5 MiB	0,004 cores
velero-b87f65b7-wrvvj	Running	1/1	0	velero-b87f65b7	185,1 MiB	0,005 cores

Figure 5-22 Monitoring the Container Backup Support pods status from the Red Hat OpenShift web console

5.2.11 Registering the Red Hat OpenShift cluster in IBM Spectrum Protect Plus server manually

During the Container Backup support instance deployment, the Red Hat OpenShift cluster normally self-registers in the IBM Spectrum Protect Plus server.

However, this task might fail; if so, the Red Hat OpenShift cluster must be added manually.

To add the cluster manually, you must complete the following steps, which are described in the following sections:

1. Retrieving the Container Back Support host address
2. Retrieving the Container Backup Support host address
3. Retrieving the Container Back Support certificate
4. Registering the Red Hat OpenShift cluster
5. Retrieving the Container Back Support certificate
6. Registering the Red Hat OpenShift cluster

Retrieving the Container Backup Support host address

If IBM Spectrum Protect Plus server and the Container Backup Support are running on the same Red Hat OpenShift cluster, the host address is `baas-rest-spp-agent.baas.svc`. Otherwise, you can retrieve the Container Backup Support host address by running the following command, as shown in Example 5-22:

```
oc get route baas-spp-agent-route --namespace baas -ojsonpath="{.spec.host}"
```

Example 5-22 Retrieving the Container Backup Support host address

```
$ oc get route baas-spp-agent-route --namespace baas -o  
jsonpath="{.spec.host}"  
baas-spp-agent-route-baas.apps.ocp4.dpr.escc.1ab
```

Retrieving the Container Back Support certificate

If the IBM Spectrum Protect Plus server and Container Backup Support are running on the same Red Hat OpenShift cluster, this step is not necessary.

On your Red Hat OpenShift cluster, if you are using the default configuration, you can retrieve the required certificate by running the following command, as shown in Example 5-23:

```
oc get secret -n openshift-ingress-operator router-ca -o  
jsonpath='{.data.tls.crt}' | base64 -d
```

Example 5-23 Retrieving the certificate when you use a default Red Hat OpenShift configuration

```
$ oc get secret -n openshift-ingress-operator router-ca -o  
jsonpath='{.data.tls.crt}' | base64 -d  
-----BEGIN CERTIFICATE-----  
MIIDDDCCAFsgAwIBAgIBATANBgkqhkiG9w0BAQsFADAmMSQwIgYDVQQDDBtpbmdy  
ZXNzLW9wZXJhdG9yQDE2NjU0MTkwNzAwHhcNMjIxMDEwMTYyNDI5WhcNMjQxMDA5  
MTYyNDMwWjAmMSQwIgYDVQQDDBtpbmdyZXNzLW9wZXJhdG9yQDE2NjU0MTkwNzAw  
ggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQL3faPyeq2VW6DbwN1bJe  
+SnFofrXGc0tBc0ph5egTZ+IzisjVeHxoxS9LUHb07dRBjA1JoX3BbX3YeEYoJ0m  
rhpFnzOTHumWLC8+pNEAgsmJlB5d60xtWfZ0Fo3jW6Jv07t09d7S0eux10xdvXYi  
OWkBa/R/nQkwxEOJ9WygYviXRDZk9AeFlAMg/JQu/VHjekKabmLQZDog7Ug/zsBE  
QzJ7+U3dipviAhe+lsIaowi0Y2BFz8iT+OdrfGG/kqDfb0B2QavUb+RneuyrPXp1  
G6o10bQpXW3C74LpVIFbkPIo0VeIQVlUqCyZnuSjMH5vKYOT+9PZja2ZUstC3NyV  
AgMBAAGjRTBDMA4GA1UdDwEB/wQEAwICpDASBgNVHRMBAf8ECDAGAQH/AgEAMBOG  
A1UdDgQWBQB3/WdhNtb1z8SI9HbgIyxt9HBtaJANBgkqhkiG9w0BAQsFAAOCAQEA  
wumL0qVb3kp5X8jQ1QKJ3n9Av1egwjJhmKqxt+UteBGr56cGjWQdZV0tWRIJBMf  
PRo512TWlGZYwQ6alogFCGZPFmm8shilwixxG9vCnR867hDiXB1HeCXBhFqPErld  
EIZvWIhurzDrvJbS0v2alRq0Y5uW0CczpXmDF01m5Ebv4vsheSUHGboQheWSMULK  
V9Dexpey7sRX47ZA4nBFUZ/djgT5maqxDERkq5MLA51V3K/Z/yZe133e4q9r  
mQwj3c9x88XI/N2avA2U4CG/C4hXSkZvSdvGXC1pUQ53SK4bP2/wsT+fj6t7ThH  
vrrQHxqjNTH1IEek8em9Zw==  
-----END CERTIFICATE-----
```

On your Red Hat OpenShift cluster, if you are using a custom certificate for external communication, for example, `server-secret` in the namespace `openshift-ingress`, you can retrieve the required certificate by running the following command, as shown in Example 5-24 on page 117:

```
oc get secret -n openshift-ingress server-secret -o jsonpath='{.data.tls.crt}' |  
base64 -d
```

Example 5-24 Retrieving the certificate when using a custom Red Hat OpenShift configuration

```
$ oc get secret -n openshift-ingress server-secret -o jsonpath='{.data.tls\.crt}'  
| base64 -d  
-----BEGIN CERTIFICATE-----  
MIIDDDCCAfSgAwIBAgIBATANBgkqhkiG9w0BAQsFADAmMSQwIgYDVQQDDBtpbmdy  
ZXNzLW9wZXJhdG9yQDE2NjU0MTkwNzAwHhcNMjIxMDEwMTYyNDI5WhcNMjQxMDA5  
MTYyNDMwWjAmMSQwIgYDVQQDDBtpbmdyZXNzLW9wZXJhdG9yQDE2NjU0MTkwNzAw  
ggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQL3faPyeq2VW6DbwN1bJe  
+SnFofrXGc0tBc0ph5egTZ+IzIsjVeHxoxS9LUHb07dRbJA1JoX3BbX3YeEYoJ0m  
rhpFnzoTHumWLC8+pNEAgsmJlB5d60xtWfZ0Fo3jW6Jv07t09d7S0eux10xdvXYi  
0WkBa/R/nQkWXE0J9WygYviXRZk9AeF1aMg/JQu/VHjekKabmLQZDog7Ug/zsBE  
QzJ7+U3dipviAhe+lsIaowi0Y2BFz8iT+0drfGG/kqDfb0B2QavUb+RneuyrPXp1  
G6o10bQpXW3C74LpVIfbkPIo0vcDP15svxxZnuSjMH5vKYOT+9PZja2ZUstC3NyV  
AgMBAAGjRTBDMA4GA1UdDwEB/wQEAwICpDASBgNVHRMBAf8ECDAGAQH/AgEAMBOG  
A1UdDgQWBBQ3/WdhNtb1z8SI9HbgIyxt9HBtaJANBgkqhkiG9w0BAQsFAA0CAQEA  
wumL0qVb3kp5X8jQ1QKJ3n9Av1egwjJhmKqxtV+UteBGr56cGjWQdZV0tWRIJBMf  
PRo512TWlGZYwQ6aDpcoSifhQt59Su78QadxG9vCnR867hDiXB1HeCXBhFqPEr1d  
EIZvWIhurzDrvJbS0v2a1RqOY5uW0CcZpXmDF01m5Ebv4vsheSUHGb0QheWSMULK  
V9Dexpey7sRX47ZA4nBFUZ/djgT5maqxDERKcq5MLA51V3K/Z/yZe133e4q9r  
mQwj3c9x88XI/N2avA2U4CG/C4hXSkZvSdvGCxc1pUQ53SK4bP2/wsT+fj6t7ThH  
vrrQHxqjNTH1IEek8em9Zw==  
-----END CERTIFICATE-----
```

Registering the Red Hat OpenShift cluster

To register the Red Hat OpenShift cluster, complete the following steps:

1. From the IBM Spectrum Protect Plus web console, select **Manage Protection** → **Containers** → **OpenShift**, click **Manage clusters**, and then click **Add Cluster**, as shown in Figure 5-23.

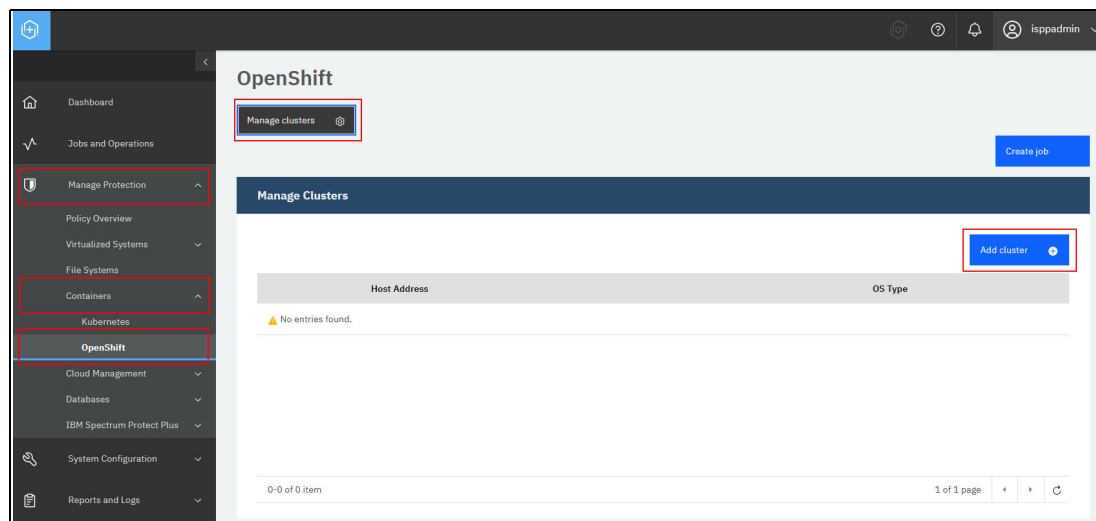


Figure 5-23 Going to the cluster registration window by using the IBM Spectrum Protect Plus web console

2. Enter the required values, as shown in Figure 5-24, by using values that are described in Table 5-6.

OpenShift

Manage clusters

Create job

Manage Clusters

Application Properties

Cluster Name: dpr

Host Address: baas-spp-agent-route-baas.apps.ocp4.d

Port Number: 443

Use existing user: ☒

Select user: dpr

Certificate: ☐ Use existing certificate ☐ Use CA certificate ☒ Copy and paste ☐ Upload

Arbitrary name of the certificate: drp-cert

Content of the certificate retrieved in the previous steps: -----BEGIN CERTIFICATE-----
MIIDODCCAFSgAwIBATANBgkqhkiG9w0BAQsFADAmMSQwIgYDVQQDDB
tpbmdy
ZXNzLW9wZXJhdG9yQDE2NjU0MTkwNzAwHhcNMjExMDExMTYyNDI5WmcNM
jQxMDA5
MTYyNDMwWjAmMSQwIgYDVQQDDBtpbmdyZXNzLW9wZXJhdG9yQDE2NjU0
MTkwNzAw

Options

Maximum concurrent PVCs: 10

Cancel Save

Figure 5-24 Adding a Red Hat OpenShift cluster by using the IBM Spectrum Protect Plus web console

Table 5-6 Red Hat OpenShift form fields description

Field	Value
Cluster Name	The Red Hat OpenShift cluster name that is specified in baas-value-cr.yaml.
Host Address	The Container Backup Support host address that was retrieved in "Retrieving the Container Backup Support host address" on page 116.
Port Number	443.
Use Existing User	Checked.
Select User	The user that corresponds to the data mover that is specified in baas-options.sh. This user was created during the Container Backup Support deployment.
Certificate	Select Copy and paste .

Field	Value
Certificate Name	An arbitrary name to save the certificate.
Certificate Content	The certificate itself, which was retrieved in “Retrieving the Container Back Support certificate” on page 116.

Important: After setting the certificate’s name and content, do not forget to click **Create** to save the certificate.

3. Click **Save** to register the cluster. After the Red Hat OpenShift cluster registration finishes, you can run an inventory to gather information from this cluster, as shown in Figure 5-25.

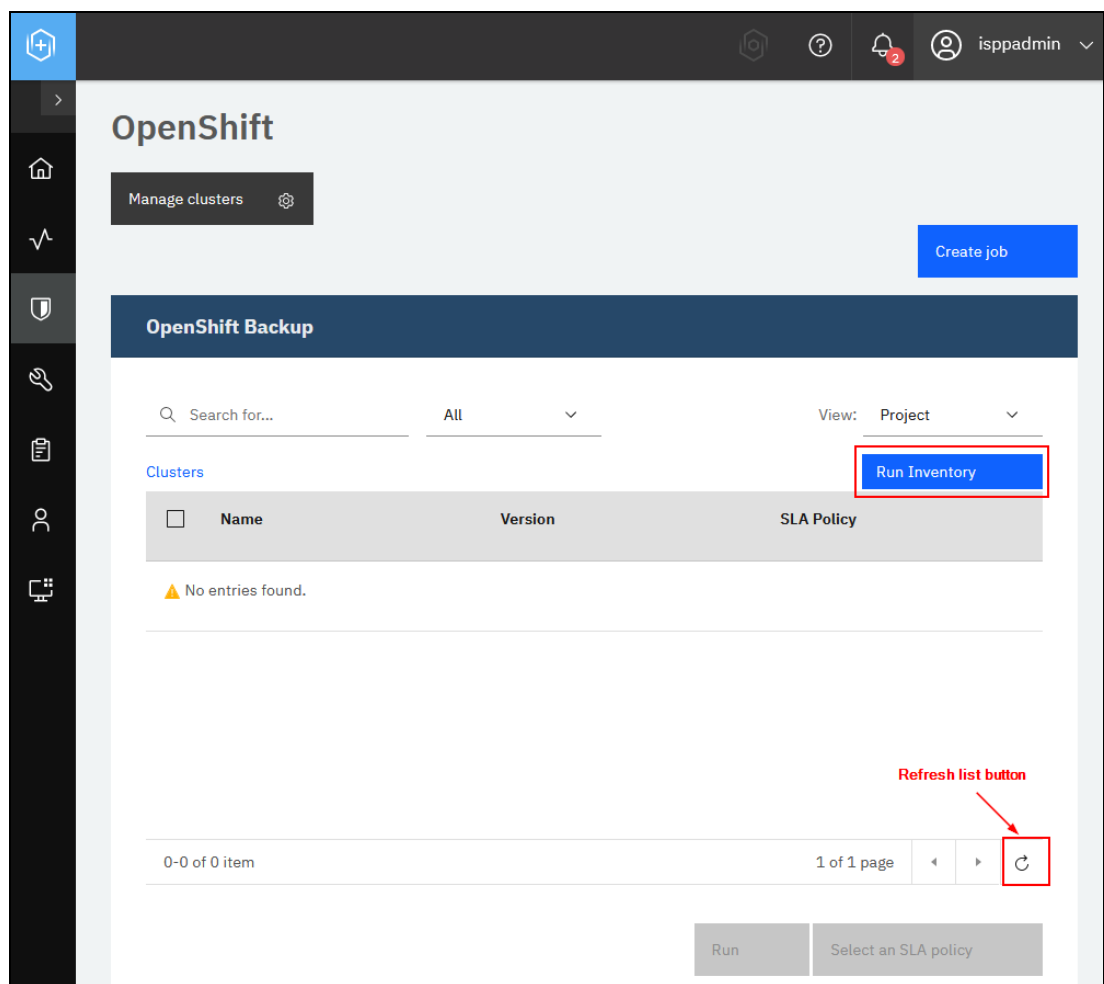


Figure 5-25 Running an inventory job for the Red Hat OpenShift cluster from the IBM Spectrum Protect Plus web console

5.2.12 Updating Container Backup Support

When updating the IBM Spectrum Protect Plus server, you might need to update other components, including Container Backup Support. To do so, click **Installed Operators**, select the IBM Spectrum Protect Plus Container Backup Support operator, click the **Subscription** tab, and click **Update Channel**, as shown in Figure 5-26.

