Redbooks
ibm.com/redbooks

IBM Storage Fusion Backup and Restore for IBM Cloud Pak for Data

Paulina Acevedo

Austen Stewart

Todd Tosseth

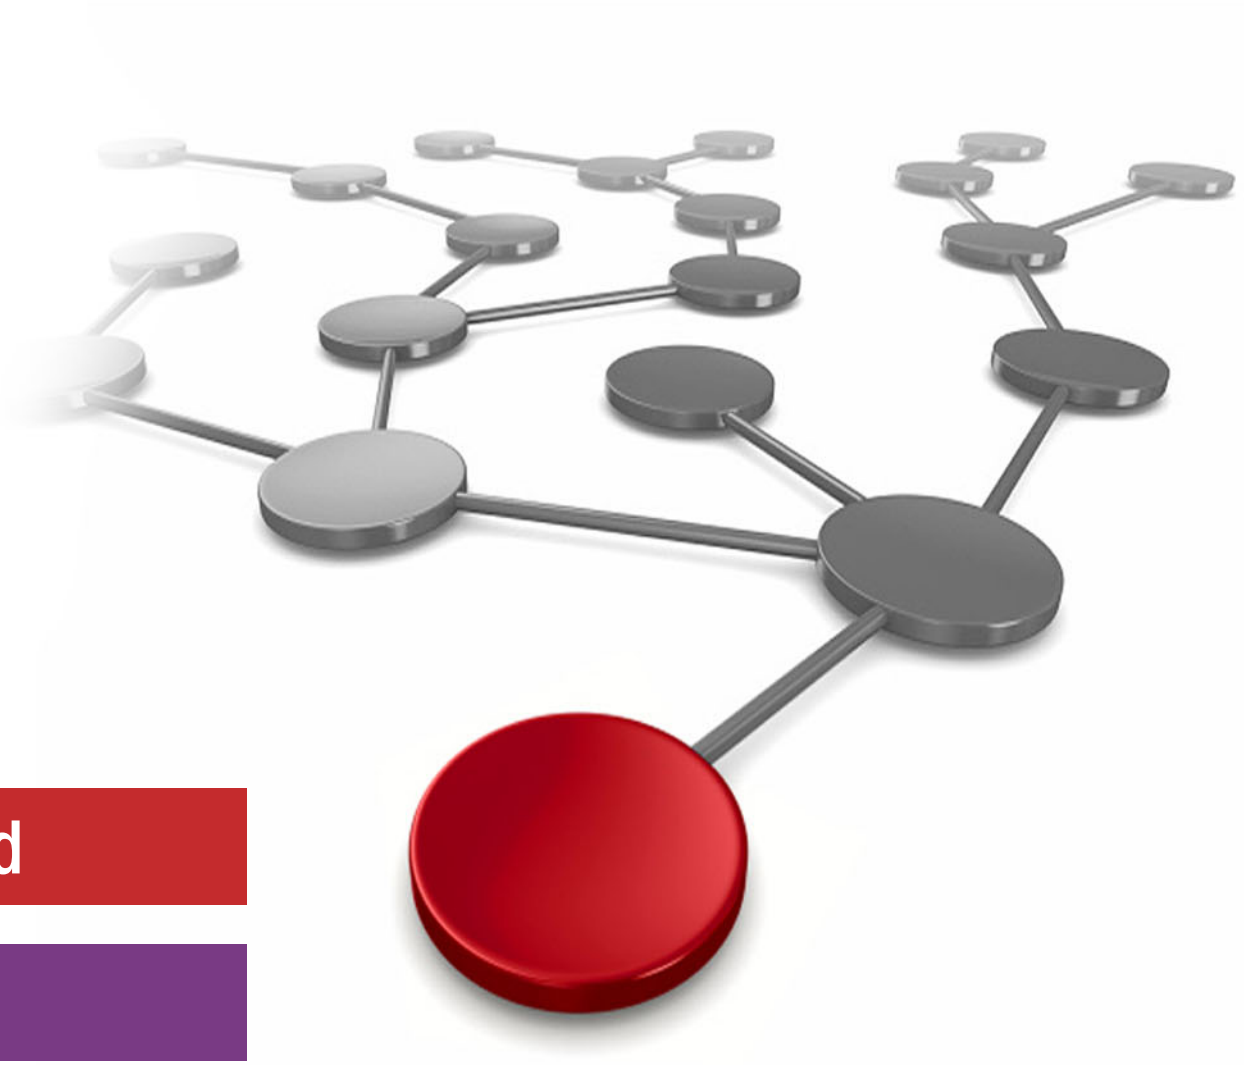

Cloud

Storage

IBM

Redpaper

IBM Redbooks

IBM Storage Fusion Backup and Restore for IBM Cloud Pak for Data

June 2023

REDP-5706-00

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

First Edition (June 2023)

This edition applies to Version 2, Release 5, of IBM Storage Fusion HCI (5771-PP7).

Contents

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 https://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®	Redbooks®
DS8000®	IBM Cloud Pak®	Redbooks (logo) ®
Enterprise Storage Server®	IBM Spectrum®	Spectrum Fusion™
IBM®	IBM Spectrum Fusion™	XIV®

The following terms are trademarks of other companies:

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

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

Preface

IBM Cloud Pak® for Data can be protected with IBM Storage Fusion. This IBM Redpaper publication covers backing up IBM Cloud Pak for Data with a non-disruptive (online) backup and then restoring to an alternate cluster. During an online backup, normal runtime operations in the Cloud Pak for Data cluster continue while the backup completes. The backup process includes creating policies and automating backups in IBM Storage Fusion, then protecting Cloud Pak for Data, protecting IBM Storage Fusion namespace and the IBM Storage Protect Plus (SPP) catalog. Backup and restore is supported from IBM Storage Fusion HCI to IBM Storage Fusion software as well as from IBM Storage Fusion Software to IBM Storage Fusion HCI.

Note: IBM Spectrum® Fusion™ HCI and IBM Spectrum Fusion™ have become IBM Storage Fusion HCI System and IBM Storage Fusion. This edition uses some of the IBM Spectrum brand names and will be updated with the next edition. See Evolving the IBM Storage Portfolio Brand Identity and Strategy to learn more about how IBM Storage Fusion HCI System and IBM Storage Fusion are key to the IBM Storage Portfolio.

IBM Spectrum Fusion must be at 2.3 or higher with “Backup” service installed. If using IBM Storage Fusion 2.5.2, the “Backup (Legacy)” service should be used.

If using Storage Fusion 2.6.x, refer to Chapter 5 Backup and Restore in this Redbooks publication here.

Authors

This paper was produced by a team of specialists from around the world working with the IBM Redbooks, Tucson Center.

Paulina Acevedo is one of two System Test architects for the Cloud Pak Storage Test team. Paulina has been with IBM® for more than 15 years and has held several different positions within the Systems organization. She is a certified Project Manager and has been the System Test Project manager for several products including IBM XIV®, A9000, Spectrum NAS, and Spectrum Virtualize.

Austen Stewart is a computer programmer and tester from Tucson, Arizona. He joined IBM as a co-op in 2020, working on IBM DS8000® storage testing automation. After returning full-time in 2021, Austen began working on storage testing for Cloud Paks. He holds a BS degree in Computer Science from University of Arizona.

Todd Tosseth is a Software Engineer for IBM in Tucson, Arizona. Joining IBM in 2001, he has worked as a test and development engineer on several IBM storage products, such as DS8000, IBM Spectrum Scale, and IBM Enterprise Storage Server®. He is working on IBM Cloud® Pak as a system test engineer, with an emphasis on Cloud Pak storage integration.

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

Larry Coyne
IBM Redbooks®, Tucson Center

Nazar Abbas, Tara Astigarraga, Henry Chiu, Karli Collins, Matt Divito, Paul Frost, Michael Fruchtman, AshaRani G R, Lisa Huston, Venkitesh Jagadeesan, Frank Lautenbach, Tom Lee, Hon Peng Leong, Boda Devi Manikanta, JJ Miller, Mitch Montanez, Savitha H N, Lu Nguyen, Shyamala Rajagopalan, Dessa Simpson, Shirley Shum, Jim Smith, Andy Streit, Hemalatha B T, Thiha Than, Jayson Tsingine, Israel Vizcarra, Henning Wilbert
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:

  https://www.linkedin.com/groups/2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/subscribe

- Stay current on recent Redbooks publications with RSS Feeds:

  https://www.redbooks.ibm.com/rss.html

Summary of changes

This section describes the technical changes made in this edition of the paper and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes
for IBM Storage Fusion Backup and Restore for Cloud Pak for Data
as created or updated on October 24, 2023.

October 2023, First Edition minor updates

This revision includes the following new and changed information.

New information

- Updated supported IBM Storage Fusion levels in 1.1, “Considerations and requirements” on page 2 with the following:

  The following must be considered when preparing to protect your Cloud Pak for Data environment on Storage Fusion 2.3.x, 2.4.x, and 2.5.x. If using Storage Fusion 2.6.x, refer to Chapter 5 Backup and Restore in this Redbooks publication here.

Changed information

- Updated IBM Spectrum to IBM Storage
- Updated urls

1

Preparation

This chapter describes the considerations, requirements, and pre-requisites to set up IBM Storage Fusion backup and recovery for the Cloud Pak for Data environment.

1.1 Considerations and requirements

The following must be considered when preparing to protect your Cloud Pak for Data environment on Storage Fusion 2.3.x, 2.4.x, and 2.5.x. If using Storage Fusion 2.6.x, refer to Chapter 5 Backup and Restore in the Redbooks publication here.

- If using Storage Fusion 2.3.x or 2.4.x, 2.5.0 or 2.5.1, the “Backup” service should be installed and utilized
- If using IBM Storage Fusion 2.5.2, the “Backup (Legacy)” service should be installed and utilized
- IBM Storage Fusion Data Protection is enabled (which deploys IBM Storage Protect Plus)
- OpenShift OCP 4.8 and OCP 4.10 are supported, but both source and target clusters must be at the same major version
- IBM Cloud Pak for Data must be at version 4.5.3 or higher
- All services are installed at the same Cloud Pak for Data release
- The Cloud Pak for Data control plane is installed in a single project (namespace)
- Backup and restore of the Cloud Pak for Data instance with tethered namespaces is supported

1.2 Getting pre-requisites ready

Before backing up and protecting the Cloud Pak for Data environment, some pre-requisites must be prepared.

1.2.1 Installing the cpdbr-oadp service

The first item to prepare is the `cpdbr-oadp` service in the Cloud Pak for Data operators and Cloud Pak for Data instance namespaces.

To install the services, prepare your environment to use cpd-cli, which can be found at the following location: https://github.com/IBM/cpd-cli/releases

Install the `cpdbr-oadp` service in the Cloud Pak for Data operators and instance namespaces. Reference the Cloud Pak for Data link:
https://www.ibm.com/docs/en/cloud-paks/cp-data/4.7.x?topic=utilities-installing-cpdbr-service-storage-fusion-integration

In the sample environment, it is an Express install, which has foundational services and Cloud Pak for Data operators in the same **ibm-common-services** namespace. The `cpdbr-oadp` service was installed by issuing the following command, as shown in Figure 1-1.

```
./cpd-cli oadp install --foundation-namespace=ibm-common-services
--operators-namespace=ibm-common-services --component=cpdbr-ops-hooks
--cpdbr-hooks-image-prefix=icr.io/cpopen/cpd --log-level=debug -verbose
```

```
$ ./cpd-cli oadp install --foundation-namespace=ibm-common-services --operators-namespace=ibm-common-services --
component=cpdbr-ops-hooks --cpdbr-hooks-image-prefix=icr.io/cpopen/cpd --log-level=debug --verbose
processing request...
foundation namespace: ibm-common-services
operators namespace: ibm-common-services

spp cpdbr-ops installed
$
```

Figure 1-1 Install the cpdbr-oadp service in the Cloud Pak for Data operators namespace

Afterwards, we issued **`oc get pod -n ibm-common-services |grep cpdbr`** to ensure the cpdbr-ops-service pod is running, as seen in Figure 1-2.

```
$ oc get pod -n ibm-common-services |grep cpdbr
cpdbr-ops-service-54c54d797c-bhb5z                              1/1     Running     0          44h
$
```

Figure 1-2 Command to ensure the cpdbr-ops-service pod is running

The install of the `cpdbr-oadp` service will also generate and apply the `ibmcpd-operators` for the IBM Storage Fusion recipe in the Cloud Pak for Data operators namespace, as seen in Figure 1-3.

```
$ oc get frcpe -n ibm-common-services
NAME                AGE
ibmcpd-operators    45h
$
```

Figure 1-3 Generate and apply the ibmcpd-operators

Next, the `cpdbr-oadp` service is installed in the Cloud Pak for Data instance namespace. In the sample environment, the Cloud Pak for Data instance is in the **czen** namespace. The `cpdbr-oadp` service was installed by issuing the following command, as seen in Figure 1-4.

```
./cpd-cli oadp install --cpd-namespace=czen --component=cpdbr-hooks
--cpdbr-hooks-image-prefix=icr.io/cpopen/cpd --log-level=debug -verbose
```

```
$ ./cpd-cli oadp install --cpd-namespace=czen --component=cpdbr-hooks --cpdbr-hooks-image-prefix=icr.io/cpopen
/cpd --log-level=debug --verbose
processing request...
cpd namespace: czen

spp cpdbr installed
$
```

Figure 1-4 Install cpdbr-oadp service in the Cloud Pak for Data instance namespace

Afterwards, we issued **`oc get pod -n czen |grep cpdbr`** to ensure the cpdbr-service pod is running, as seen in Figure 1-5.

```
$ oc get pod -n czen |grep cpdbr
cpdbr-service-795956fd87-zsbg8   1/1     Running   0          45s
$
```

Figure 1-5 Ensure the cpdbr-service pod is running

The install of the `cpdbr-oadp` service will also generate and apply the **ibmcpd** IBM Storage Fusion recipe in the Cloud Pak for Data instance namespace, as seen in Figure 1-6.

```
$ oc get frcpe -n czen
NAME      AGE
ibmcpd    20m
$
```

Figure 1-6 ibmcpd IBM Storage Fusion recipe in the CP4D instance namespace

Moreover, the installation of `cpdbr-oadp` service also triggers the IBM Storage Fusion application to then add additional information to the *spec* area of the Cloud Pak for Data instance application yaml, such as the variables that will be used for the recipe to be able to perform the backup and the restore.

We can view the `czen` yaml by issuing **`oc get fapp czen -n ibm-spectrum-fusion-ns -o yaml`**, as depicted in Figure 1-7.

```
$ oc get fapp czen -n ibm-spectrum-fusion-ns -o yaml
apiVersion: application.isf.ibm.com/v1alpha1
kind: Application
metadata:
  annotations:
    isf.ibm.com/created-event: created
    spp-data-protection.isf.ibm.com/recipe: '{"name":"ibmcpd","namespace":"czen","kind":"Recipe","apiVersion":"s
pp-data-protection.isf.ibm.com/v1alpha1"}'
    usedcapacity: "725614592"
  creationTimestamp: "2022-11-09T20:12:20Z"
  generation: 2
  name: czen
  namespace: ibm-spectrum-fusion-ns
  resourceVersion: "83221257"
  uid: 22556b0f-f50b-4929-b9b1-d12d01c87593
spec:
  enableDR: false
  homeCluster: apps.rackb-fusion.tuc.stglabs.ibm.com
  includedNamespaces:
  - czen
  variables:
  - name: PARENT_NAMESPACE
    value: czen
status:
  activeRestore: {}
  availableBackups: 0
  backupStatus: □
  drStatus:
    aggregateStatus: not-protected
    conditions:
    - lastTransitionTime: "2022-11-09T20:12:20Z"
      message: ""
      observedGeneration: 1
      reason: AsExpected
      status: "False"
      type: Protected
```

Figure 1-7 View the czen yaml

1.2.2 Setting up Object Storage

The next step is to add a backup storage location which is required before we set up backup policies. To add a storage location, we go to the **Backup policies** page and click on **Add location** as seen in Figure 1-8 on page 5.

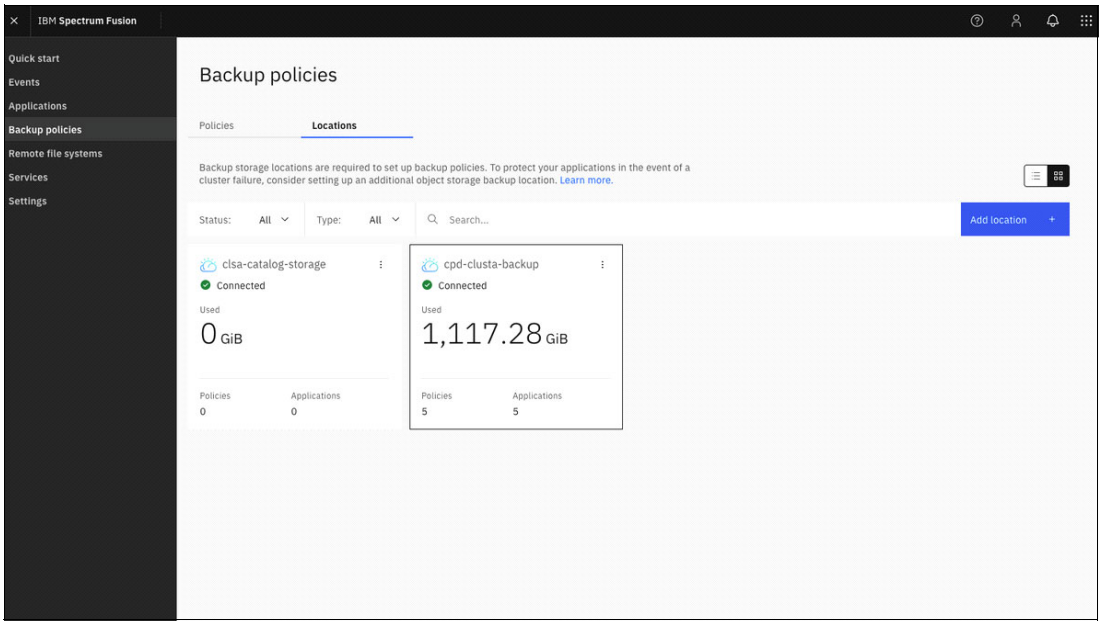

Figure 1-8 Add storage location from Backup policies page by selecting Add location

In our example, we named our location **cpst-cos-hciops** and we will be using IBM Cloud Object Storage. We then entered the endpoint and bucket and added the access key and secret key as seen in Figure 1-9.

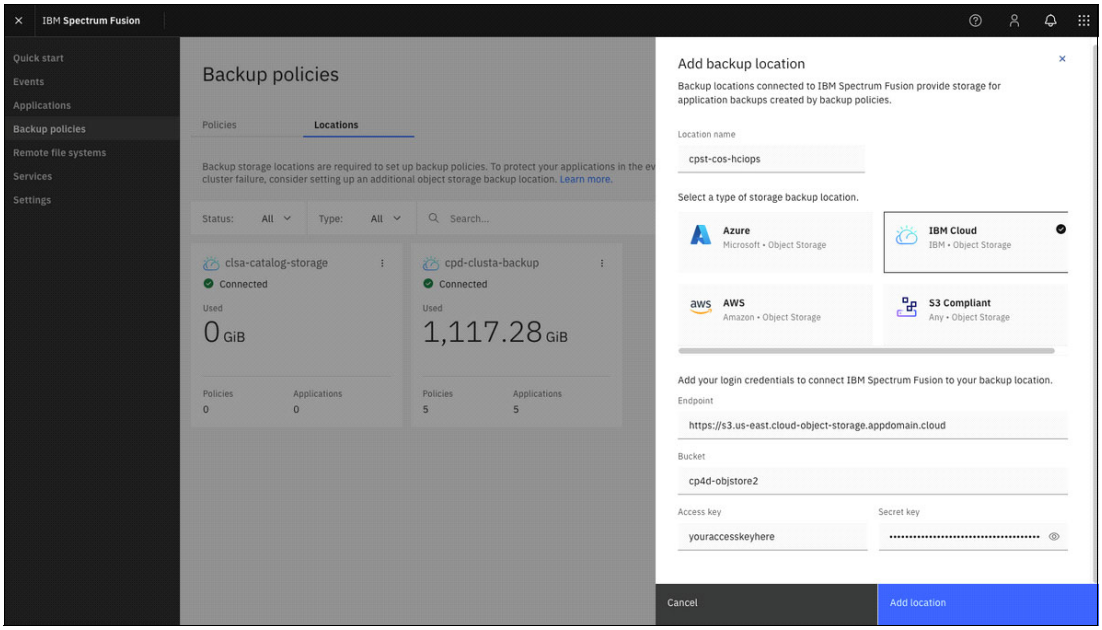

Figure 1-9 Complete Add backup location credentials then click on Add location

Once we entered the login credentials, we clicked on **Add location** and we can now see it listed as one of the Locations within the **Backup policies** page as seen in Figure 1-10.