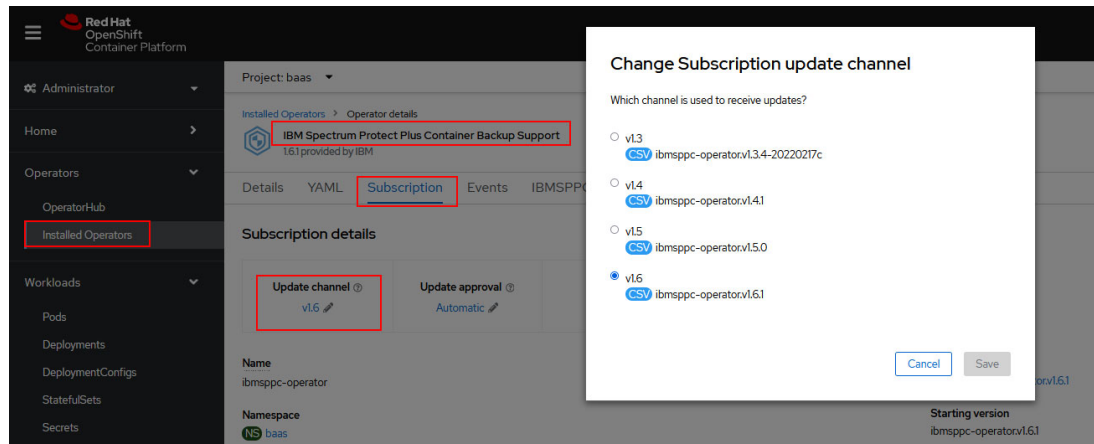


Figure 5-26 Changing the Container Backup Support Update Channel

However, depending on which version you are upgrading, some extra tasks might be required before the update. For an exhaustive description of these tasks by version, see [Updating Container Backup Support operator and instance](#).

5.3 Installing in an air-gapped environment

In some situations, the Red Hat OpenShift cluster is deployed in a restricted network without connectivity to internet. If so, performing an online installation is impossible because the Red Hat OpenShift cluster cannot pull the required images. Instead, the cluster must pull images from a local registry where the required images are stored.

To perform an air-gapped installation, see [Installing Container Backup Support on Red Hat OpenShift by using a CASE package in an air-gapped environment](#).



Using Container Backup Support

This chapter provides information about how to use the IBM Spectrum Protect Plus backup as a service (BaaS) component in a Red Hat OpenShift Container environment.

This chapter includes the following topics:

- ▶ 6.1, “Sample application” on page 122
- ▶ 6.2, “Application-consistent backups” on page 127
- ▶ 6.3, “Creating a service-level agreement policy” on page 130
- ▶ 6.4, “Red Hat OpenShift Container Platform backups” on page 136
- ▶ 6.5, “Performing restores” on page 142

6.1 Sample application

To demonstrate the container backup and restore capabilities of IBM Spectrum Protect Plus, you need a sample application as the base. In our example, we use a sample WordPress application. This section describes how you can create an instance of the sample application.

Complete the following steps:

1. From a command shell with access to your Red Hat OpenShift cluster, log in to the cluster with an account that has Red Hat OpenShift cluster administrator privileges. The login mechanism for your cluster might use a different authentication mechanism, but in this example the cluster is set up to allow the user to log in with a token, as shown in Example 6-1.

Example 6-1 Logging in to the cluster

```
oc login --token=<token> --server=https://<cluster api url>:6443
```

2. Create a project in the cluster for your WordPress application. In the examples that follow, we use `spp-wordpress` for the name of our project, as shown in Example 6-2.

Example 6-2 Creating a project to contain our sample application

```
oc new-project spp-wordpress
```

3. Before you create an application, set up the environment variable for installation, as shown in Example 6-3. This environment variable is used in the following scripts.

Example 6-3 Setting up the environment variable

```
export WORDPRESS_NAMESPACE=spp-wordpress
```

4. The WordPress application needs a secret, so create a secret by cutting and pasting the code in Example 6-4 to the command shell where you are logged in to Red Hat OpenShift.

Example 6-4 Creating the secret

```
cat <<EOF |oc apply -f -
apiVersion: v1
kind: Secret
metadata:
  name: mysql-pass
  namespace: ${WORDPRESS_NAMESPACE}
type: Opaque
data:
  password: Passw0rd
EOF
```

5. Deploy the `mysql` deployment that is used as the back end for the WordPress application by cutting and pasting the code that is in Example 6-5 to the command shell.

Example 6-5 Creating the mysql deployment

```
cat <<EOF |oc apply -f -
apiVersion: v1
kind: Service
metadata:
  name: wordpress-mysql
  namespace: ${WORDPRESS_NAMESPACE}
```

```

    labels:
      app: wordpress
spec:
  ports:
    - port: 3306
  selector:
    app: wordpress
    tier: mysql
  clusterIP: None
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: mysql-pvc
  namespace: ${WORDPRESS_NAMESPACE}
  labels:
    app: wordpress
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 2Gi
---
apiVersion: apps/v1
kind: Deployment
metadata:
  name: wordpress-mysql
  namespace: ${WORDPRESS_NAMESPACE}
  labels:
    app: wordpress
spec:
  selector:
    matchLabels:
      app: wordpress
      tier: mysql
  strategy:
    type: Recreate
  template:
    metadata:
      labels:
        app: wordpress
        tier: mysql
    spec:
      containers:
        - image: mysql:5.6
          name: mysql
          env:
            - name: MYSQL_ROOT_PASSWORD
              valueFrom:
                secretKeyRef:
                  name: mysql-pass
                  key: password
          ports:
            - containerPort: 3306

```

```

        name: mysql
    volumeMounts:
    - name: mysql-persistent-storage
      mountPath: /var/lib/mysql
    volumes:
    - name: mysql-persistent-storage
      persistentVolumeClaim:
        claimName: mysql-pvc
EOF

```

EOF

-
6. Create the wordpress deployment by cutting and pasting the code in Example 6-6 to the command shell.

Example 6-6 Creating the wordpress deployment

```

cat <<EOF |oc apply -f -
apiVersion: v1
kind: Service
metadata:
  name: wordpress
  namespace: ${WORDPRESS_NAMESPACE}
  labels:
    app: wordpress
spec:
  ports:
  - port: 80
  selector:
    app: wordpress
    tier: frontend
  type: ClusterIP
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: wordpress-pvc
  namespace: ${WORDPRESS_NAMESPACE}
  labels:
    app: wordpress
spec:
  accessModes:
  - ReadWriteOnce
  resources:
    requests:
      storage: 1Gi
---
apiVersion: apps/v1
kind: Deployment
metadata:
  name: wordpress
  namespace: ${WORDPRESS_NAMESPACE}
  labels:
    app: wordpress
spec:
  selector:
    matchLabels:
      app: wordpress

```

```

    tier: frontend
strategy:
  type: Recreate
template:
  metadata:
    labels:
      app: wordpress
      tier: frontend
  spec:
    containers:
    - image: wordpress:4.8-apache
      name: wordpress
      env:
      - name: WORDPRESS_DB_HOST
        value: wordpress-mysql
      - name: WORDPRESS_DB_PASSWORD
        valueFrom:
          secretKeyRef:
            name: mysql-pass
            key: password
      ports:
      - containerPort: 80
        name: wordpress
      volumeMounts:
      - name: wordpress-persistent-storage
        mountPath: /var/www/html
    volumes:
    - name: wordpress-persistent-storage
      persistentVolumeClaim:
        claimName: wordpress-pvc

```

EOF

-
7. The application is deployed. Ensure that the two pods for the application are running by running the command that is shown in Example 6-7.

Example 6-7 Check that the wordpress pods are running

```

oc project spp-wordpress
oc get pods

```

8. The output should be like what is shown in Example 6-8. The last part of the names of the pods differ.

Example 6-8 The wordpress pods are running

NAME	READY	STATUS	RESTARTS	AGE
wordpress-6755c4ccdd-ptplq	1/1	Running	0	5d16h
wordpress-mysql-74c97488d6-trcdl	1/1	Running	0	5d16h

9. By default, the wordpress route is not created, so create the route by running the command in Example 6-9.

Example 6-9 Exposing the service

```

oc expose svc wordpress

```

10.A new route is created. Details of the route can be obtained by running the command that is shown in Example 6-10.

Example 6-10 Getting the route details for the WordPress application

```
oc get route
```

The output is to what is shown in Example 6-11.

Example 6-11 Route details

NAME	HOST/PORT	PATH	SERVICES	PORT	TERMINATION	WILDCARD
wordpress	wordpress-spp-wordpress.apps.ocp4.dpr.escc.1ab		wordpress	80		None

11.Check that the WordPress application is correctly deployed by accessing the HOST/PORT URL that is defined in the route that we created. You should see the WordPress window, as shown in Figure 6-1.




Figure 6-1 Initial window of the WordPress site

6.2 Application-consistent backups

With the advent of cloud-native applications, microservices architecture, and containerization, applications might be composed of a multitude of moving parts, each potentially with its own storage technology. In order for a backup to be taken, the data for the whole application must be consistent. IBM Spectrum Protect Plus includes the open-source Velero distribution that provides *hooks*, which allow application developers to ensure that the data of an application is in a good state before a backup is taken.

6.2.1 Velero hooks

To support application-consistent backups, use the capabilities of Velero to annotate pods with backup pre- and post-hooks so that you can put the application into a state where a consistent backup is possible:

► Pre-hooks

- `pre.hook.backup.velero.io/container`

The container where the command should be run. Defaults to the first container in the pod. Optional.

- `pre.hook.backup.velero.io/command`

The command to run. If you need multiple arguments, specify the command as a JSON array, such as `["/usr/bin/uname", "-a"]`.

- `pre.hook.backup.velero.io/on-error`

Action to take if the command returns a nonzero exit code. Defaults to `Fail`. Valid values are `Fail` and `Continue`. Optional.

- `pre.hook.backup.velero.io/timeout`

How long to wait for the command to run. The hook is considered in error if the command exceeds the timeout. Defaults to 30s. Optional.

► Post-hooks

- `post.hook.backup.velero.io/container`

The container where the command should be run. Defaults to the first container in the pod. Optional.

- `post.hook.backup.velero.io/command`

The command to run. If you need multiple arguments, specify the command as a JSON array, such as `["/usr/bin/uname", "-a"]`.

- `post.hook.backup.velero.io/on-error`

Action to take if the command returns a nonzero exit code. Defaults to `Fail`. Valid values are `Fail` and `Continue`. Optional.

- `post.hook.backup.velero.io/timeout`

How long to wait for the command to run. The hook is considered in error if the command exceeds the timeout. Defaults to 30s. Optional.

For more information about Velero hooks, see [Backup Hooks](#).

6.2.2 Configuring the sample application with backup hooks

The example WordPress application contains a MySQL database. This database uses a persistent volume to store the database data. We create Velero hooks to ensure that the persistence volume is in an application consistent state while we create snapshots of the volume.

Note: This example is a simple one to demonstrate the usage of hooks. It is the application developers responsibility to provide the detailed hooks that permit an application consistent backup.

Complete the following steps:

1. Log in to the Red Hat OpenShift cluster as a cluster administrator. Set up the environment for creating the hooks by running the commands that are shown in Example 6-12. The WordPress pod contains a single container called `mysql`.

Example 6-12 Setting up the environment for creating the hooks

```
oc project wordpress
export WORDPRESS_MYSQL_POD=$(oc get pods | grep wordpress-mysql | awk '{ print $1 }' )
export MYSQL_CONTAINER=mysql
```

2. Create the pre-hooks by running the commands that are shown in Example 6-13. We set up a command that flushes and locks the tables. If the command fails, we fail the backup. We set a timeout of 30 seconds for the command to run, which is the default and can be omitted.

Example 6-13 Velero pre-hooks

```
oc annotate pod ${WORDPRESS_MYSQL_POD}
pre.hook.backup.velero.io/container=${MYSQL_CONTAINER}
oc annotate pod ${WORDPRESS_MYSQL_POD}
pre.hook.backup.velero.io/command='["/bin/bash", "-c", "mysql
--password=${MYSQL_ROOT_PASSWORD} -e \"FLUSH TABLES WITH READ LOCK\""]'
oc annotate pod ${WORDPRESS_MYSQL_POD} pre.hook.backup.velero.io/on-error=Fail
oc annotate pod ${WORDPRESS_MYSQL_POD} pre.hook.backup.velero.io/timeout=30s
```

3. Set up the post-hooks by running the commands that are shown in Example 6-14. These hooks will run.

Example 6-14 Velero post-hooks

```
oc annotate pod ${WORDPRESS_MYSQL_POD} post.hook.backup.velero.io/container=mysql
oc annotate pod ${WORDPRESS_MYSQL_POD}
post.hook.backup.velero.io/command='["/bin/bash", "-c", "mysql
--password=${MYSQL_ROOT_PASSWORD} -e \"UNLOCK TABLES\""]'
oc annotate pod ${WORDPRESS_MYSQL_POD} post.hook.backup.velero.io/on-error=Fail
oc annotate pod ${WORDPRESS_MYSQL_POD} post.hook.backup.velero.io/timeout=30s
```

4. You can check whether the annotations were applied correctly by running the command that is shown in Example 6-15.

Example 6-15 Getting the YAML for the WordPress pod

```
oc get pods ${WORDPRESS_MYSQL_POD} -o yaml
```

The output should contain the details of the newly applied hooks in the annotations section of the pod, as shown in Example 6-16.