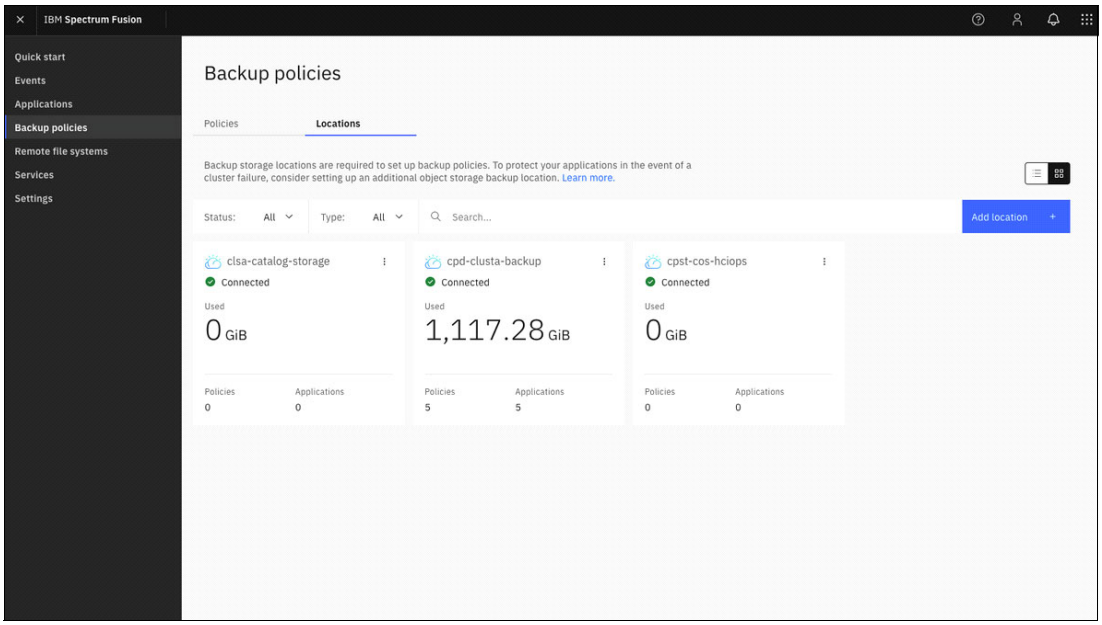

Figure 1-10 cpst-cos-hciops listed as backup location

1.2.3 Backup policies for Cloud Pak for Data applications

Prior to creating and applying backup policies to the Cloud Pak for Data applications, you must create an S3 compliant backup location in the IBM Storage Fusion UI, which will be used to store the backups and restore data from.

The next step is to log into IBM Storage Fusion UI and go to the **Backup policies** page to create a backup policy which will use an S3 object storage location. From that page, we clicked on **Add policy** and we named the policy **czen-policy** and selected a monthly frequency as seen in Figure 1-11.

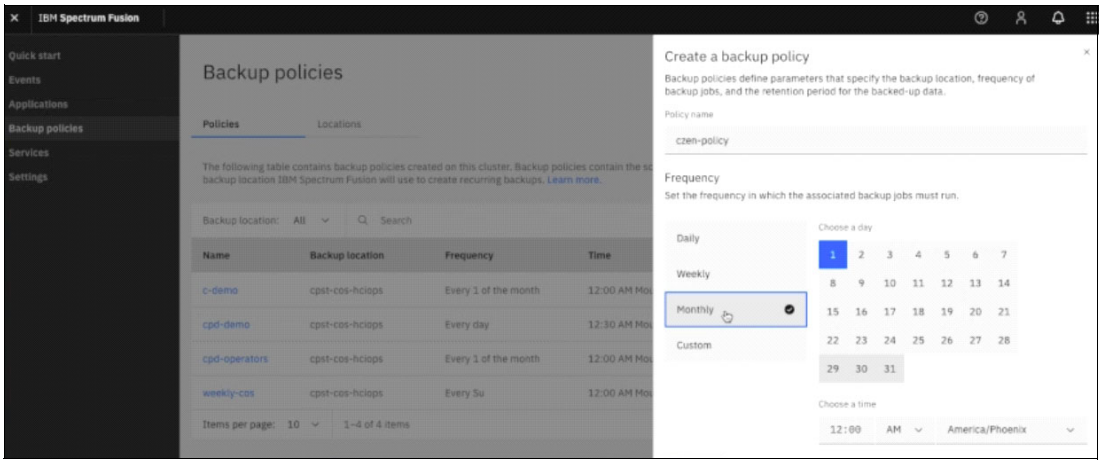

Figure 1-11 Create the backup policy czen-policy

Under the **Backup Locations** section, we selected **Object storage** as seen in Figure 1-12.

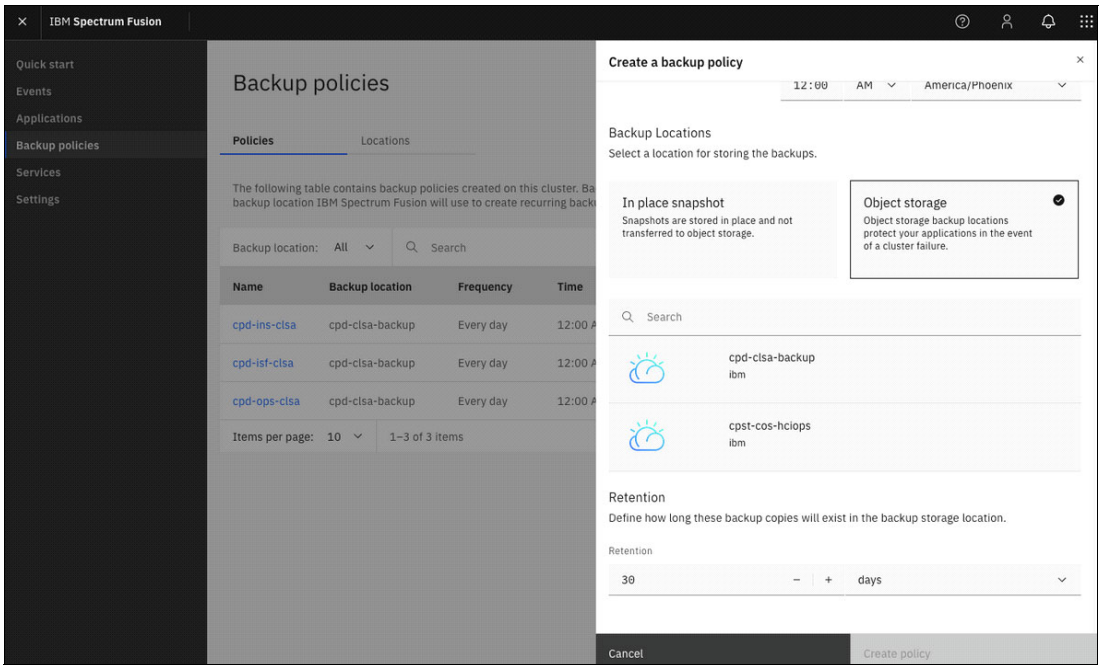

Figure 1-12 Select Object storage for backup location

Once the policy is created, it will appear in the **Backup Policies** page under the **Policies** section as seen in Figure 1-13.

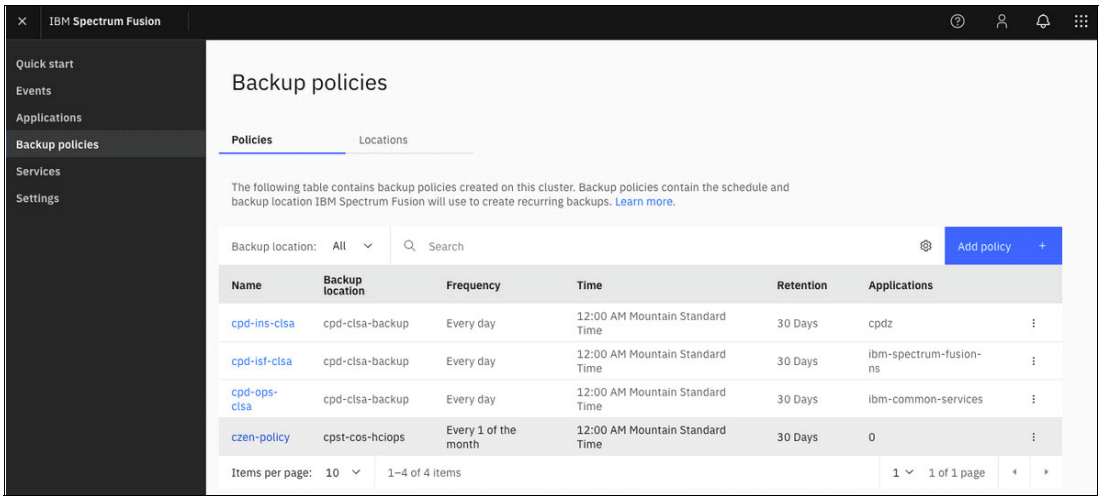

Figure 1-13 czen-policy now listed in Backup Policies page

The next step is to go to the IBM Storage Fusion **Applications** page and assign the policy to your application. In our example, we assigned the **czen-policy** we just created to the **czen** application by clicking on **Assign policy** which is under the **Backup policy** column as seen in Figure 1-14.

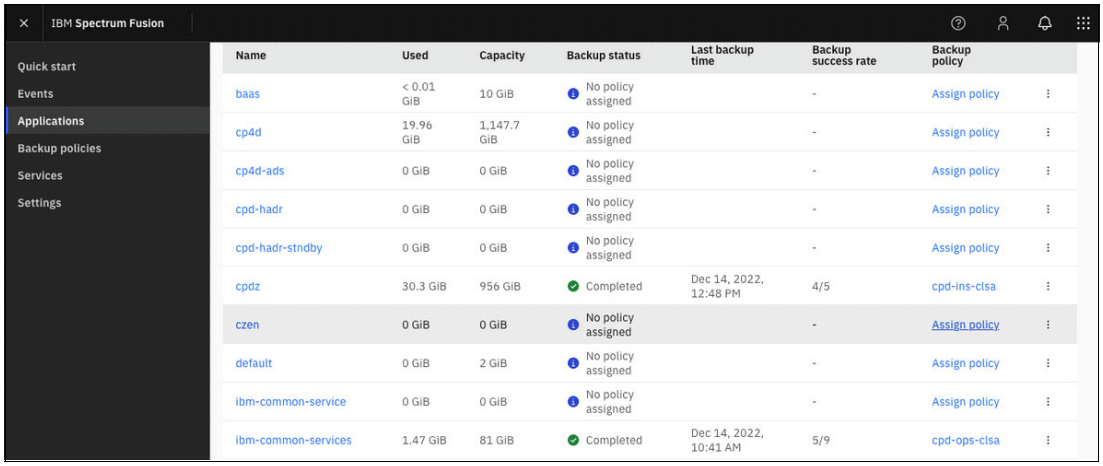

Figure 1-14 Select the Assign policy for the czen application

From the **Assign a backup policy** menu, select the policy you want to apply and click **Assign policy**. In our example, we selected the newly created **czen-policy** as show in Figure 1-15.