Example 6-16 The WordPress mysql pod with annotations created

```
kind: Pod
metadata:
  annotations:
    k8s.v1.cni.cncf.io/network-status: |-
      [{
        "name": "openshift-sdn",
        "interface": "eth0",
        "ips": [
          "10.254.5.234"
        ],
        "default": true,
        "dns": {}
      }]
    k8s.v1.cni.cncf.io/networks-status: |-
      [{
        "name": "openshift-sdn",
        "interface": "eth0",
        "ips": [
          "10.254.5.234"
        ],
        "default": true,
        "dns": {}
      }]
    openshift.io/scc: anyuid
    post.hook.backup.velero.io/command: '["/bin/bash", "-c", "mysql
--password=$MYSQL_ROOT_PASSWORD
-e \"UNLOCK TABLES\""]'
    post.hook.backup.velero.io/container: mysql
    post.hook.backup.velero.io/on-error: Fail
    post.hook.backup.velero.io/timeout: 30s
    pre.hook.backup.velero.io/command: '["/bin/bash", "-c", "mysql
--password=$MYSQL_ROOT_PASSWORD
-e \"FLUSH TABLES WITH READ LOCK\""]'
    pre.hook.backup.velero.io/container: mysql
    pre.hook.backup.velero.io/on-error: Fail
    pre.hook.backup.velero.io/timeout: 30s
    .....
```

The application is now configured to run the hooks when we take a backup with IBM Spectrum Protect Plus. For more information, see 6.4, “Red Hat OpenShift Container Platform backups” on page 136.

6.3 Creating a service-level agreement policy

You must create a service-level agreement (SLA) policy before you back up and restore the sample application. You can manage SLAs by selecting **Manage Protection** → **Policy Overview**, as shown in Figure 6-2.

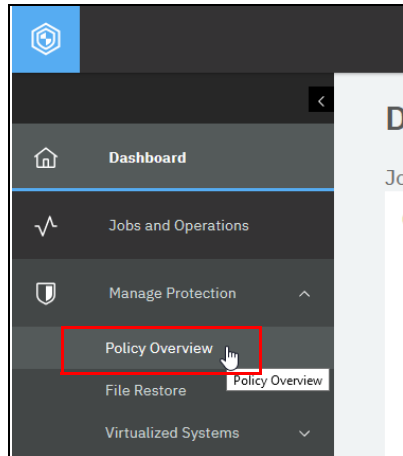


Figure 6-2 Policy Overview menu

The list of available SLAs is shown in Figure 6-3. Click **Add SLA Policy** to create a SLA.

SLA Policies					
					Add SLA Policy
<input type="text" value="Search Policy by name..."/>					
Name		Job type	Frequency		Retention
		Container	Snapshot Backup	Every 6 Hours Every 1 Days at 7:10:00 AM	1 Days 1 Months
		HANA_backup	Backup	Every 1 Days at 12:00:00 PM	3 Days
		SPPCatalogtoS3	Backup	Every 1 Days at 1:00:00 AM	2 Days
		selfprotect	Copy Backup	Every 1 Days at 5:00:00 AM Every 1 Days at 4:00:00 AM	Same retention as source selection. 5 Days
1-4 of 4 items					1 of 1 page

Figure 6-3 SLA overview

6.3.1 General parameters

In IBM Spectrum Protect Plus, different types of SLAs are available, depending on the type of application data to be backed up. In our example, we set up a Containers policy type by completing the following steps:

1. In the Add SLA Policy window, select Category drop-down menu and select **Containers**, as shown in Figure 6-4.

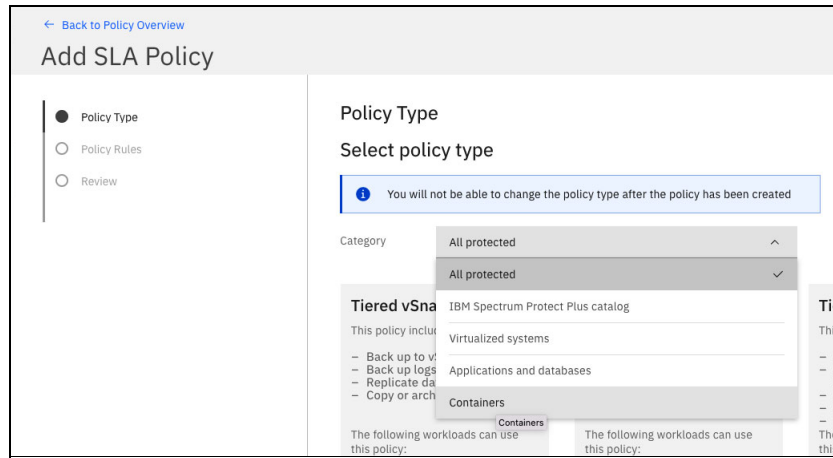


Figure 6-4 Selecting the policy type

2. A set of appropriate policy types are shown. Because we selected the **Containers** policy type, the only option is Tiered snapshot, as shown in Figure 6-5. Select this template and click **Next** to configure the policy.

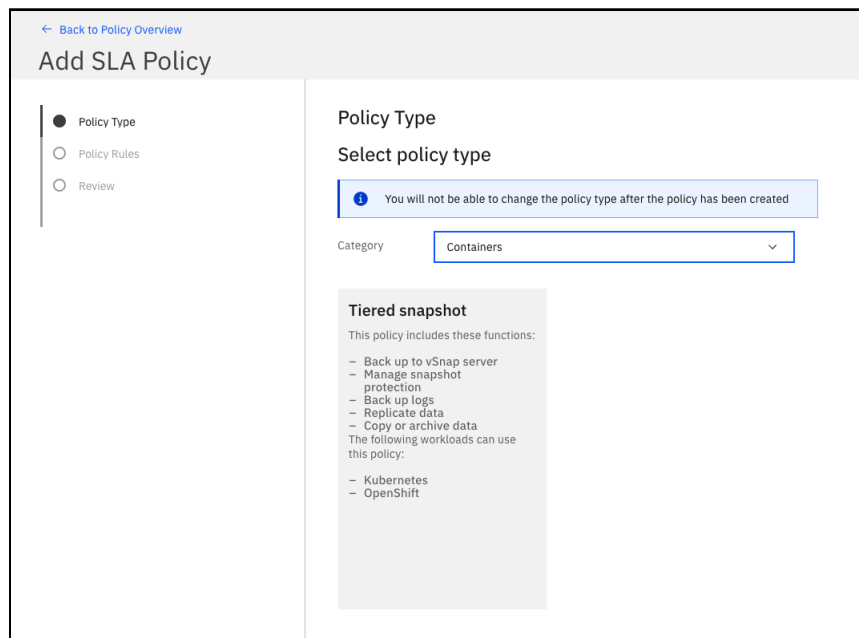


Figure 6-5 Available policy templates

- Now, configure the rules for the policy. First, define the name of the SLA. A best practice is to provide a specific name that provides information about the SLA. In our example, we defined the name `spp-containers`, as shown in Figure 6-6.

This SLA applies to containers. In our example, we create a policy that allows for an ad hoc backup, so we do not provide a schedule. This policy is fine for testing or when the snapshots or backups that use this SLA are triggered by an external scheduler by using a REST API. For production environments that use the IBM Spectrum Protect Plus, complete the details for the Scheduling capability.

The screenshot displays the 'Add SLA Policy' interface. On the left, a sidebar contains navigation icons. The main area is titled 'Add SLA Policy' and includes a 'Back to Policy Overview' link. Below the title, there are three tabs: 'Policy Type', 'Policy Rules' (which is active), and 'Review'. The 'Policy Rules' section contains the following configuration options:

- Current Policy Type:** Tiered snapshot
- Name:** spp-container-policy
- ☐ Disable all Schedules
- Snapshot Protection:**
 - Retention:** 15 Days
 - ☒ Disable Schedule
 - Repeats:** Daily
 - Every:** 1 day(s)
 - Start Time:** 11/02/2022 00:00 Europe/Amsterdam
 - Snapshot Prefix:** spp-containers

Figure 6-6 Configuring Snapshot Protection

6.3.2 Configuring a backup policy

Optionally, you can activate and define a backup policy. A backup policy backs up the resources and persistent volume claims (PVCs) so that they are stored outside the Red Hat OpenShift Container Platform cluster in the IBM Spectrum Protect Plus vSnap servers. In our example, we configure our policy to have no schedule (testing purposes only). We configure the following items:

- ▶ The target site to which you want to store the backups (primary or secondary). In our example, the target site is the primary site.
- ▶ Whether the backups are stored in encrypted vSnap servers. In our examples, we use an unencrypted vSnap server.

Figure 6-7 on page 133 shows our example backup policy.

Backup Policy

☒ Backup Storage

Retention 15 - + Days

☒ Disable Schedule

Repeats Daily Every: 1 day(s)

Start Time 11/01/2022 09:11 Europe/Amsterdam

Target Storage Type Site Target Primary

☐ Only use encrypted disk storage.

Figure 6-7 Backup policy

Important: Snapshots and backups are performed for resources and PVCs, which are triggered by one or many SLAs. However, snapshots and backups of a resource or PVC cannot occur concurrently. If two jobs address the same resource, the second job fails.

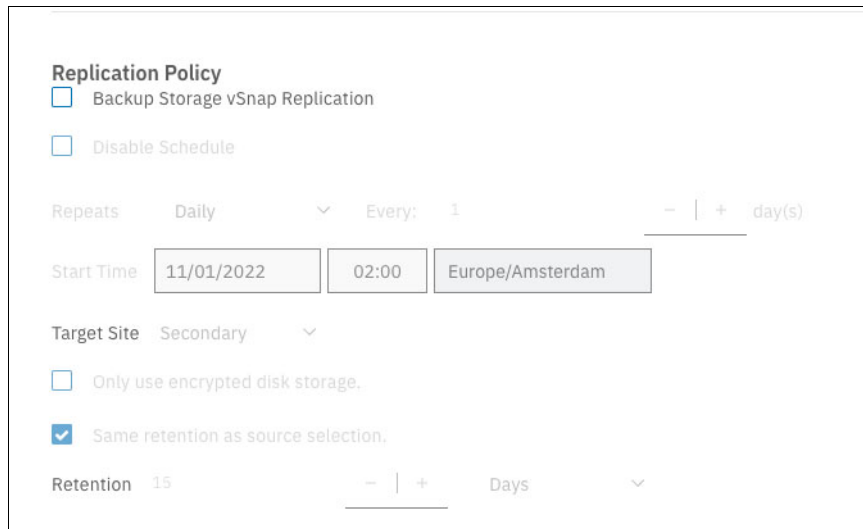
Therefore, you must be careful about defining SLA snapshots and backup schedules, and how SLAs are assigned to resources or PVCs to avoid overlap.

Important: When a backup is performed, for each resource and PVC that uses this SLA, the backup uses the latest snapshot that is available in the Red Hat Red Hat OpenShift Container Platform cluster.

Therefore, having backups more frequently than snapshots results in many identical backups. The backups can use the same latest snapshots as source.

6.3.3 Defining a replication policy

Optionally, you also can define a replication policy. A replication policy defines how resources and PVCs are backed up by a backup policy to another site. A replication policy configuration is like a backup policy, but it also provides another parameter, which is the target site (see Figure 6-8). With this parameter, you can choose the site where the target vSnaps are located. For our example, we do not create a replication policy, so leave the **Backup Storage vSnap Replication** checkbox clear.



The screenshot shows the 'Replication Policy' configuration window. It includes several settings: a checkbox for 'Backup Storage vSnap Replication' which is unchecked, a checkbox for 'Disable Schedule' which is unchecked, a 'Repeats' dropdown set to 'Daily', an 'Every' field set to '1', and a unit dropdown set to 'day(s)'. The 'Start Time' section contains three fields: '11/01/2022', '02:00', and 'Europe/Amsterdam'. The 'Target Site' dropdown is set to 'Secondary'. There are two checkboxes for storage: 'Only use encrypted disk storage.' (unchecked) and 'Same retention as source selection.' (checked). The 'Retention' field is set to '15' with a unit dropdown set to 'Days'.

Figure 6-8 Configuring a replication policy

Note: To allow a replication policy, first create the appropriate replication partnerships between the VSnap server, as described in *IBM Spectrum Protect Plus Practical Guidance for Deployment, Configuration, and Usage*, REDP-5532.

6.3.4 Additional Copies

You also can choose to copy data to a standard object store (cloud services) or to an IBM Spectrum Protect instance (repository servers). These additional copies can be used, for example, the following reasons:

- ▶ Provide long-term retention backup.
- ▶ Store backups on a low-cost tier.
- ▶ Outsource backups off-site.

Note: To create extra copies, you must enable a backup policy to vSnap.

The additional copies can be incremental or full copies. For incremental copies, the first copy is always a full copy. The next copies are incremental, so only changes since the last copy are transferred. Configuring the **Repeats** and **Retention** parameters are like the previous snapshot and backup policies. For our example, we do not use a schedule.

A source for the copy must be selected, which is either the backup policy destination or the replication policy destination. For example, because we have not configured a replication policy destination, only a single option is highlighted.

After selecting the source, select a destination for the copy. For our example, we use the S3 compatible MinIO server that we created in Chapter 5, “Implementing Container Backup Support” on page 85. The parameters for the destination are configured as follows:

- Destination: A cloud service or a repository server.
- Target: One of the cloud services or repository servers that are available.

Figure 6-9 shows Additional Copies window.

Additional Copies

☒ Standard object storage (incremental copy)

☐ Disable Schedule

Repeats: **Daily** Every: **1** day(s)

Start Time: **11/01/2022** **00:00** **Europe/Amsterdam**

☒ Same retention as source selection

Source: ☒ Backup Policy Destination ☐ Replication Policy Destination

Destination: **Cloud services**

Target: **MinIO**

Figure 6-9 Additional Copies

Note: The target should be defined before the SLA configuration by selecting **System Configuration** → **Backup Storage** → **Cloud Storage** or **Repository Server**.

For more information about more copies, see *IBM Spectrum Protect Plus Practical Guidance for Deployment, Configuration, and Usage*, REDP-5532.

6.3.5 Configuring object storage archival

To configure object storage archival, complete the following steps:

1. It is possible to configure archiving for the object storage, as shown in Figure 6-10. For our example, we do not configure this capability.

☐ Archive object storage (full copy)

Specify to run once Weekly, Monthly, or Yearly

☐ Disable Schedule

Repeats: **Weekly** Every: **1** week(s)

On: ☐ S ☐ M ☒ T ☐ W ☐ Th ☐ F ☐ Sa ☐ Su

Start Time: **11/01/2022** **00:00** **Europe/Amsterdam**

Retention: **1** Months

Source: ☒ Backup Policy Destination ☐ Replication Policy Destination

Destination: **Cloud services**

Target: **Select**

Figure 6-10 Configuring object storage archival

2. Click **Next**. Figure 6-11 shows a summary of the policy that will be created.

← [Back to Policy Overview](#)

Add SLA Policy

✓ Policy Type

✓ Policy Rules

● Review

Review

SLA Policy
SLA Policy Name: spp-container-policy
SLA Policy Type: Tiered snapshot

Snapshot Protection
Retention: 15 Days (Snapshot)
Snapshot Prefix: spp-containers

Backup Policy
Retention: 15 Days (Backup)
Target Site: Primary
Encrypted Disk Storage:

Additional Copies

Standard object storage (incremental copy)
Retention: Same retention as source selection. (Object Storage Copy)
Destination Target: MinIO
Schedule: Every 1 Days at 12:00:00 AM Europe/Amsterdam from Nov 1, 2022

Figure 6-11 SLA policy review

3. Click **Submit** to create the SLA policy.

6.4 Red Hat OpenShift Container Platform backups

In this section, we use our sample application to show how we can create a snapshot, a backup copy of the application, and the cluster-scoped resources that are needed to restore the application to a new namespace on our cluster (which also can be used to restore the application to a new cluster). We use the policy that was created in 6.3, “Creating a service-level agreement policy” on page 130.

To back up Red Hat OpenShift Container Platform objects, you must understand what kind of objects exist and what is backed up if such an object is assigned to an SLA. The following Red Hat OpenShift Container Platform objects can be backed up with IBM Spectrum Protect Plus:

- **Cluster:** Assigning an SLA to a cluster backs up all resources in the cluster. These resources include all metadata (cluster-scoped and namespace-scoped resources) and all PVCs.

Note: Depending on the size and the amount of namespaces, pods, and PVC data that is managed by the Red Hat OpenShift Container Platform cluster, it might not be practical to assign the entire cluster to a single SLA. Instead, split the workload into smaller pieces and assign the various namespaces to multiple SLAs.

- **Cluster Resources:** Assigning a SLA to Cluster Resources backs up cluster-scoped resources, such as PersistentVolumes, ClusterRoles, StorageClasses, CSIDrivers, VolumeSnapshotClasses, and CustomResourceDefinitions.

Note: This type of backup might help in disaster recovery (DR) situations when a full cluster must be rebuilt. However, this type does not include persistent data in PVCs or namespace-scoped resources, so these resources also must be protected.

- ▶ **Project:** A project in the cluster, including all namespace-scoped resources (metadata) and the PVCs (persistent data).
- ▶ **Project Resources:** Project-scoped resources, such as PersistentVolumeClaims, pods, containers, ConfigMaps, secrets, services, and deployments.

Note: Selecting only the project resources does not back up the persistent data. To also protect the persistent data, assign the entire project to an SLA or include dedicated PVCs to the SLA.

- ▶ **PVC:** PVCs in a project (the persistent data).

As a prerequisite for the tasks that are described in the next sections, the following assumptions can be made:

- ▶ An application server inventory job was performed and discovered all resources in the Red Hat OpenShift Container Platform cluster.
- ▶ One or more SLAs of type Containers were created, as defined in 6.3.1, “General parameters” on page 131.

6.4.1 Performing container backups

To back up Red Hat OpenShift Container Platform resources, you must assign those resources to an SLA. This requirement also is true for ad hoc backups.

Complete the following steps:

1. Log in to the IBM Spectrum Protect Plus GUI and select **Manage Protection** → **Containers** → **OpenShift**, as shown in Figure 6-12.




Figure 6-12 Accessing the backup capability

The clusters that IBM Spectrum Protect Plus are configured with appear, as shown in Figure 6-13.

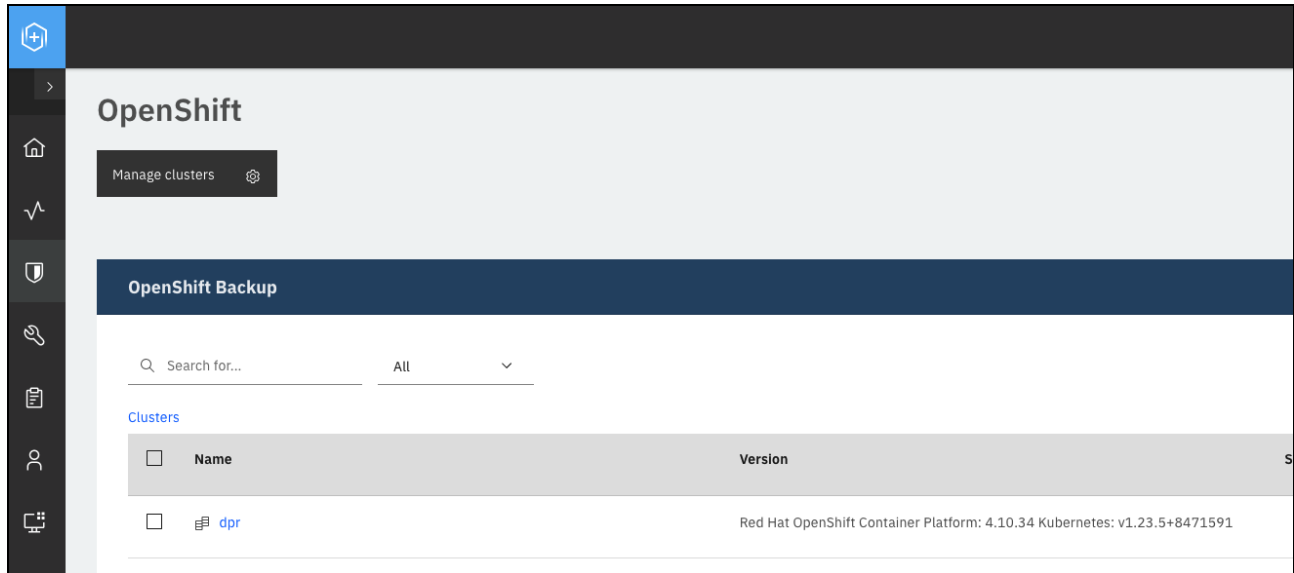


Figure 6-13 Red Hat OpenShift Container Platform Clusters that are configured for IBM Spectrum Protect Plus

2. Filter the resources by searching for the example project, and the project resources and any cluster-scoped resources are shown. Select the project and the cluster resources, as shown in Figure 6-14. The cluster resources are needed only if you are backing up to a new cluster.

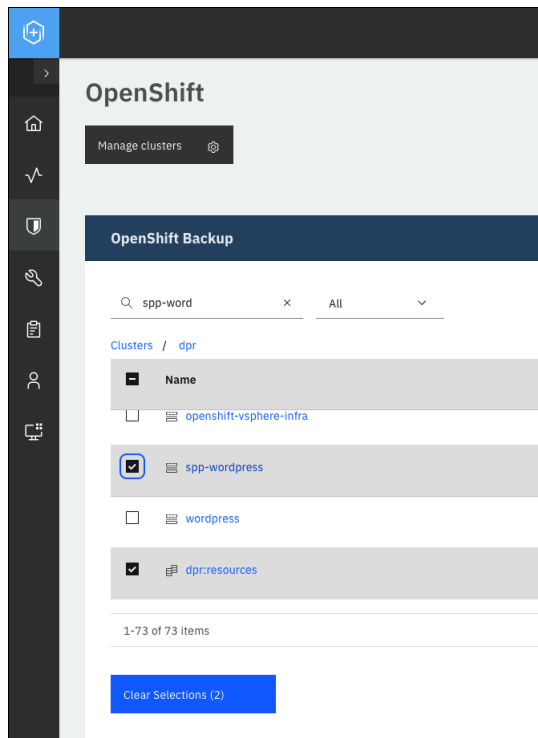


Figure 6-14 Selecting the resources to back up

3. Select the policy for the backup. In our example, we select the policy that we created in 6.3.1, “General parameters” on page 131, as shown in Figure 6-15. Click **Save**.

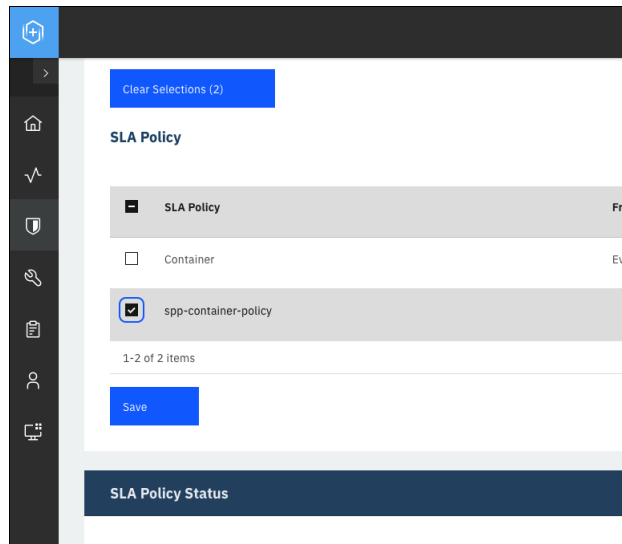


Figure 6-15 Selecting the policy to use for the backup

4. From the SLA Policy Status section, click **Actions for your policy** and click **Start**, as shown in Figure 6-16.

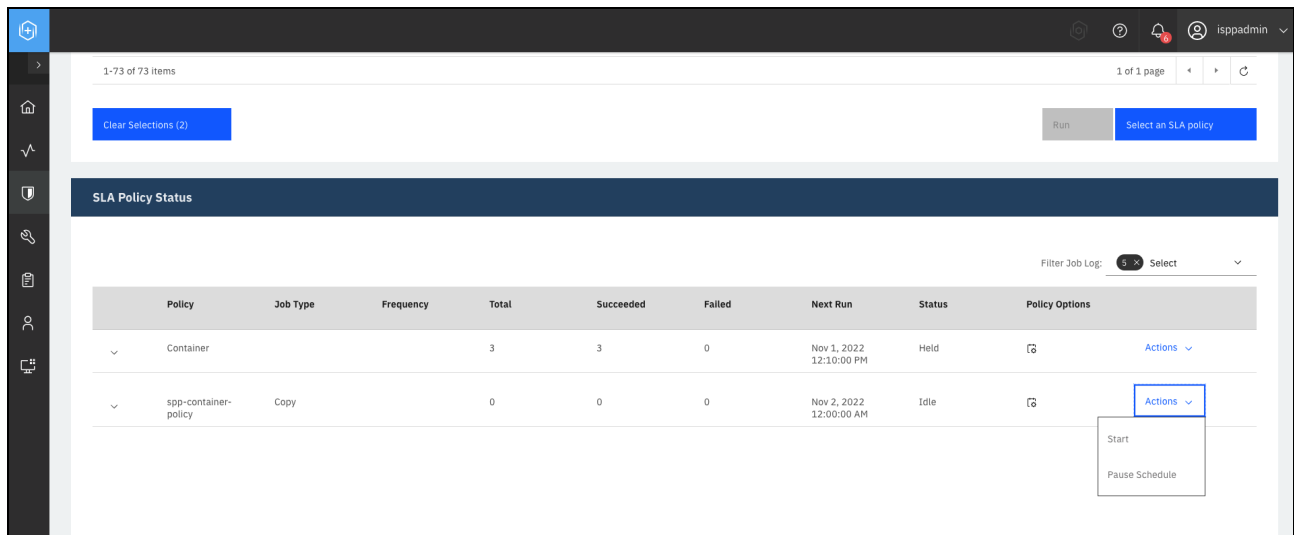


Figure 6-16 Running the policy

5. A dialog window opens and prompts you for an action. In this example, we first create a snapshot of the application and cluster resources, so select **Snapshot** and click **OK**, as shown in Figure 6-17.

Note: This procedure is required only for manual backups. Scheduled backups perform the snapshot and backup automatically.




Figure 6-17 Starting the backup snapshot process

A dialog opens and indicates that the snapshot started, as shown in Figure 6-18.

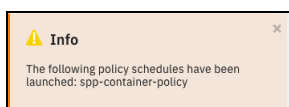


Figure 6-18 Dialog to indicate that a job is running

6. When the job is complete, a snapshot of the resources at the cluster and project level are created. Now, we create a backup of the data to our S3 storage so that the application can be restored to a new namespace. Return to the backup window and find the container SLA policy that we used to create the snapshot. Select **Actions** → **Start** to start a new job. In the dialog window that opens, select **Backup**, as shown in Figure 6-19.




Figure 6-19 Backup option

6.4.2 Monitoring backup jobs

To see an overview of all jobs that are active, select **Jobs and Operations** → **Running Jobs** in the IBM Spectrum Protect Plus GUI, as shown in Figure 6-20.