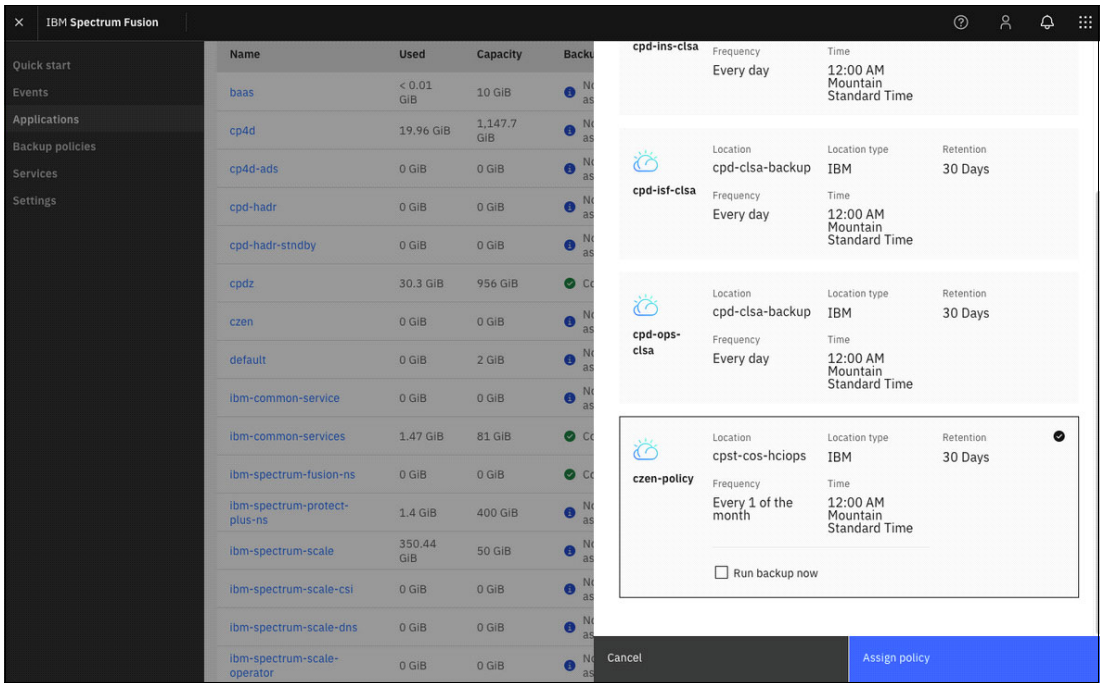

Figure 1-15 Assign the czen-policy to the czen namespace

After assigning the policy, now we can see that the **czen-policy** is applied to the **czen** namespace from the Applications page as seen in Figure 1-16 on page 9.

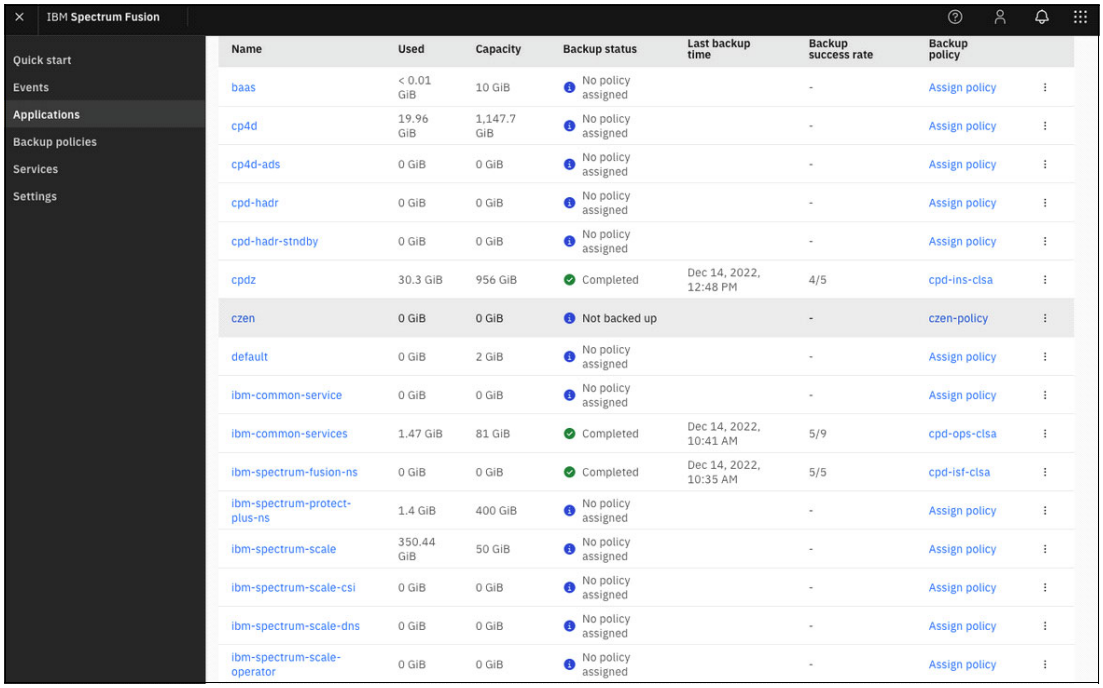

Figure 1-16 Page shows czen-policy assigned to czen namespace

In the **Backup Polices** page, we can also verify that the backup policy has been applied to **czen** namespace/application as seen in Figure 1-17.

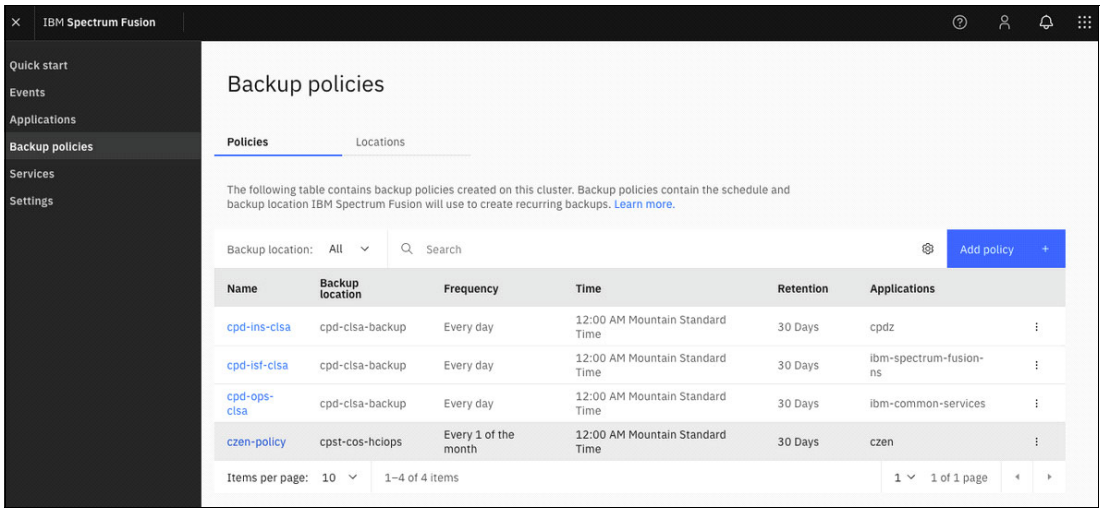

Figure 1-17 Verify that the backup policy has been applied to czen namespace

We can also list the policies and view the objects that get created via command line by issuing the following command: **`$ oc get fbp, fpa`** as shown in Figure 1-18.

```
20220831-21:43:58 [ocp-a:ibm-spectrum-fusion-ns] ~ $ oc get fbp,fpa
NAME                                                PROVIDER    BACKUPSTORAGELOCATION   SCHEDULE      RETENTION   RETENTIONUNIT
backuppolicy.data-protection.isf.ibm.com/c-demo     isf-ibmspp  cpst-cos-hciops         00 0 1 * *    30          days
backuppolicy.data-protection.isf.ibm.com/cpd-demo   isf-ibmspp  cpst-cos-hciops         30 0 * * *    5           days
backuppolicy.data-protection.isf.ibm.com/cpd-operators isf-ibmspp cpst-cos-hciops       00 0 1 * *    30          days
backuppolicy.data-protection.isf.ibm.com/czen-policy isf-ibmspp cpst-cos-hciops         00 0 1 * *    30          days
backuppolicy.data-protection.isf.ibm.com/weekly-cos isf-ibmspp  cpst-cos-hciops         00 0  * * 0   30          days

NAME                                                                    PROVIDER    APPLICATION          BACKUPPOLICY    RECIPE          RECIPENAMESPACE     PHASE     LASTBA
CKUPTIMESTAMP   CAPACITY
policyassignment.data-protection.isf.ibm.com/cpd-demo-c-demo            isf-ibmspp  cpd-demo             c-demo          ibmcpd          cpd-demo            Assigned
                <no value>
policyassignment.data-protection.isf.ibm.com/cpst-zen-demo-cpd-demo     isf-ibmspp  cpst-zen-demo        cpd-demo        ibmcpd          cpst-zen-demo       Assigned  12h
                <no value>
policyassignment.data-protection.isf.ibm.com/czen-czen-policy           isf-ibmspp  czen                 czen-policy     ibmcpd          czen                Assigned
                <no value>
policyassignment.data-protection.isf.ibm.com/ibm-common-services-cpd-operators isf-ibmspp ibm-common-services cpd-operators ibmcpd-operators ibm-common-services Assigned  34h
                <no value>
```

Figure 1-18 Command line list of policies and objects created

From that command we can see the backup policy we created which is called **czen-policy** and the policy assignment of the **czen** namespace. Furthermore, we can also see the **Recipe** column lists the **ibmcpd** recipe.

Additionally, the **ibm-common-services** namespace was assigned the **cpdops-policy** as seen in Figure 1-19.

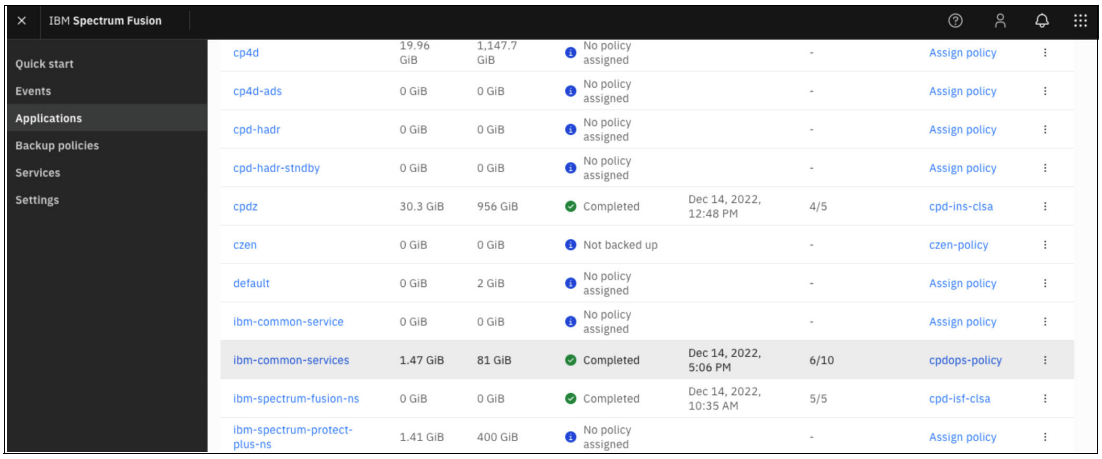

Figure 1-19 cpdops-policy was also assigned to the ibm-common-services namespace

1.2.4 Backup policy for IBM Storage Fusion application

In the Fusion UI, create a new Backup policy for the IBM Storage Fusion application backup and restore resources. In this case, policy **cpd-isf** was created, as seen in Figure 1-20 on page 11.