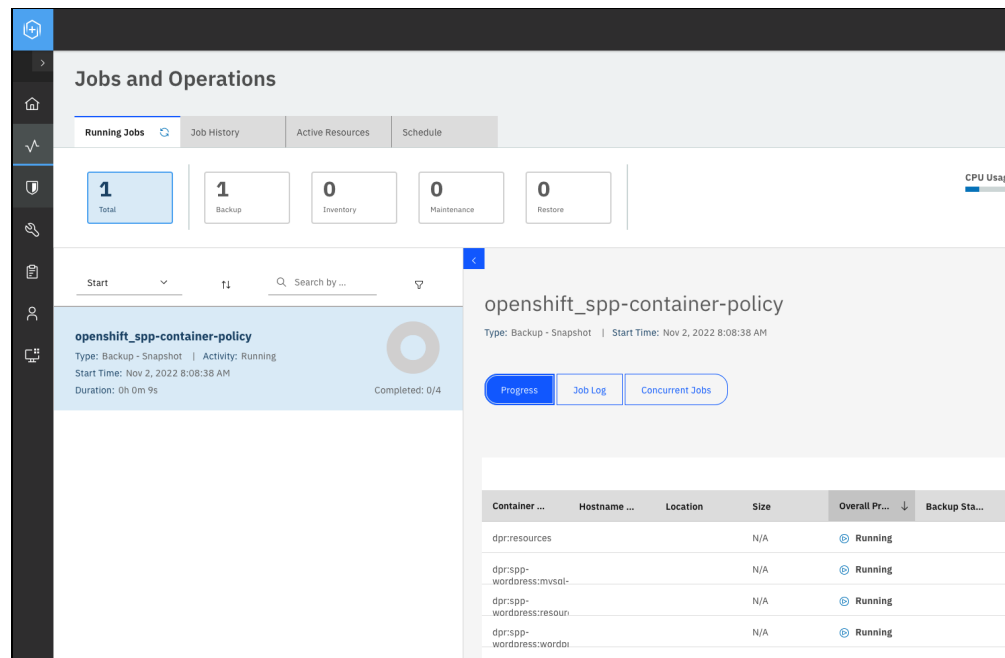


Figure 6-20 Jobs and operations

In 6.2.2, “Configuring the sample application with backup hooks” on page 128, we created hooks to ensure that an application consistent backup was taken. The job log for the backup should show that the hooks that we created ran, as shown in Figure 6-21.

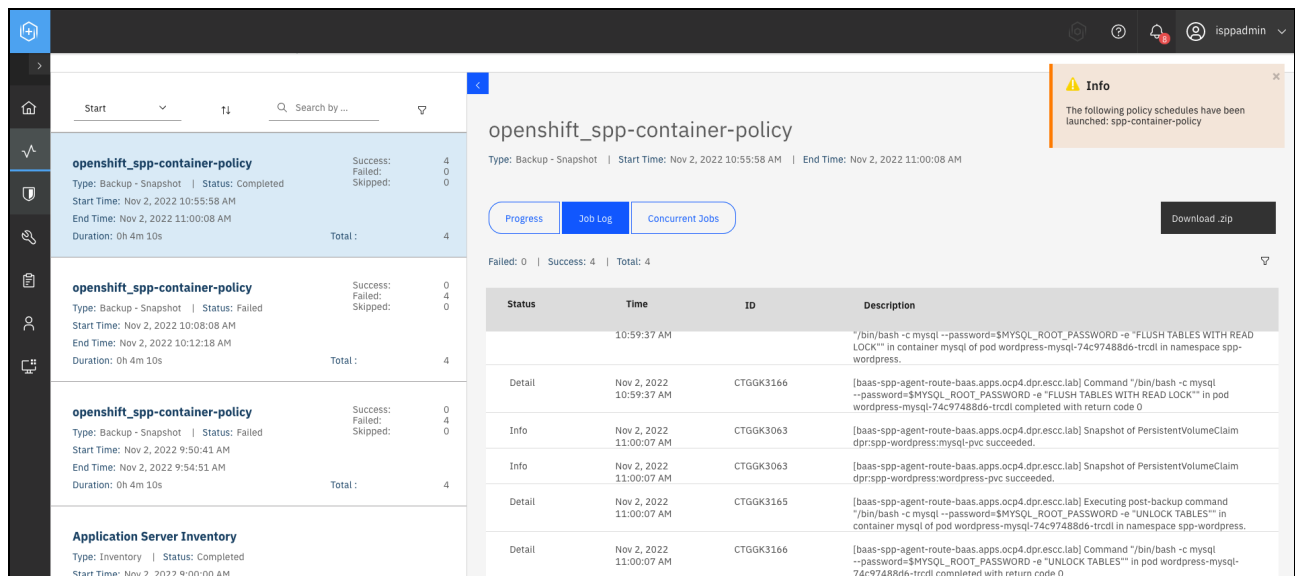


Figure 6-21 Backup logs showing running of the hooks

6.5 Performing restores

This section describes how to perform restores of Red Hat OpenShift objects by using the IBM Spectrum Protect Plus GUI. We use the backup that is we created in 6.4, “Red Hat OpenShift Container Platform backups” on page 136 to demonstrate the restore procedure to a new namespace on the same cluster. We will restore all of the resources that we backed up.

6.5.1 Deciding which objects must be restored

Before starting a restore job, you must know which objects can be restored and which objects might be required in a specific situation.

The following potential restore scenarios are available:

- ▶ Cluster or namespace configurations mistakenly were changed, or objects (such as custom resource definitions (CRDs) or deployments) were deleted: This issue might be solved by restoring cluster- or namespace-scoped resources.

Note: Restoring cluster or namespace metadata does not overwrite objects.

- ▶ A logical error occurred to persistent data, for example, a database table was dropped or deleted: This issue might be solved by restoring one or more PVCs.
- ▶ A stateful application must be cloned to another namespace or cluster for development or testing purposes: This issue might be solved by restoring a full namespace backup, including metadata resources and PVCs. For our example, we clone the application to a new namespace on the same cluster.

6.5.2 Restoring the application

To restore the application, complete the following steps:

1. To start the restore, select **Manage Protection** → **Containers** → **OpenShift**, and then click **Create job**, as shown in Figure 6-22 on page 143.

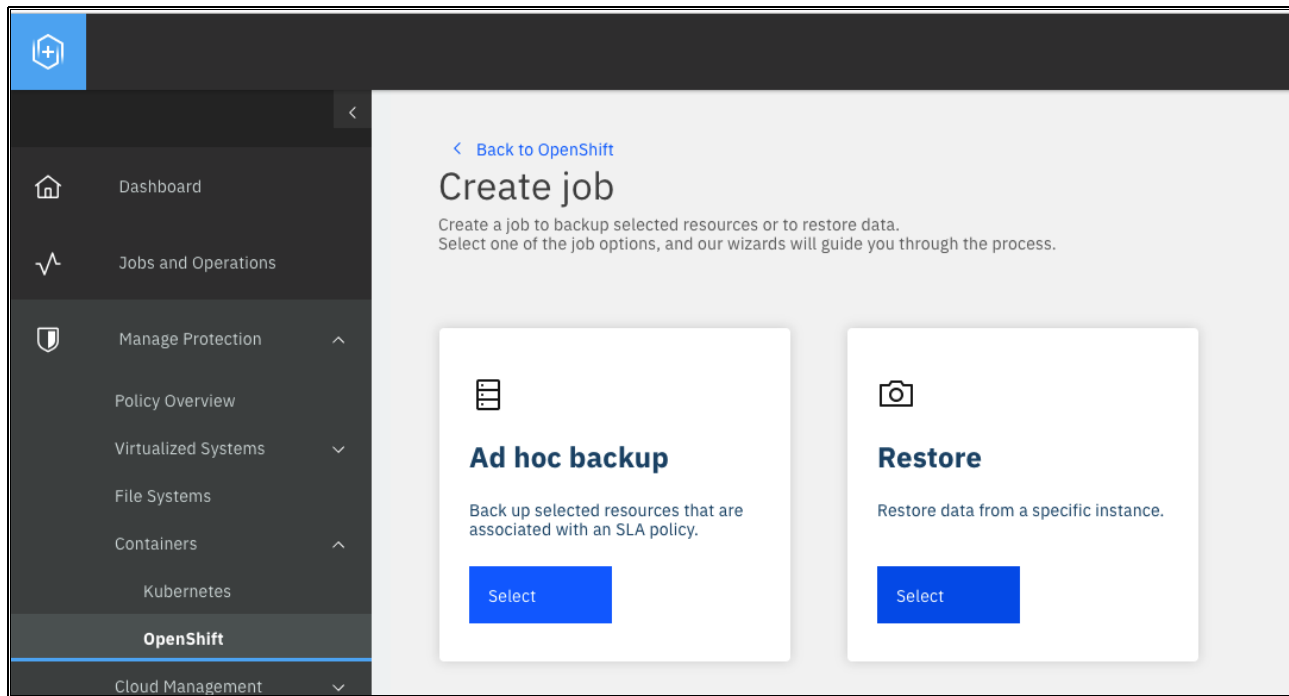


Figure 6-22 Creating a job

2. Select the **Restore** job to start the configuration. Select the source for the restore by selecting the resources that we backed up from the spp-wordpress project, as shown in Figure 6-23. Click **Next**.

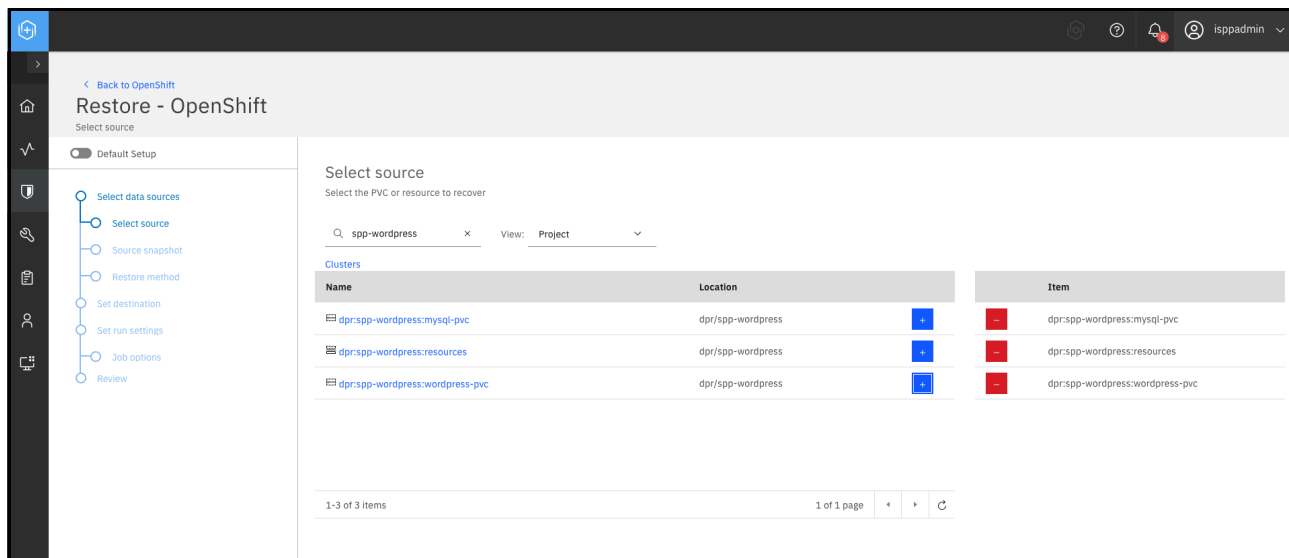


Figure 6-23 Selecting the source for the restore job

3. Select the snapshot to restore. We restore from the backup copy that we created on the MinIO storage. In the columns of the Source snapshot window, select **From Copy**, **On-Demand**, **Cloud service copy**, and **MinIO**, and then for each of the resources, we want to select the restore point from the available restore points, as shown in Figure 6-24. Click **Next**.

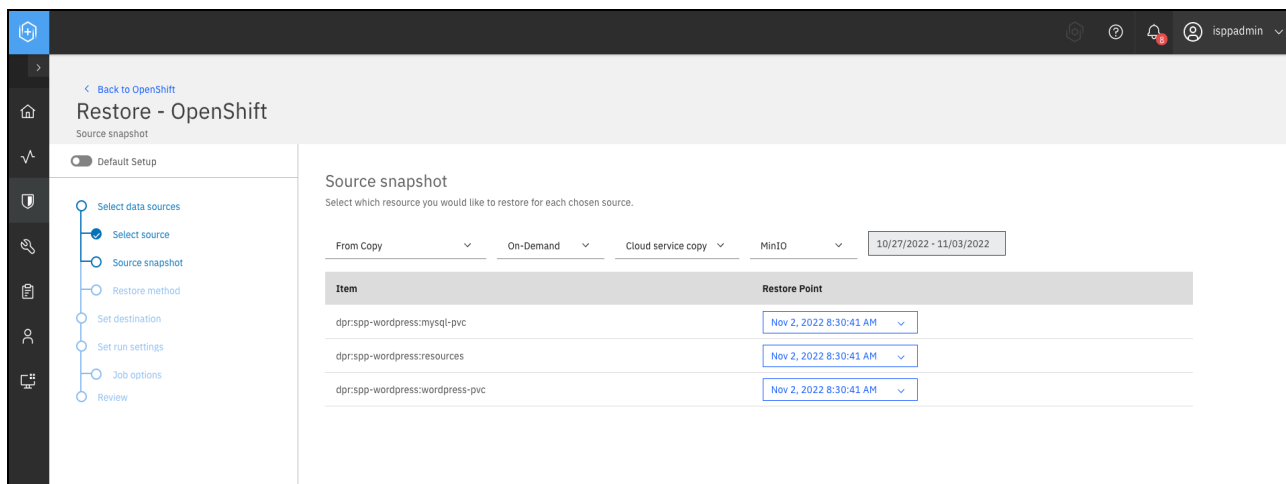


Figure 6-24 Selecting a restore job backup

4. Select the restore method, and ensure that the radio button for the same cluster is selected along with original storage class. For the namespace, specify an alternative namespace (for our example, we use spp-wordpress-restore). Click **Next**. Figure 6-25 shows the window for the restore job destination.