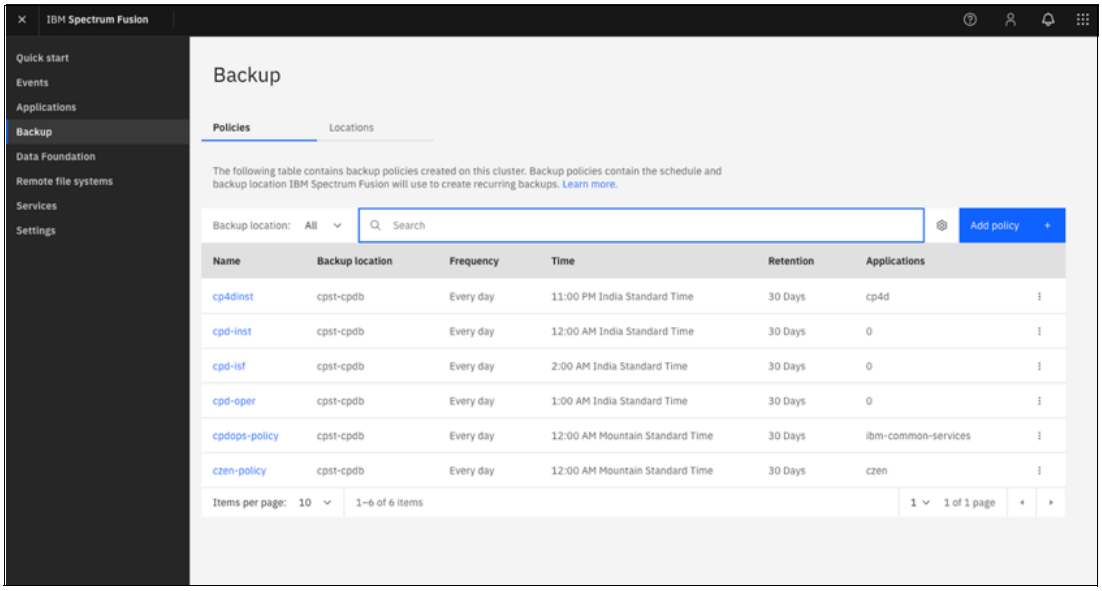

Figure 1-20 Backup policy cpd-isf created to backup the IBM Storage Fusion resources

Then assign the policy to the IBM Storage Fusion application. In this example, ibm-spectrum-fusion-ns. Seen in Figure 1-21.

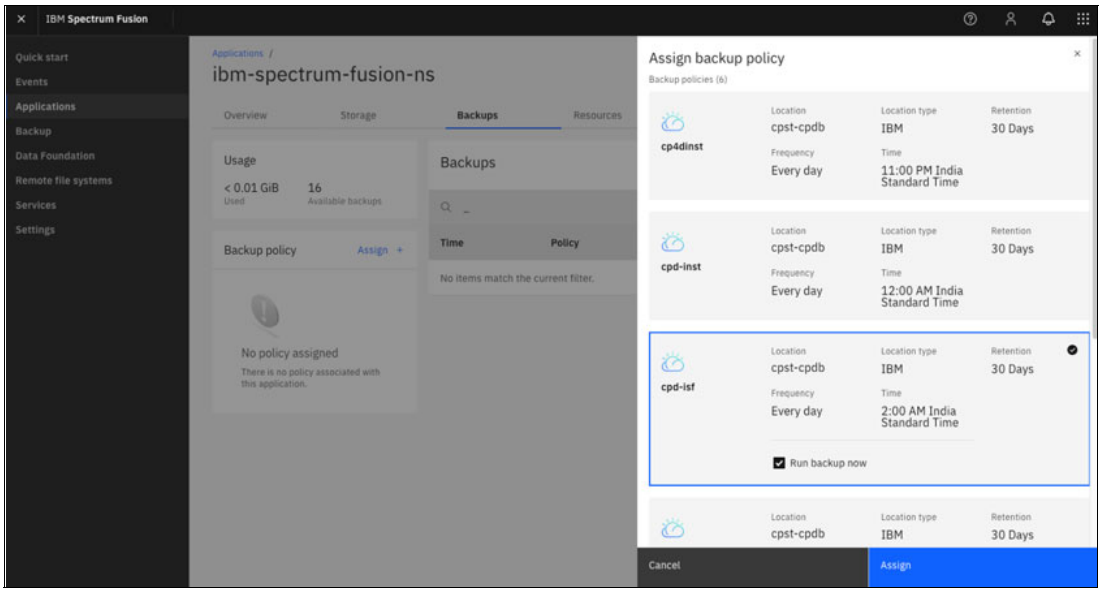

Figure 1-21 Assign backup policy for ibm-spectrum-fusion-ns Application

1.2.5 Backup policy for the IBM Storage Protect Plus catalog

Log into the IBM Storage Protect Plus user interface (SPP UI) to create the SPP catalog backup policy. To get the credentials, you need to get them from the OCP console and go to **Workloads** → **Secrets** → **spp-connection**. That page lists the URL for IBM Storage Protect Plus as well as the username and password.

From the SPP UI, go to **Manage Protection** → **Policy Overview** and click **Add SLA Policy**. Create new SLA policy of type IBM Storage Protect Plus catalog, as seen in Figure 1-22.

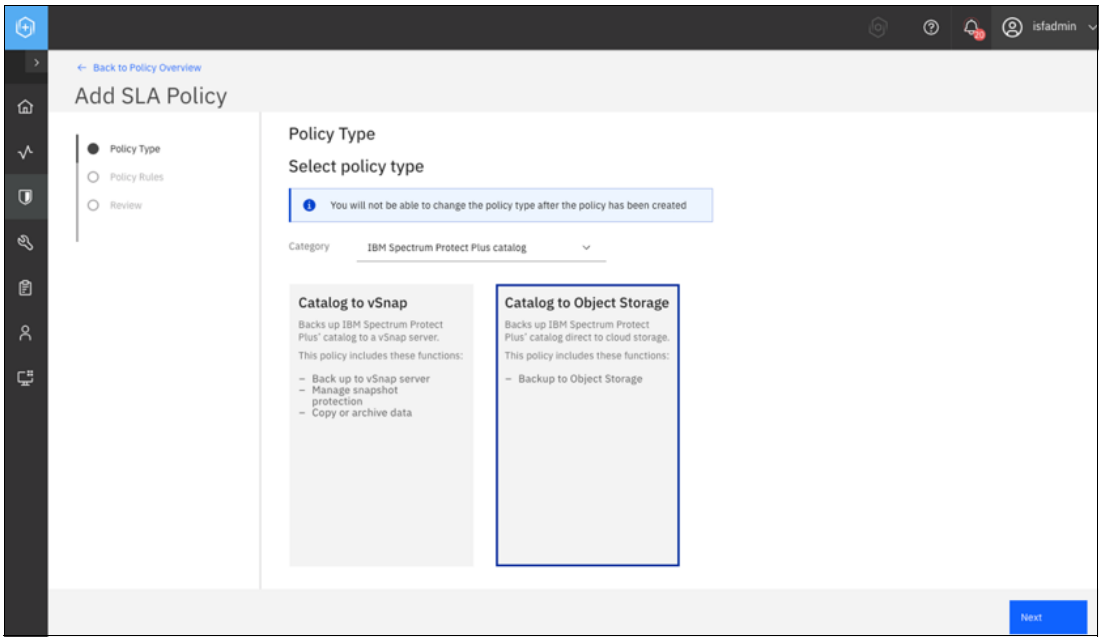

Figure 1-22 Create SLA Policy to backup the SPP catalog to Object Storage

Select the desired policy rules, such as frequency and retention. Note, the start time must be after the completion time of the Cloud Pak for Data operators, Cloud Pak for Data instance, and IBM Storage Fusion backups, to capture the most recent backups. Select the desired object storage location, which can be accessed by the target cluster. Seen in Figure 1-23.

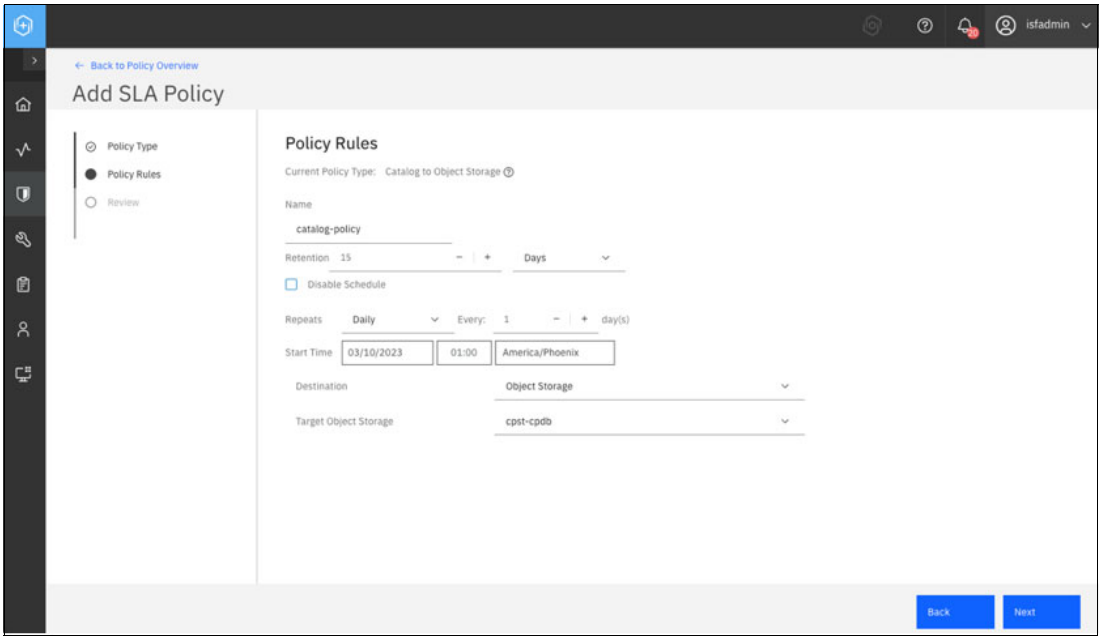

Figure 1-23 Define the policy rules and storage destination for the SPP catalog backup policy

Now assign the policy for backing up the SPP catalog. In the SPP console, navigate to **Manage Protection** → **IBM Spectrum Protect Plus** → **Backup**. Find the newly created policy and select it. Click **Save**. Seen in Figure 1-24 on page 13.