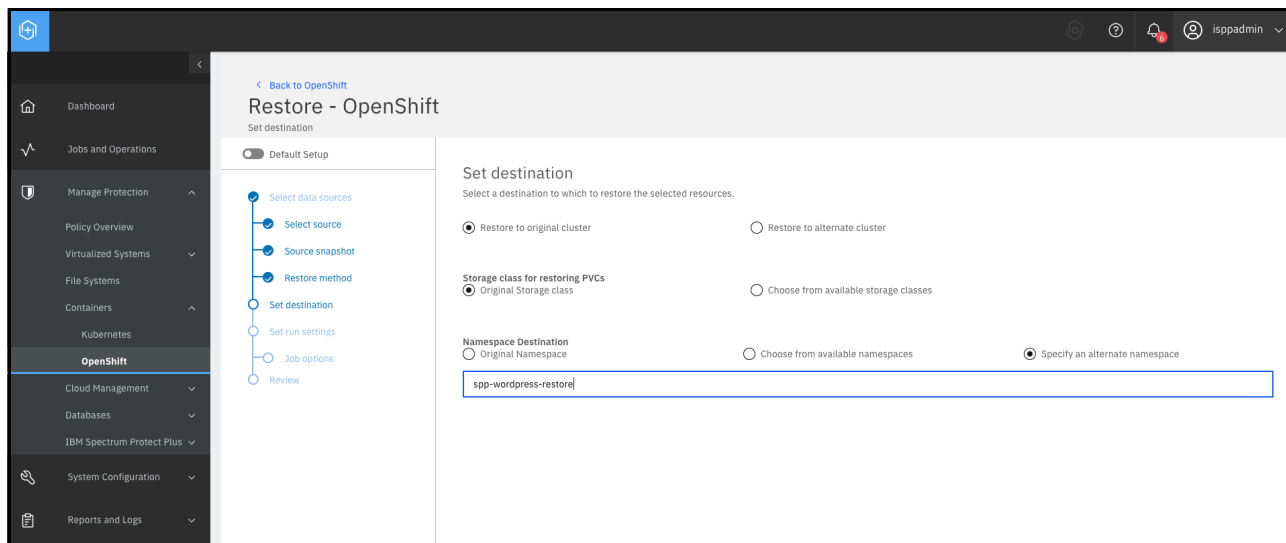


Figure 6-25 Restore job destination

5. Next, complete the restore options window. Select **Do not Overwrite PVCs** and check the remaining checkboxes for the options, as shown in Figure 6-26 on page 145. Click **Next**.

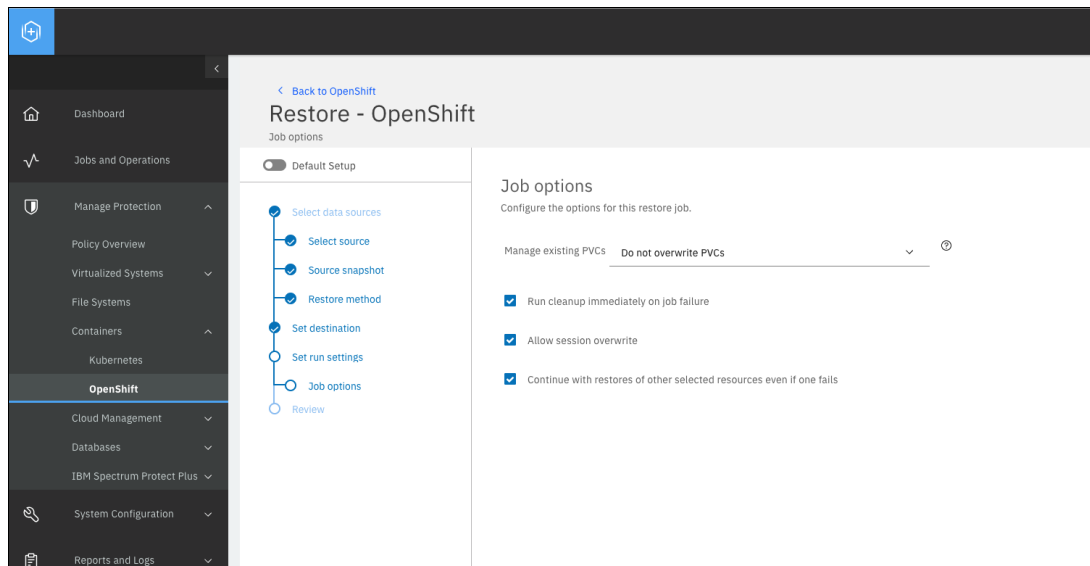


Figure 6-26 Restore job options

Figure 6-27 shows a summary of the restore job.

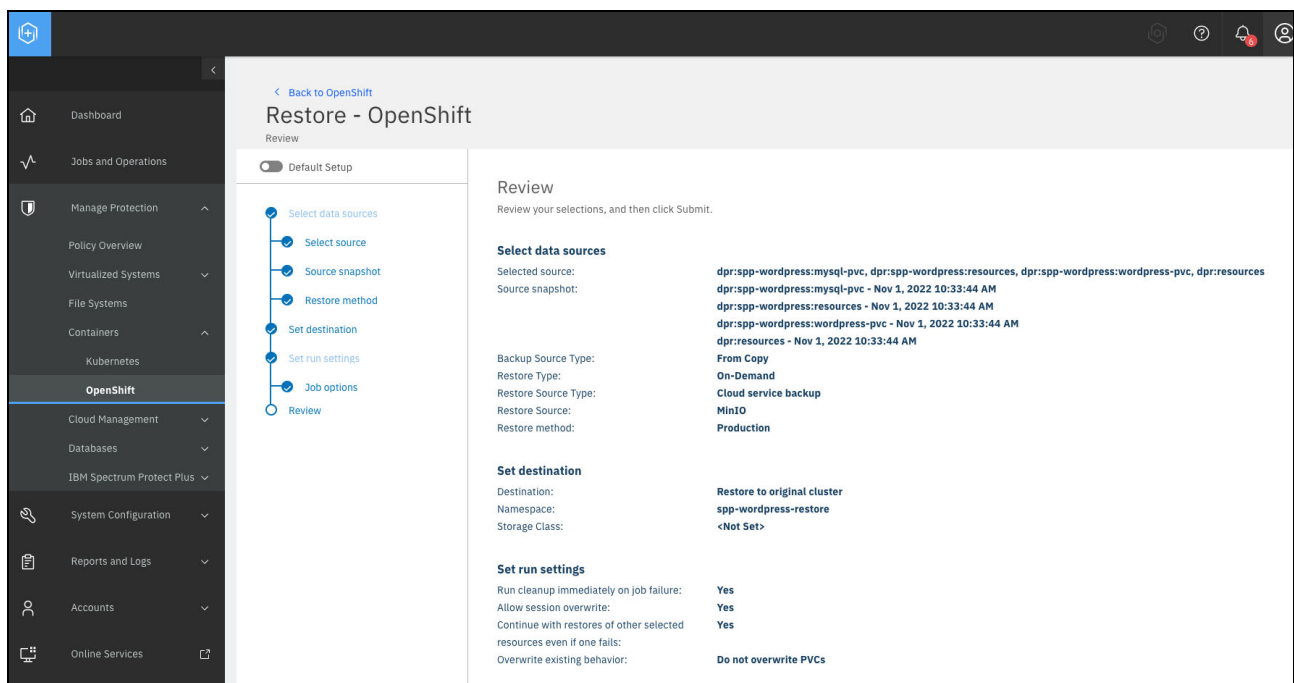


Figure 6-27 Restore job summary

6. The restore job can be monitored in the same way as the backup job, as described in 6.4.2, “Monitoring backup jobs” on page 141. When the restore is complete, the new application is available in the `spp-wordpress-restore` namespace. Log in to the cluster as a cluster administrator. The route to the application can be retrieved by running the commands that are shown in Example 6-17.

Example 6-17 Retrieving the route for the application

```
oc project spp-wordpress-restore
oc get route
```

The host or port entry of the returned route is the route to the application.

6.5.3 Application considerations for necessary rework after a restore

In the WordPress application, some information is hardcoded in it that relates to the original workspace. For the application to work correctly, this information must be corrected for the new workspace. To accomplish this task, complete the following steps:

1. Log in to the Red Hat OpenShift cluster as a cluster administrator, and then run the commands that are shown in Example 6-18 to get the name of the `mysql` pod.

Example 6-18 Getting the name of the mysql pod

```
oc project wordpress
export MYSQL_POD=$(oc get pods | grep wordpress-mysql | awk '{ print $1 }' )
```

2. Run `exec`, as shown in Example 6-19.

Example 6-19 Running exec

```
oc exec -it ${MYSQL_POD} --
```

3. Run the commands that are shown in Example 6-20. The domain portion of the URL should be replaced with the domain that you are using for Red Hat OpenShift.

Example 6-20 Updating the mysql wp-options table

```
export TARGETURL=http://wordpress-wordpress-restore.<domain>

mysql -u root -p${MYSQL_ROOT_PASSWORD} -e "use wordpress; UPDATE wp_options SET
option_value='${TARGETURL}' WHERE option_name='siteurl'; UPDATE wp_options SET
option_value='${TARGETURL}' WHERE option_name='home';"
```

4. Restart the front-end pod. This process forces the WordPress application to read the configuration options again and stop the unwanted redirects to the old home URL.



Red Hat OpenShift cluster disaster recovery solution

This chapter describes the usage of IBM Spectrum Protect Plus as a disaster recovery (DR) solution for Red Hat OpenShift clusters.

This chapter includes the following topics:

- ▶ 7.1, “General disaster recovery considerations” on page 148
- ▶ 7.2, “Protecting the environment” on page 149
- ▶ 7.3, “Preparing for a DR restore with IBM Spectrum Protect Plus” on page 152
- ▶ 7.4, “Cluster-scoped resources recovery” on page 153
- ▶ 7.5, “Application namespaces recovery” on page 158

7.1 General disaster recovery considerations

When discussing DR for a Red Hat OpenShift cluster, You should understand the different components and objects that make up a cluster environment and which ones must be recovered if a disaster occurs.

A generic setup of a Kubernetes based container environment is shown in Figure 7-1 (the red circles show what is in scope of IBM Spectrum Protect Plus, and the gray boxes indicate elements that are out of scope).

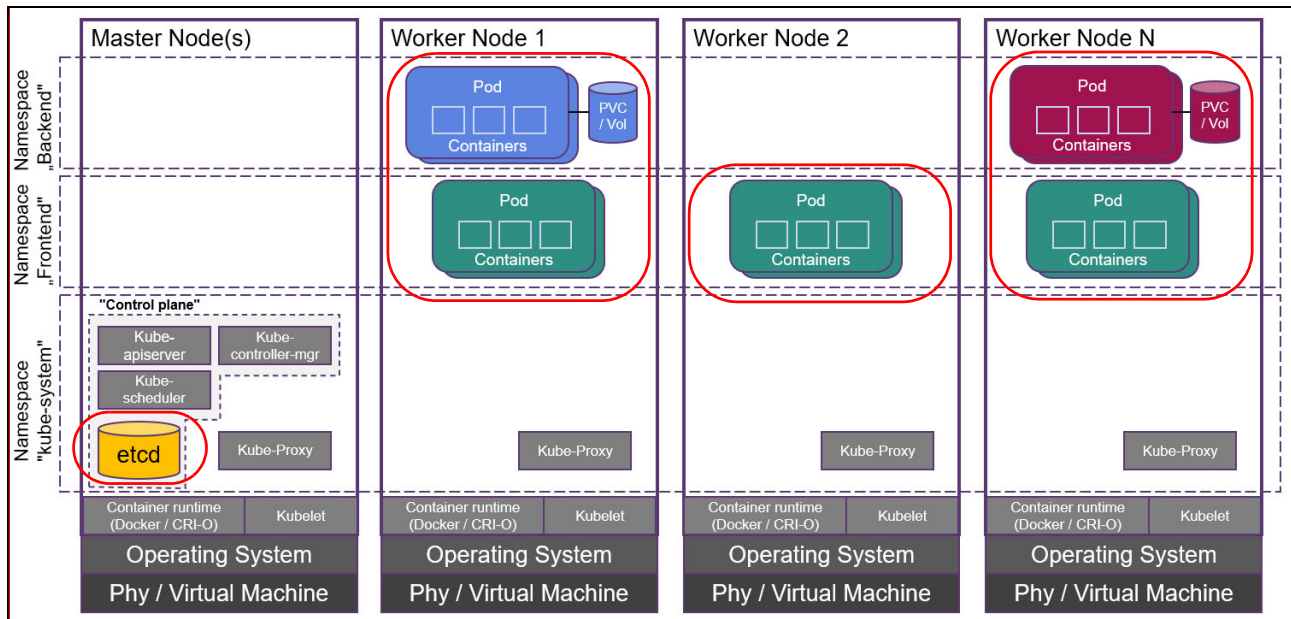


Figure 7-1 A generic Red Hat OpenShift cluster overview diagram

The environment consists of the following essential components:

- ▶ A set of physical or virtual machines (VMs) that host the control and worker nodes of the environment, network components, and storage.
- ▶ The host operating system, a container run time, Kubelet, and Kube-Proxy.
- ▶ The control plane with the distributed configuration database (etcd).
- ▶ Pods, containers, persistent volume claims (PVCs), and so on.

Important: IBM Spectrum Protect Plus Container Backup Support does not protect all these components. Therefore, if a disaster occurs, a basic infrastructure must be reestablished manually before an IBM Spectrum Protect Plus restore can be performed.

The following components are not protected by IBM Spectrum Protect Plus Container Backup Support and cannot be restored with the product:

- ▶ The operating system of the nodes.
- ▶ Kubernetes or Red Hat OpenShift system components and their configuration (control plane components, kubelet, kube-proxy, and the container run time).
- ▶ Any artifacts that are outside of the Kubernetes environment, such as non IBM Container Storage Interface (CSI) volume drivers or node-related network components and configurations.

Tip: To be prepared for a disaster, it is a best practice to perform regular recovery tests to a fenced environment. As applications change over time, ensure that all components are included in the backup and available for recovery.

7.2 Protecting the environment

As described in 7.1, “General disaster recovery considerations” on page 148, a comprehensive strategy must be developed to protect the Red Hat OpenShift environment against a disaster.

7.2.1 Protecting the infrastructure, and complementary systems and services

Because not all components of a Red Hat OpenShift container environment are in the scope of an IBM Spectrum Protect Plus backup, more arrangements must be made to prepare for a recovery of the base infrastructure and external services on which the Red Hat OpenShift cluster relies. These arrangements include the following examples:

- ▶ Collect and save infrastructure-related configuration information, such as the following items:
 - The number of control and worker nodes in the cluster
 - Network interfaces, hostnames, and IP addresses
 - CPU and memory resources per node
 - Storage capacities, CSI storage-classes, and provider types

This information must be saved in a way so that you can quickly rebuild an “empty” Red Hat OpenShift cluster from scratch. This preparation can be facilitated by using configuration management systems like GitHub, or automation frameworks like Ansible.