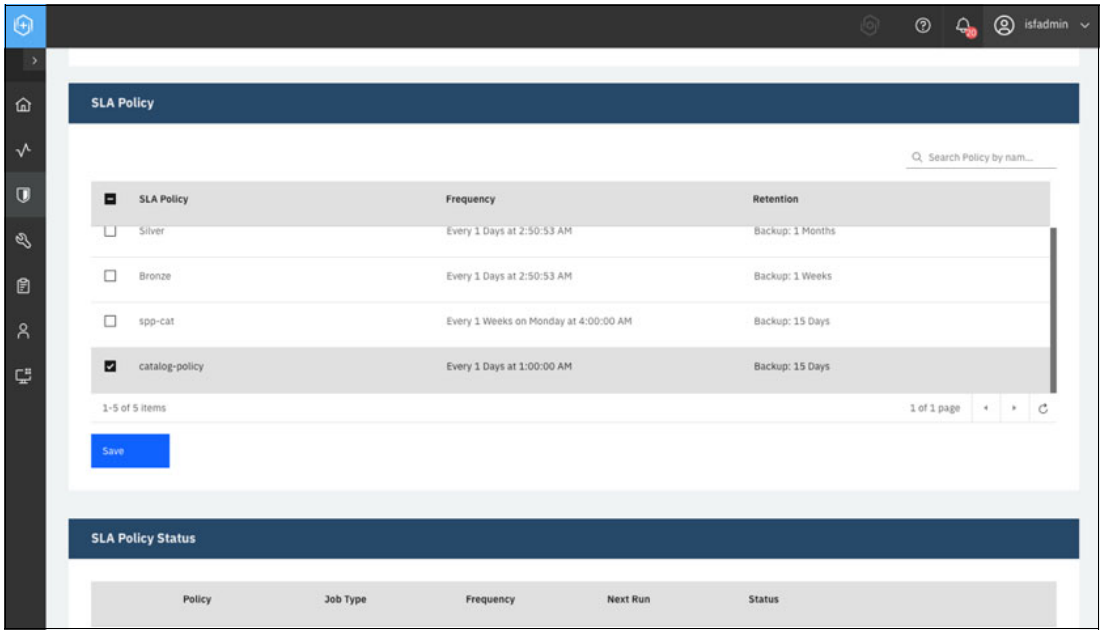

Figure 1-24 Assign the SLA policy to backup the SPP catalog

At this point the SPP catalog will backup on the selected schedule.

1.2.6 IBM Storage Protect Plus access credentials for source cluster

The user must document the access credentials for IBM Storage Protect Plus for the source cluster, in the event the source cluster goes down and will be restored to an alternate cluster.

The credentials reside in the **spp-connection** secret in the IBM Storage Fusion namespace. An example command to obtain the credentials is the following:

```
oc extract secret/spp-connection --to=- -n ibm-spectrum-fusion-ns
```

2

Backing up the source cluster

Now that the Cloud Pak for Data and the IBM Storage Fusion applications, as well as the IBM Storage Protect Plus catalog, have backup policy assignments, they will begin backing up according to their respective backup policy frequencies. However, an on demand backup of all of the elements can be performed.

This chapter describes the steps to backup the Cloud Pak for Data operators and instance, IBM Storage Fusion namespace, and the IBM Storage Protect Plus catalog. In our example, we installed Cloud Pak for Data platform 4.5.3 with IBM Db2® service. Timings for backup and restore are based on our specific deployment.

2.1 Backing up the Cloud Pak for Data operators

From the IBM Storage Fusion UI, we can start the backup of the IBM Cloud Pak for Data operators. In the sample case, we used the IBM Common Services express installation. To start the backup, we went to **Applications** and selected **ibm-common-services**. Then clicked on the *Backups* tab at the top of the screen. Then, clicked the **Actions** drop-down menu and selected **Backup now** as seen in Figure 2-1.

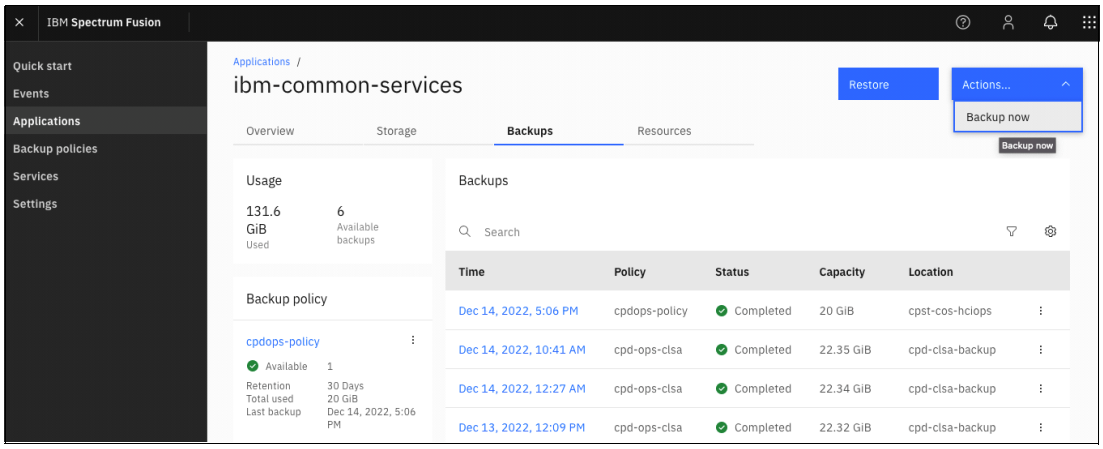

Figure 2-1 Select a backup

Then from the **Run a backup now** window, click on **Run backup policy** as shown in Figure 2-2.

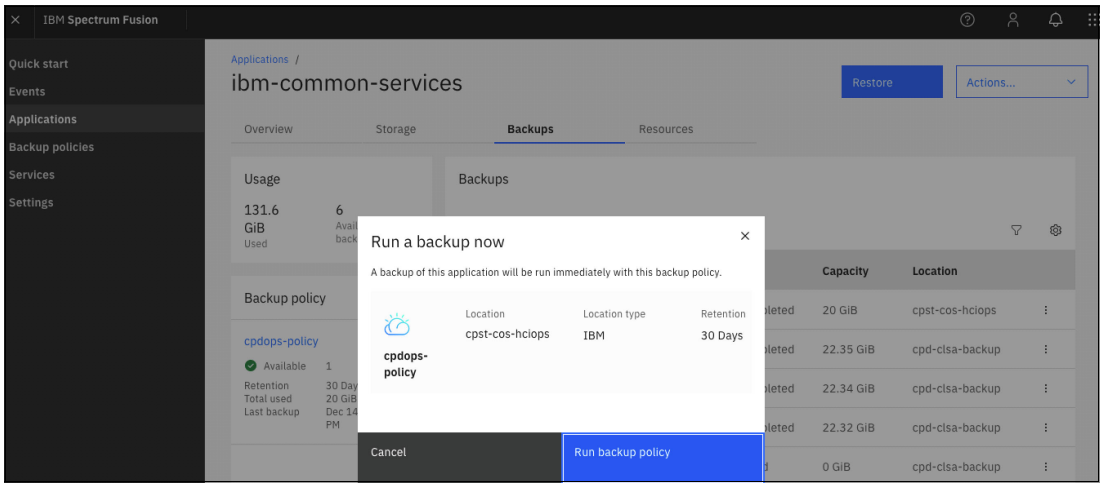

Figure 2-2 Run the backup policy cdpops-policy immediately

The next step is to log into the IBM Storage Protect Plus user interface (SPP UI) to watch the progress of the backup job. To get the credentials, you need to get them from the OCP console and go to **Workloads** → **Secrets** → **spp-connection**. That page lists the URL for IBM Storage Protect Plus as well as the username and password.

After logging into the SPP user interface, click on the **Jobs and Operations** icon on the left-hand side, and in the **Running Jobs** tab, we can see the progress of the scheduled backup job as seen in Figure 2-3 on page 17.

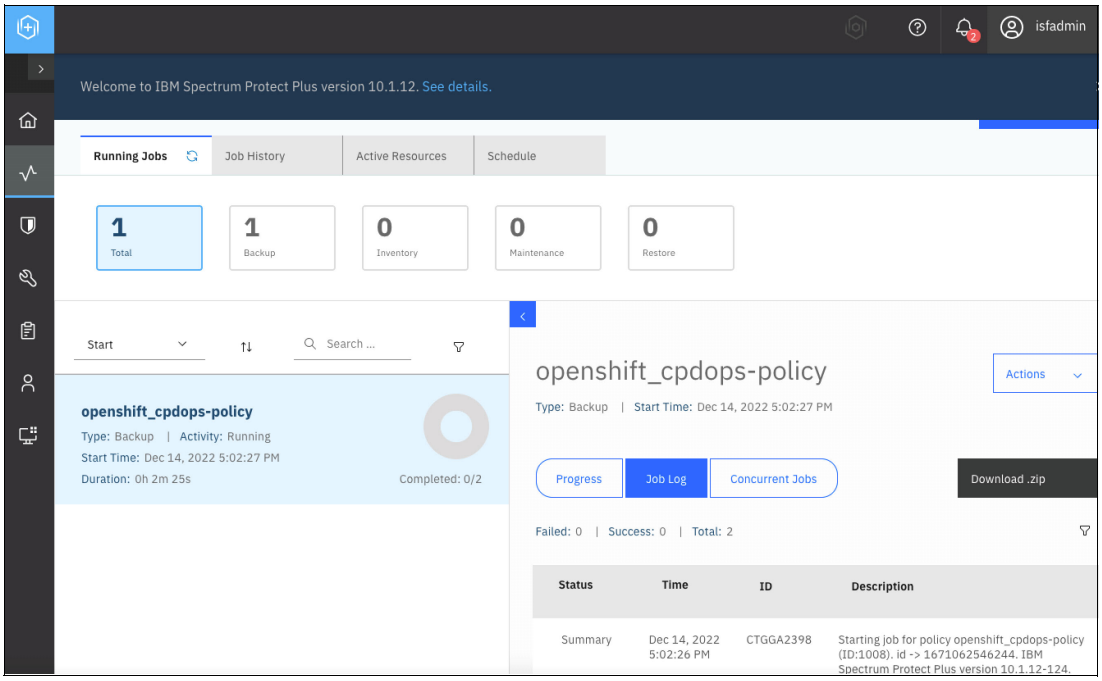

Figure 2-3 View the jobs running

After the job completes, we can go to the **Job History** tab and from there, we can see there are two phases to the backup. The first phase is the **Snapshot** which performs the local snapshot of all the PVCs and resources. In our case, this phase took four minutes and 21 seconds. The second phase is the **Backup** which is a copy of the backup and this phase uploads all information to the selected storage location and this phase took six minutes and 22 seconds. In total, the backup of Cloud Pak for Data operators took ten minutes and 43 seconds to complete as seen in Figure 2-4.