- ▶ Save the identity provider configuration (and potentially secrets) to reestablish user authentication quickly on a rebuilt cluster.
- ▶ Use an external registry to store and manage your images. An example for a distributed and highly available container image registry is Red Hat Quay, which can be implemented as a private registry for the enterprise on-premises. It also offers a hosted option through [quay.io](#). For more information, see [Red Hat Quay](#).

Work with the Red Hat OpenShift administrators to develop a strategy to reattach the registry after a cluster is rebuilt from scratch.

- ▶ Back up the complementary systems and services, such as DNS, NTP, LDAP, or file servers by using a backup software product.

Tip: If these services are provided by systems that are running on VMware or Microsoft Hyper-V VMs, IBM Spectrum Protect Plus can be used to back up these systems.

- ▶ Ensure that the latest version of the IBM Spectrum Protect Plus server, vSnap servers, and the backup as a service (BaaS) component is installed.
- ▶ Ensure that the backup systems are running on a separate infrastructure. For example, the IBM Spectrum Protect Plus server and vSnaps must not be hosted on the same storage systems that host volumes for the Red Hat OpenShift container storage providers.

Tip: Create extra copies of backup data in a different location, such as by implementing vSnap server replication or adding copies from a vSnap server to IBM Spectrum Protect or a supported IBM Cloud Object Storage System.

7.2.2 Protecting cluster resources and persistent data

As described in 4.5, “Container backup and restore types” on page 77, IBM Spectrum Protect Plus Container Backup provides options to back up Red Hat OpenShift cluster metadata (cluster-scoped and namespace-scoped resources), and persistent application data that is stored within PVCs.

Figure 7-2 shows various cluster-scoped (blue boxes) and namespace-scoped resources and objects (green boxes) within a Red Hat OpenShift cluster.

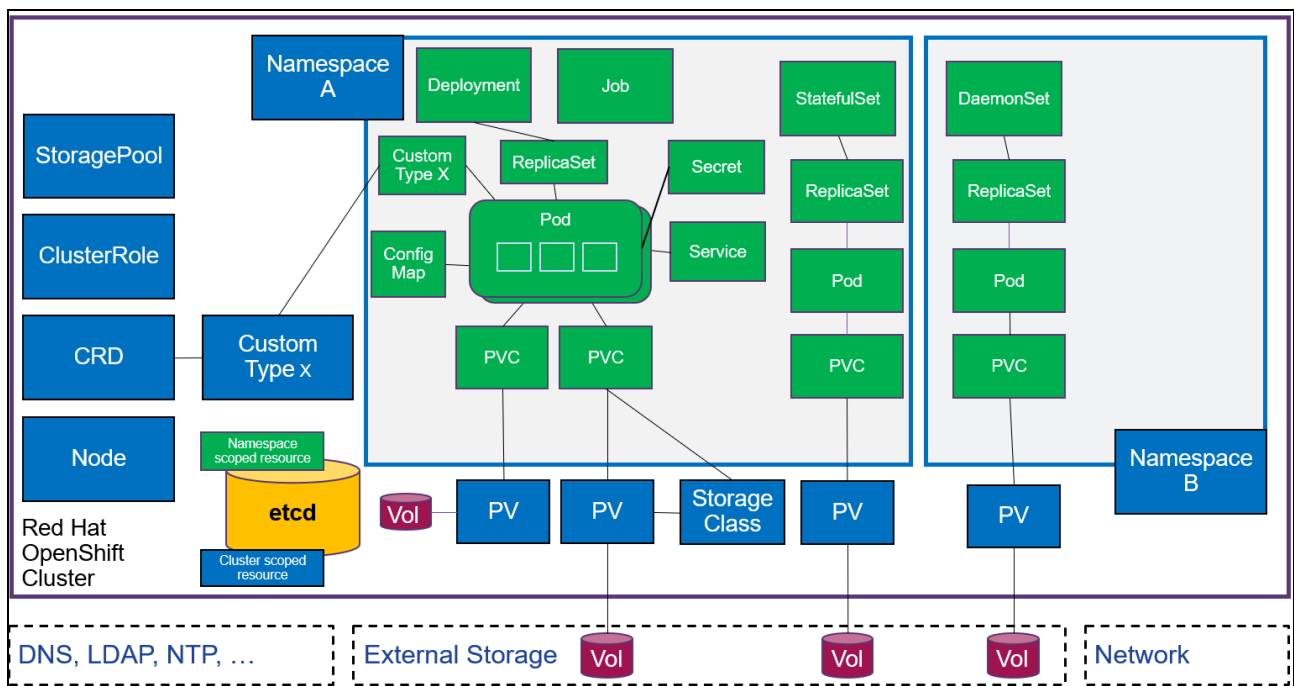


Figure 7-2 Cluster- and namespace-scoped resources and objects

Backing up cluster-scoped resources

Create backups of the cluster-scoped resources (the blue boxes in Figure 7-2) to allow restores of cluster metadata if an accidental deletion or a disaster occurs. Ensure that the policies that are used perform backups to vSnap servers or cloud storage so that backup copies are stored outside of the Red Hat OpenShift cluster.

Figure 7-3 on page 151 shows the assignment of cluster-scoped resources of Red Hat OpenShift cluster DPR-0CP to a service-level agreement (SLA) that is named ESCC-DPR-RESOURCE.




Figure 7-3 Assignment of cluster-scoped resources to an SLA

For more information about how to create SLA policies and assigning them to Red Hat OpenShift resources, see Chapter 6, “Using Container Backup Support” on page 121.

Note: IBM Spectrum Protect Plus Container Backup Support uses the Velero tool to back up and restore cluster- or namespace-scoped resources.

Velero does not overwrite existing cluster resources during a restore. A restore requires manually reconstructing some configuration items that are created after a cluster is rebuilt from scratch in a DR situation.

For example, when redeploying an empty Red Hat OpenShift cluster as part of a DR procedure, the OAuth server is implemented with a default configuration. This configuration might not contain all identity providers that were configured in the original cluster. Even if the former OAuth configuration was backed up by IBM Spectrum Protect Plus, it is not applied as part of a restore process because Velero must overwrite a configuration.

Therefore, the OAuth configuration must be reestablished manually by using information that was saved and prepared separately (see 7.2.1, “Protecting the infrastructure, and complementary systems and services” on page 149).

Backing up important namespaces

Create backups of all namespaces that contain critical applications. These backups must include the metadata (namespace-scoped resources) and the PVCs. For more information about how to create backups of namespaces, see Chapter 6, “Using Container Backup Support” on page 121.

Ensure that the SLA policies that are used perform backups to vSnap servers externally store backup copies to the Red Hat OpenShift cluster.

The example that is shown in Figure 7-4 shows the assignment of namespace-scoped resources and PVCs of the namespaces mariadb and wordpress in Red Hat OpenShift cluster DPR-OCF to an SLA that is named ESCC-OCF-DPR.

OpenShift Backup

Search for...

ESCC-OCF-DPR

View: Project

Clusters

Last inventory completed May 19, 2021 10:01:57 AM

Run Inventory

<input type="checkbox"/>	Name	Location	SLA Policy
<input type="checkbox"/>	DPR-OCF:mariadb:mariadb	DPR-OCF/mariadb	ESCC-OCF-DPR
<input type="checkbox"/>	DPR-OCF:mariadb:resources	DPR-OCF/mariadb	ESCC-OCF-DPR
<input type="checkbox"/>	DPR-OCF:wordpress:mysql-pvc	DPR-OCF/wordpress	ESCC-OCF-DPR
<input type="checkbox"/>	DPR-OCF:wordpress:resources	DPR-OCF/wordpress	ESCC-OCF-DPR
<input type="checkbox"/>	DPR-OCF:wordpress:wordpress-pvc	DPR-OCF/wordpress	ESCC-OCF-DPR

Total: 5

Figure 7-4 Assignment of namespace-scoped resources and PVCs to an SLA

Tip: For more complex applications (for example, spanning multiple namespaces or requiring the recovery of single components in a sequence of multiple steps), it might be helpful to assign labels to these components and SLA policies that are based on labels instead of namespaces. An example is an application with a front end (web server) and a back end (database) that share a namespace. By using the labels “tier=frontend” and “tier=backend” and assigning the SLA policy to these labels, you can do more granular restores later instead of restoring the entire namespace.

7.3 Preparing for a DR restore with IBM Spectrum Protect Plus

If a disaster occurs, the basic infrastructure must be reinstalled. The IBM Spectrum Protect Plus BaaS component must be deployed as a foundation for subsequent restore tasks. The following steps must be completed before the restore can start:

1. Reinstall a new Red Hat OpenShift cluster by using the information that was gathered in the production environment (see 7.2.1, “Protecting the infrastructure, and complementary systems and services” on page 149). Then, deploy the same cluster version that was used before the disaster occurred.
2. Configure (or reconfigure) the OAuth configuration and the identity providers.
3. Configure the Red Hat OpenShift cluster to reattach to an image registry, for example, Red Hat Quay.
4. Install (or reinstall) the container storage providers that were used by the applications to store PVCs.

Note: Although the storage-class names can be different from the formerly used names, the storage providers must be of the same type.

5. Implement the prerequisites to install the IBM Spectrum Protect Plus Container Backup support (BaaS) component. For more information about implementing the BaaS prerequisites, see 4.3, “Red Hat OpenShift prerequisites and supported environments” on page 68.

6. Install and configure the BaaS component and register the Red Hat OpenShift cluster to the IBM Spectrum Protect Plus server by using the same name that was used before the disaster occurred. For more information about installing the BaaS component, see Chapter 5, “Implementing Container Backup Support” on page 85.
7. Verify that the connection between the IBM Spectrum Protect Plus GUI server and the BaaS component was correctly established. You can verify the connection by running and monitoring an application server inventory job. In the IBM Spectrum Protect Plus GUI, select **Manage Protection** → **Containers** → **OpenShift** → **Manage Clusters** → **Actions**. Then, select **Inventory** from the drop-down menu.

7.4 Cluster-scoped resources recovery

During the preparatory measures of rebuilding a new cluster from scratch, most configurations were re-created already (see 7.2.1, “Protecting the infrastructure, and complementary systems and services” on page 149).

However, some configuration items such as Namespaces, ClusterRoles, or CustomResourceDefinitions might be missing after a scratch installation or an unintended deletion and can be restored through IBM Spectrum Protect Plus.

Note: Recovering an entire Red Hat OpenShift cluster is a two-fold process. Although all *configuration data* should be tracked in and restored from a configuration management platform (like GitHub), IBM Spectrum Protect Plus should be leveraged to back up and restore the *user data*. Usually, this data consists of namespace-scoped resources and PVCs that were included in a namespace backup. Thus, restoring user data on a namespace level (or even more granularly based on labels) should be sufficient to recover applications to a new cluster.

7.4.1 Temporarily pausing the machine-config operator

Note: Consider the following points:

- ▶ A restore of cluster-scoped resources also recovers machine-config objects if those objects existed before the disaster and were not re-created manually. Examples for such machine-configs are NTP configurations or hardware-related CSI drivers.
- ▶ When a new machine-config object is applied as part of a restore process, a restart of the Red Hat OpenShift control and worker nodes is triggered, which interrupts the restore operation and causes the operation to hang.

To prevent restored machine-configs from interrupting or hanging the restore process, the machine-config operator must be paused before starting a restore of cluster-scoped resources, as shown in Example 7-1.

Example 7-1 Disabling autoreboot for the machine-config operator

```
[root@helper ~]# oc patch --type=merge --patch='{\"spec\":{\"paused\":true}}'
machineconfigpool/master
machineconfigpool.machineconfiguration.openshift.io/master patched
[root@helper ~]# oc patch --type=merge --patch='{\"spec\":{\"paused\":true}}'
machineconfigpool/worker
machineconfigpool.machineconfiguration.openshift.io/worker patched
```

When the restore of cluster-scoped resources by using IBM Spectrum Protect Plus is complete, re-enable the machine-config operator (see Example 7-2).

Example 7-2 Enabling autoreboot for the machine-config operator

```
[root@helper ~]# oc patch --type=merge --patch='{ "spec": { "paused": false } }'
machineconfigpool/master
machineconfigpool.machineconfiguration.openshift.io/master patched
[root@helper ~]# oc patch --type=merge --patch='{ "spec": { "paused": false } }'
machineconfigpool/worker
machineconfigpool.machineconfiguration.openshift.io/worker patched
```

For more information about how to disable an autoreboot that is triggered by the Red Hat OpenShift machine-config operator, see [Disable autoreboot after a change with the machine-config operator in OpenShift Container Platform 4](#).

7.4.2 Restoring cluster-scoped resources

To start the process of restoring cluster-scoped resources, complete the following steps:

1. In the IBM Spectrum Protect Plus GUI, select **Manage Protection** → **Containers** → **OpenShift**, and then click **Create Job** in the upper right (see Figure 7-5).

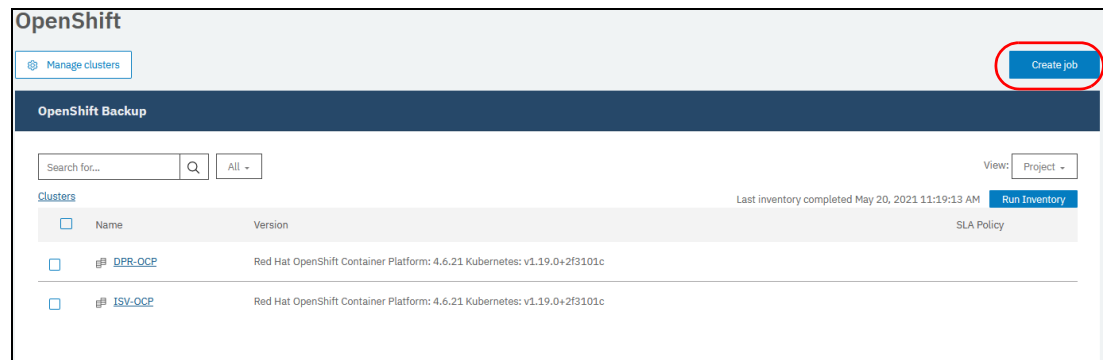


Figure 7-5 Creating an IBM Spectrum Protect Plus job

2. Click **Restore** to start the restore wizard (see Figure 7-6).




Figure 7-6 Creating a restore job

3. In the Restore wizard, select the name of the Red Hat OpenShift cluster that must be recovered. Then, select the cluster-scoped resources item and add it to the restore item list by clicking the plus sign (+), as shown in Figure 7-7 on page 155.

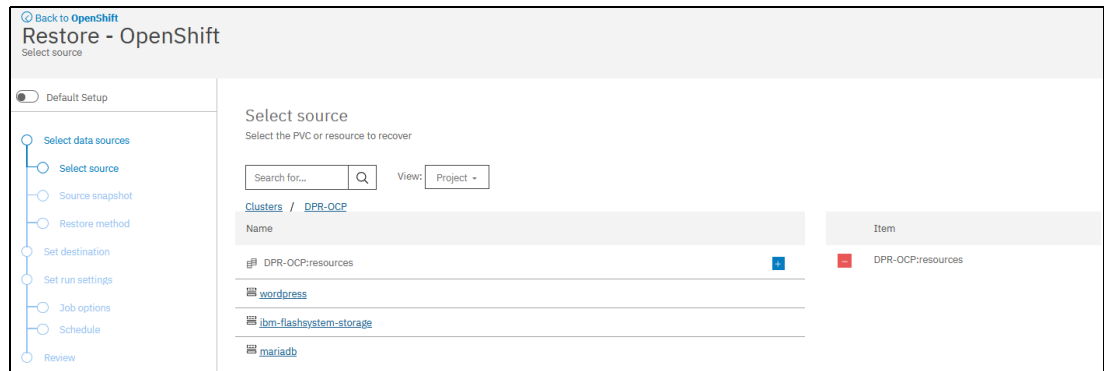


Figure 7-7 Selecting the cluster-scoped resources for restore

4. To restore data from an external copy (a vSnap server or cloud storage), select **From Copy** → **On-Demand** and select the latest good backup timestamp from the list, as shown in Figure 7-8.

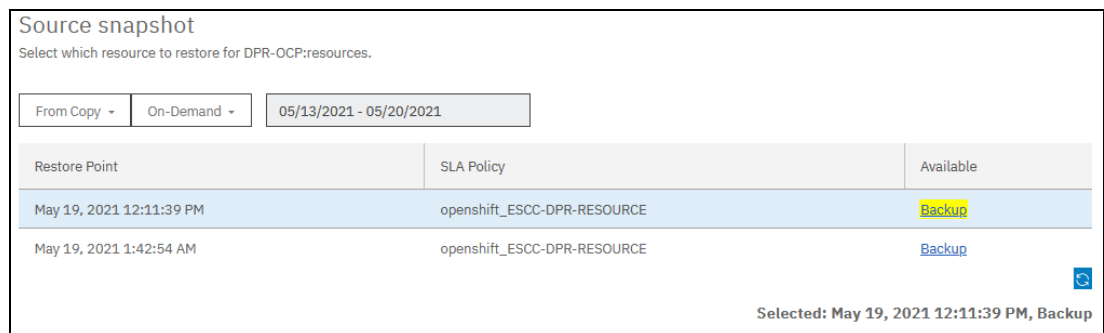


Figure 7-8 Selecting a backup timestamp to be restored

5. Follow the instructions in the wizard by making the following selections:
 - Do not change the PVC name
 - Restore to original cluster
 - Use the original storage class
 - Use the original namespace
 - Use the proposed default job options

- Review the summary, and if all the information is correct, click **Submit**, as shown in Figure 7-9.

Restore - OpenShift
Review

☐ Default Setup

Review
Review your selections, and then click Submit.

Select data sources

Source type: **OpenShift**
 Selected source: **DPR-OCPr-resources**
 Source snapshot: **DPR-OCPr-resources - May 19, 2021 12:11:39 PM**
 Backup Source Type: **From Copy**
 Restore Type: **On-Demand**
 Restore Source Type: **Site**
 Restore Source: **Primary**
 Restore method: **Production**

Set destination

Destination: **Restore to original cluster**
 Namespace: **<empty>**
 Storage Class: **<empty>**

Set run settings

Run cleanup immediately on job failure: **Yes**
 Allow session overwrite: **Yes**
 Continue with restores of other selected resources even if one fails: **Yes**

Preview Restore **Back** **Submit**

Figure 7-9 Submitting the cluster-scoped resources restore job

7.4.3 Monitoring the progress of the restore job

To monitor the progress of a restore job, complete the following steps:

- In the IBM Spectrum Protect Plus GUI, select **Jobs and Operations** → **Running Jobs** and validate that the restore job started correctly (see Figure 7-10).

Jobs and Operations

Running Jobs | Job History | Active Resources | Schedule

1 Total Jobs | **0** Backup | **0** Inventory | **0** Maintenance | **1** Restore

CPU Usage: 6%
IBM Spectrum Protect Plus Host Machine

Sort By: Start | 11 | Search by name: [Q]

onDemandRestore_1621506182752
 Type: Restore | Activity: Running
 Start Time: May 20, 2021 12:23:03 PM
 Duration: 0h 0m 58s | Databases Completed: 0/1

onDemandRestore
 Type: Restore | Start Time: May 20, 2021 12:23:03 PM

Job Log | Concurrent Jobs | Download .zip

Failed: 0 | Success: 0 | Total: 1

Status	Time	ID	Description
Summary	May 20, 2021 12:23:03 PM	CT06A2398	Starting job for policy onDemandRestore_1621506182752 (ID:1106). id -> 1621506183097. IBM Spectrum Protect Plus version 10.1.8-4083.
Detail	May 20, 2021 12:23:03 PM	CT06A2109	Policy has (1) destination database mappings.
Detail	May 20, 2021 12:23:03 PM	CT06A2674	Recovery policy options: {sourceType=OpenShift, selectedSource=[DPR-OCPr-resources], sourceSnapshot=[DPR-OCPr-resources - May 19, 2021 12:11:39 PM], backupSourceType=copy, restoreType=on-demand, restoreSource=Primary, restoreMethod=Production, destination=Restore to original cluster, autocleanupOnJobFailure=true, allowSessionOverwrite=true, continueWithRestoreOnError=true, continueScriptsOnError=false}
Detail	May 20, 2021 12:23:03 PM	CT06A1527	Resolved policy to (restore).
Summary	May 20, 2021 12:23:04 PM	CT06A2490	Resuming PVC/Resource restore.
Detail	May 20, 2021 12:23:05 PM	CT06A2179	Load inventory data: In progress
Info	May 20, 2021 12:23:05 PM	CT06A1095	Loading data for restore points.

Figure 7-10 Monitoring the restore job in the IBM Spectrum Protect Plus GUI

After the backup data is copied from the vSnap to the BaaS component, a Velero restore job reads the cluster-scoped metadata from the backup image and applies it to the Red Hat OpenShift cluster.

Restoring cluster metadata can be a long-running task, and no updates appear in the job log in the IBM Spectrum Protect Plus GUI for some time.

To get an overview of the detailed progress, the Velero job can be queried by using native Velero commands:

1. To query the restore job progress with native Velero commands, log in to the Administrator view in the Red Hat OpenShift Container Platform GUI and select **Workloads** → **Pods**. Select project **baas** from the **Project** drop-down menu, as shown in Figure 7-11.

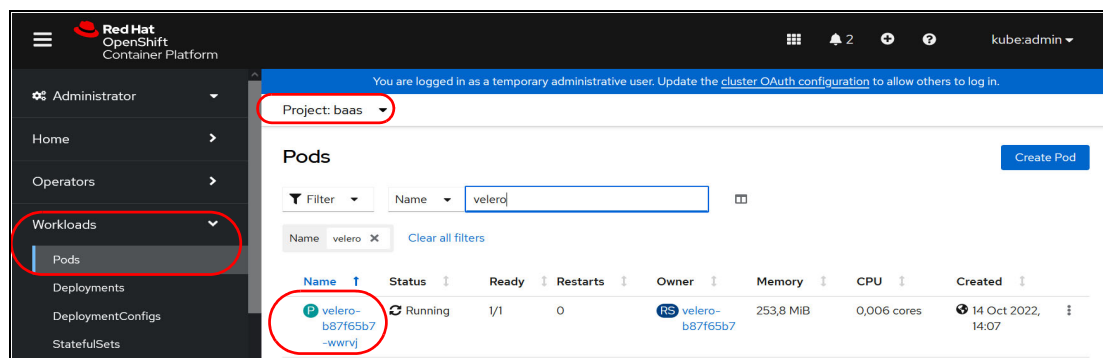


Figure 7-11 Selecting the Velero pod in the Red Hat OpenShift GUI

2. Click the name of the Velero pod and select the **Terminal** tab, as shown in Figure 7-12.




Figure 7-12 Opening a terminal session to the Velero pod

3. Use the following command to determine the ID of the running restore job:

```
./velero restore get
```

The command and its output are shown in Example 7-3.

Example 7-3 Identifying the Velero restore job

```
sh-4.4$ ./velero restore get
NAME BACKUP STATUS STARTED COMPLETED ERRORS WARNINGS CREATED SELECTOR
spprestore-k8s-b168a23b-4193-4475-b8c6-7d12bb6251c5
sppbackup-k8s-e8342fb0-2df5-4b54-9679-bf5c752e7da7 Completed 2022-11-01
10:45:17 +0000 UTC 2022-11-01 10:50:17 +0000 UTC 0 443
2022-11-01 10:45:17 +0000 UTC <none>
```

4. Use the following command to get an estimate about the total number of metadata objects to be restored and how many objects were restored so far:

```
./velero restore describe <restore-job-name> | grep restored
```

The command and its output are shown in Example 7-4.

Example 7-4 Listing the current Velero restore job status

```
sh-4.4$ ./velero restore describe
spprestore-k8s-b168a23b-4193-4475-b8c6-7d12bb6251c5 |grep restored
Total items to be restored: 2880
Items restored: 523
```

Tip: In some situations, a Velero restore job can hang. This issue occurs if the Red Hat OpenShift cluster node that is running the Velero pod is restarted during the restore job (see 7.4.1, “Temporarily pausing the machine-config operator” on page 153).

In such a situation, the CLI can be used to delete a Velero restore job before triggering a new restore task through the IBM Spectrum Protect Plus server GUI.

The following command can be used to delete a hanging Velero restore:

```
./velero restore delete <restore-job-name>
```

7.5 Application namespaces recovery

After the Red Hat OpenShift cluster is restored, the important applications can be recovered. The namespaces objects (projects) were recovered by restoring cluster-scoped resources before, but without any content.

To recover the namespace contents (namespace-scoped resources and PVCs), complete the following steps:

1. Log in to the IBM Spectrum Protect Plus GUI and select **Manage Protection** → **Containers** → **OpenShift**, and then click **Create Job** in the upper right. Click **Restore** to start the restore wizard.
2. Click a namespace to be restored, select the namespace-scoped resources and the PVCs, add them to the restore item list by clicking the plus sign (+), and then click **Next** (see Figure 7-13 on page 159).

Figure 7-13 Selecting namespace resources and objects to be restored

3. To restore namespace data from a vSnap server, select **From Copy** → **On-Demand** → **Site**, select the name of the vSnap site with the backup, and select the latest good backup timestamp from the list, as shown in Figure 7-14.

Note: To restore namespace data from cloud storage, select **From Copy** → **On-Demand** → **Cloud service backup**, select the name of the cloud provider with the backup, and select the latest good backup timestamp from the list.

Figure 7-14 Selecting a backup timestamp to be restored

4. Follow the instructions in the wizard by making the following selections:
 - **Do not change the PVC name**
 - **Restore to original cluster**
 - **Use the original storage class**
 - **Use the original namespace**
 - **Use the proposed default job options**

5. Review the summary. If all the information is correct, click **Submit**, as shown in Figure 7-15.

Review

Review your selections, and then click Submit.

Select data sources

Selected source: DPR-OCp:wordpress:mysql-pvc, DPR-OCp:wordpress:resources, DPR-OCp:wordpress:wordpress-pvc

Source snapshot: DPR-OCp:wordpress:mysql-pvc - May 19, 2021 12:18:15 PM
DPR-OCp:wordpress:resources - May 19, 2021 12:18:15 PM
DPR-OCp:wordpress:wordpress-pvc - May 19, 2021 12:18:15 PM

Backup Source Type: From Copy

Restore Type: On-Demand

Restore Source Type: Site

Restore Source: Primary

Restore method: Production

Set destination

Destination: Restore to original cluster

Namespace: <empty>

Storage Class: <empty>

Set run settings

Run cleanup immediately on job failure: Yes

Allow session overwrite: Yes

Continue with restores of other selected resources even if one fails: Yes

Back

Submit

Figure 7-15 Submitting the namespace restore job

For more information about how to perform restores of namespace resources and objects, see Chapter 6, “Using Container Backup Support” on page 121.

Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide more information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only

- ▶ *IBM Spectrum Protect Plus Practical Guidance for Deployment, Configuration, and Usage*, REDP-5532
- ▶ *IBM Spectrum Protect Plus Protecting Database Applications*, REDP-5640
- ▶ *Using the IBM Block Storage CSI Driver in a Red Hat OpenShift Environment*, REDP-5613

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, drafts, and additional materials at the following website:

ibm.com/redbooks

Online resources

The following websites are also relevant as further information sources:

- ▶ IBM block storage Container Storage Interface (CSI) driver:
<https://www.ibm.com/docs/en/blockstg-csi-driver>
- ▶ IBM Container website:
<https://www.ibm.com/cloud/learn/containers>
- ▶ IBM Hybrid cloud website:
<https://www.ibm.com/cloud/learn/hybrid-cloud>
- ▶ IBM Spectrum Protect Plus documentation:
<https://www.ibm.com/docs/en/spp>
- ▶ Red Hat OpenShift Velero plug-in:
<https://github.com/konveyor/openshift-velero-plugin>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Abbreviations and acronyms

ACL	access control list
BYOS	bring your own storage
BaaS	backup as a service
CASE	Container Application Software for Enterprises
CIDR	Classless Inter-Domain Routing
CLI	command-line interface
CP4D	Cloud Pak for Data
CRD	custom resource definition
CSI	Container Storage Interface
CephFS	Ceph File System
DR	disaster recovery
ESCC	EMEA Storage Competence Center
FQDN	Fully Qualified Domain Name
HNAS	Hitachi NAS
IBM	International Business Machines Corporation
OCI	Open Container Initiative
OSSM	Open Snap Store Manager
PVC	persistent volume claim
PV	persistent volume
RBD	Rados Block Device
RHEL	Red Hat Enterprise Linux
RWO	read_write_once
SLA	service-level agreement
VM	virtual machine
VPA	Vertical Pod Autoscaler



REDP-5636-01

ISBN 0738460958

Printed in U.S.A.

Get connected