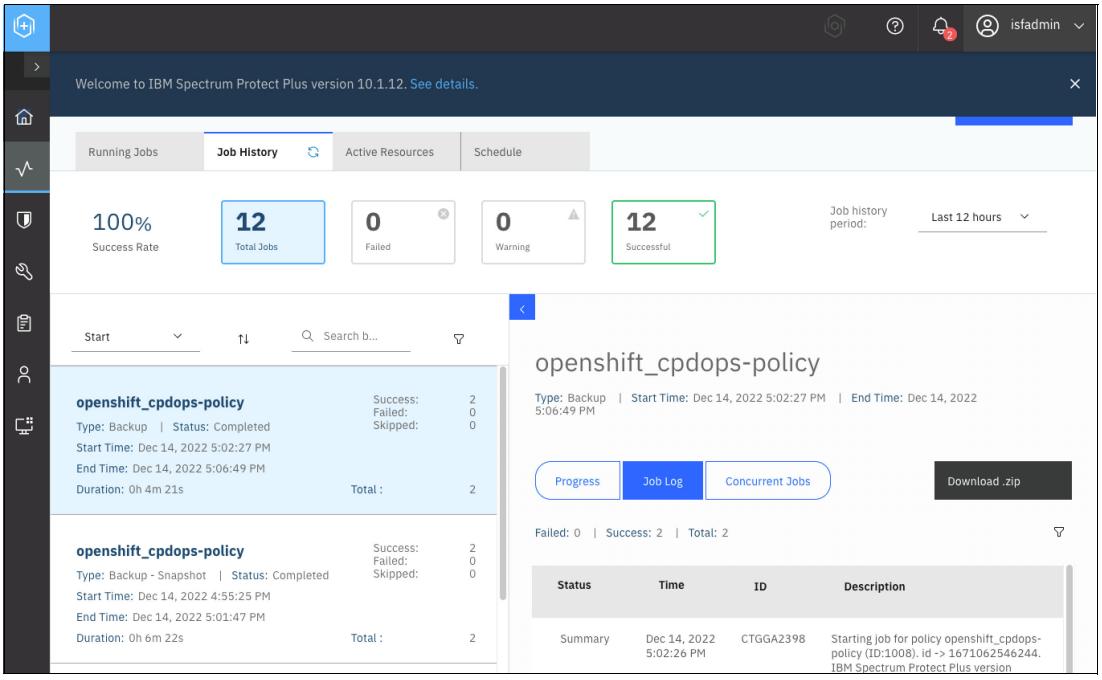

Figure 2-4 View statistics for snapshot and backup

Once the backup completes, we can see the **Backup status** listed as **Completed** in the **Applications** page of IBM Storage Fusion which shows the status of the most recent backup as shown in Figure 2-5.

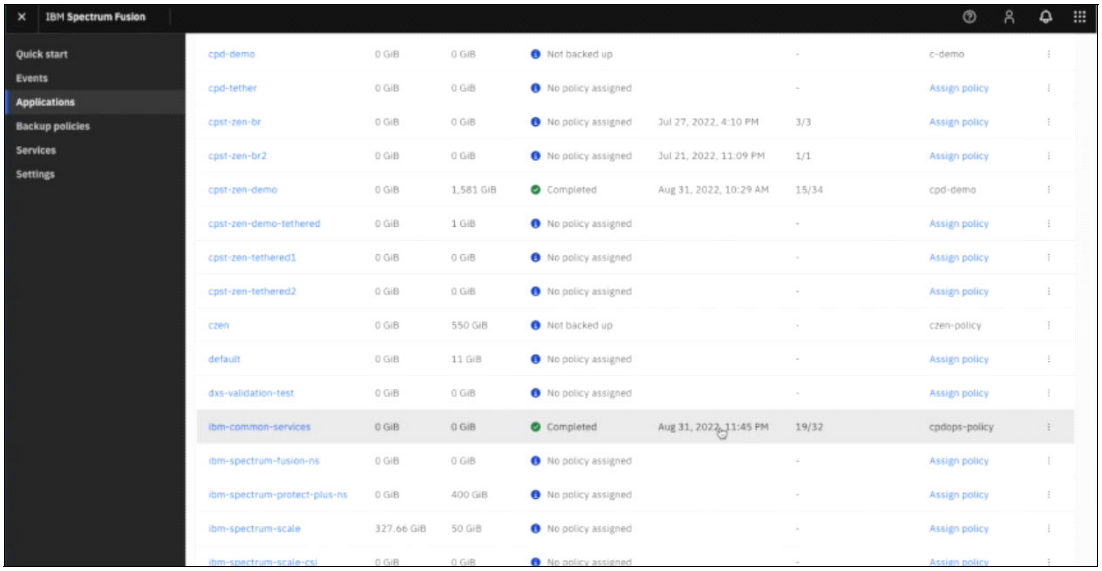

Figure 2-5 Status of most recent backups

2.2 Backing up the Cloud Pak for Data instance

Now we are ready to do the backup of the Cloud Pak for Data instance, which in our example, the application/namespace is called **czen**. To do so, we went to the **Applications** page within IBM Storage Fusion and clicked on **czen** as shown in Figure 2-6.

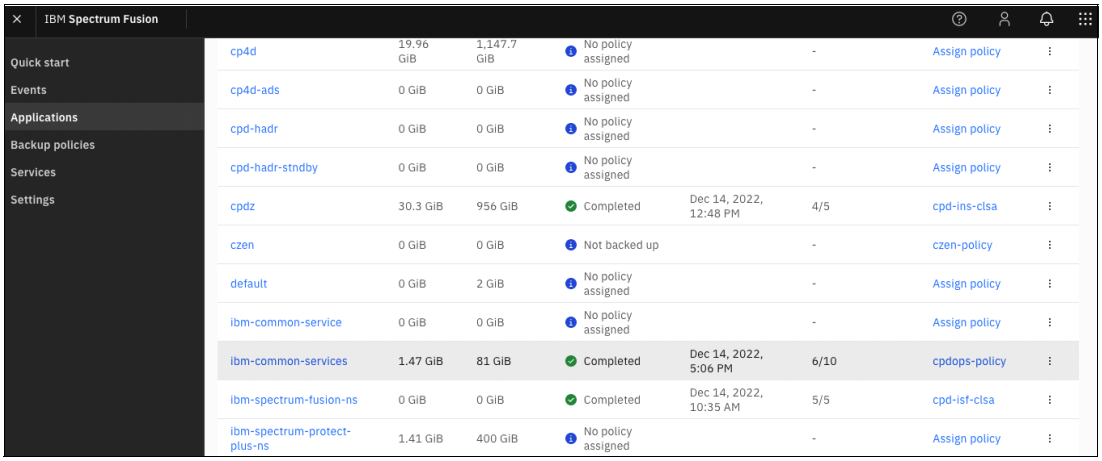

Figure 2-6 Select namespace czen

Once, inside the **czen Applications** page, we clicked on the **Backups** tab and from the **Actions** drop-down menu, selected **Backup now** as show in Figure 2-7 on page 19.

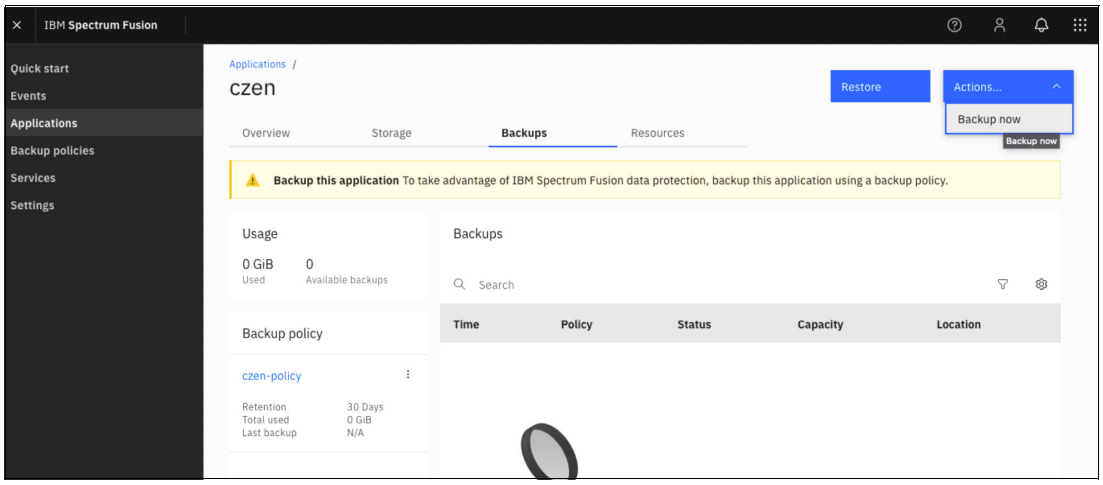

Figure 2-7 Selected Backup now from the czen Applications page

Then from the **Run a backup now** window, we clicked on **Run backup policy** which will utilize the backup policy that was just created as shown in Figure 2-8.

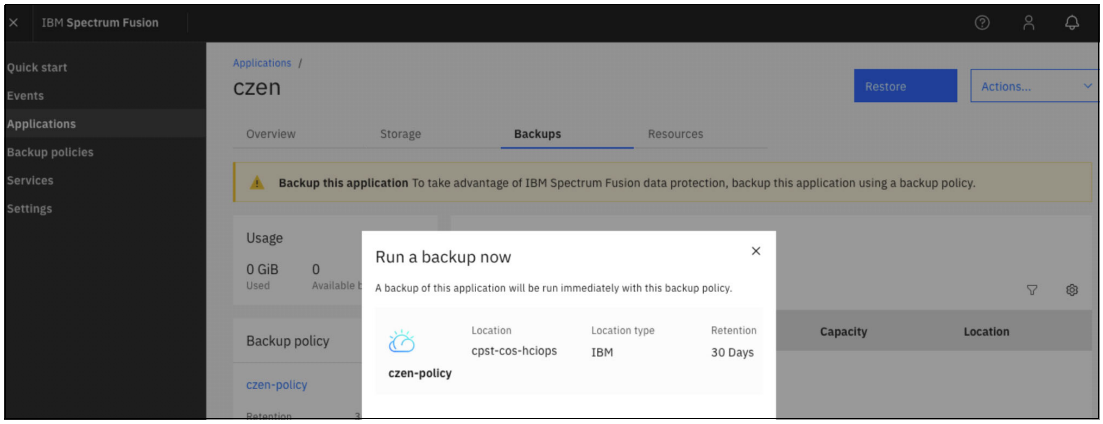

Figure 2-8 Run the backup now for czen-policy

Once it starts, we can see the snapshot creation is in progress as depicted in Figure 2-9.

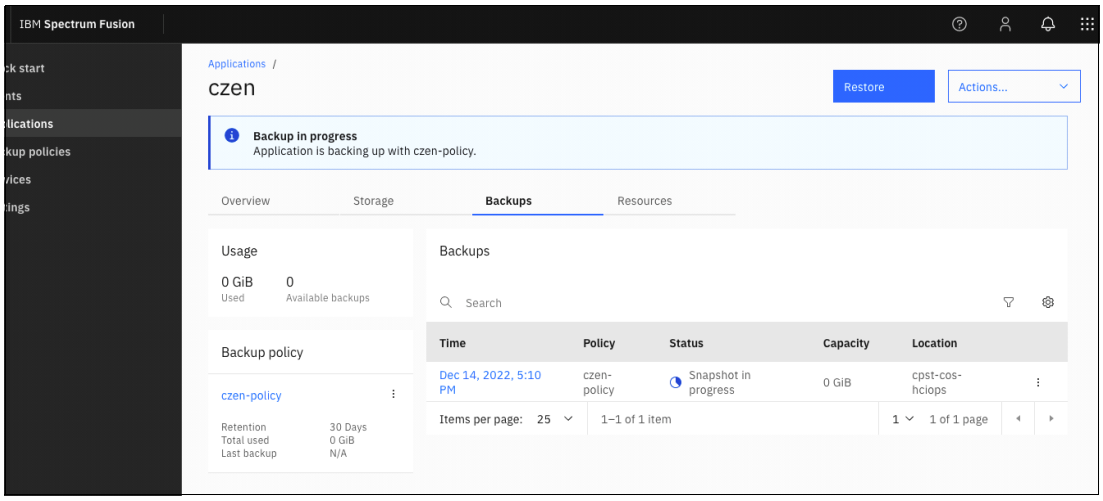

Figure 2-9 View the Snapshot in progress status

To see the progress of the backup, we went to the IBM Storage Protect Plus UI then went to the **Jobs and Operations** page and clicked the **Job History** tab. In our example, the snapshot took 3 minutes 52 seconds to complete and the copy backup took 3 minutes 51 seconds as depicted in Figure 2-10.

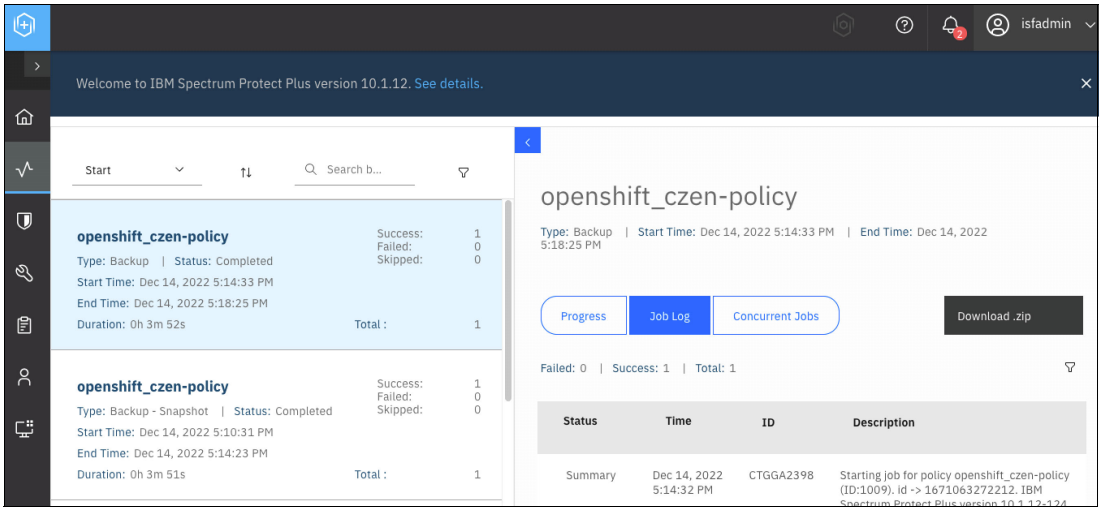

Figure 2-10 Review the progress status of the jobs

So overall, it took under 20 minutes to complete all the phases of the backups for our deployment of **ibm-common-services** and **czen** applications.

We can also see that the SPP job log shows that the execution of recipe hooks and activities from the recipe completed successfully as seen in Figure 2-11.

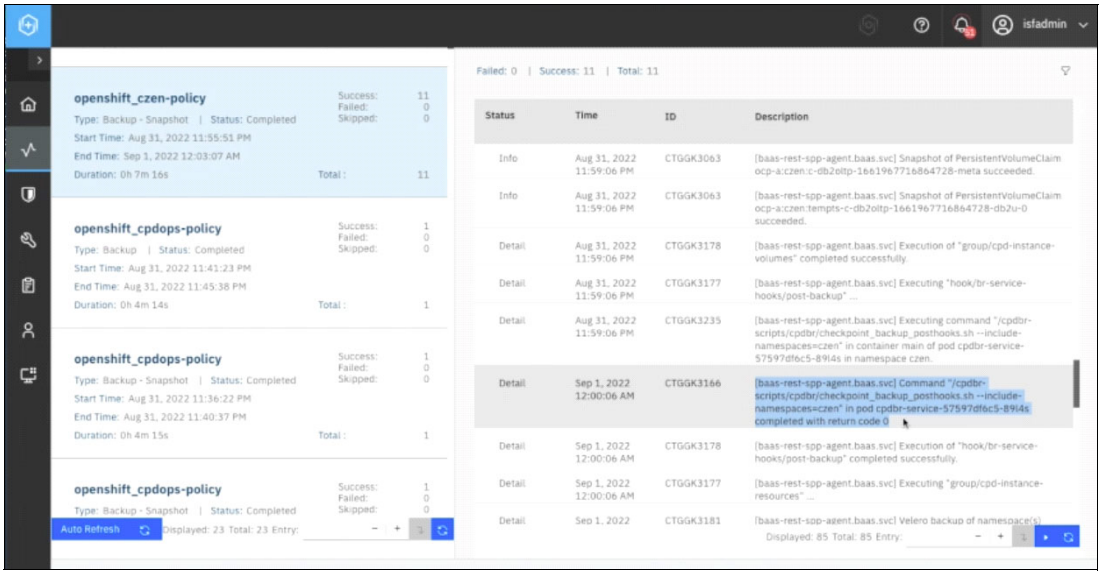

Figure 2-11 SPP job log shows detail of completed phases of the backups

2.3 Backing up the IBM Storage Fusion namespace

From the IBM Storage Fusion UI, we can start the backup of the IBM Storage Fusion namespace. Note: The IBM Storage Fusion application backup is also a recipe based backup, in which only a subset of the Kubernetes resources are backed up, which include the backup storage locations (BSLs), the secrets for the BSLs, the backup policies created in Fusion, and the Fusion applications. To start the backup, we went to **Applications** and selected **ibm-spectrum-fusion-ns**. Then clicked on the Backups tab at the top of the screen. Then, clicked the **Actions** drop-down menu and selected **Backup now** as seen in Figure 2-12.

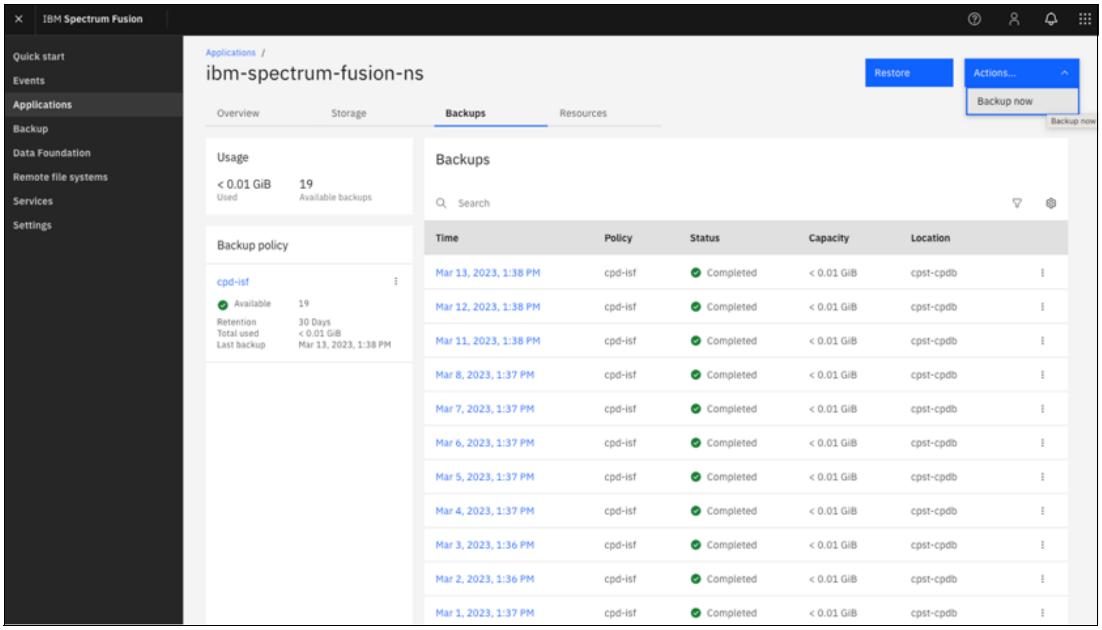

Figure 2-12 Select Backup now for the ibm-spectrum-fusion-ns Fusion application to initiate an On-Demand backup

The next step is to log into the IBM Storage Protect Plus user interface (SPP UI) to watch the progress of the backup job. To get the credentials, you need to get them from the OCP console and go to **Workloads** → **Secrets** → **spp-connection**. That page lists the URL for IBM Storage Protect Plus as well as the username and password.

After logging into the SPP user interface, click on the **Jobs and Operations** icon on the left-hand side, and in the **Running Jobs** tab, we can see the progress of the scheduled backup job as seen in Figure 2-13.

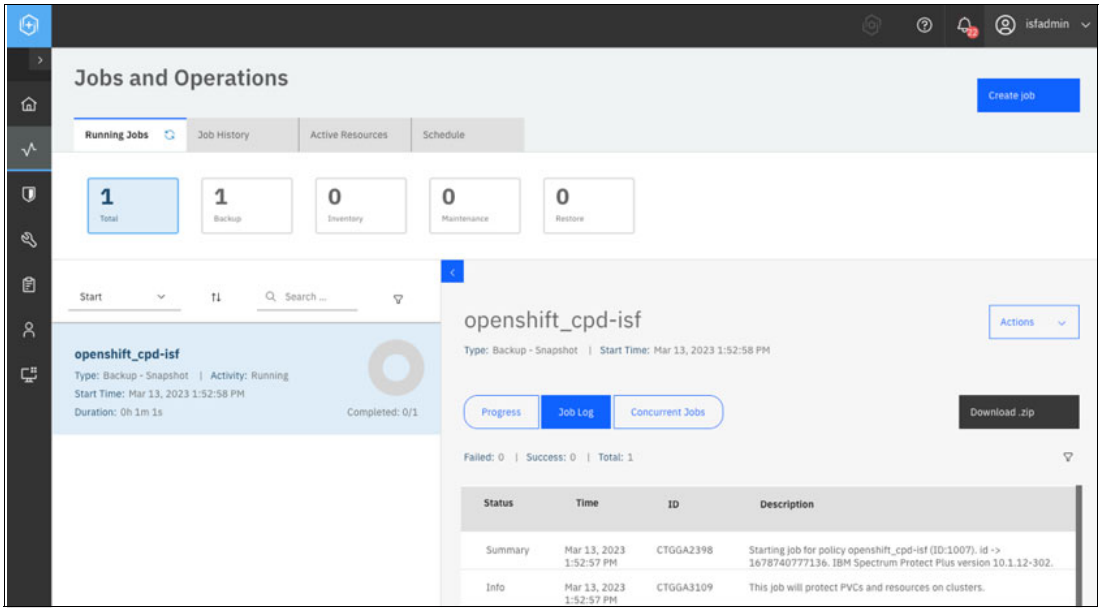

Figure 2-13 Monitor SPP running jobs

After the job completes, we can go to the **Job History** tab and from there, we can see there are two phases to the backup. The first phase is the **Snapshot** which performs the local snapshot of all the resources. In our case, this phase took three minutes and 17 seconds. The second phase is the **Backup** which is a copy of the backup and this phase uploads all information to the selected storage location and this phase took four minutes and 17 seconds. In total, the backup of IBM Storage Fusion resources took seven minutes and 34 seconds to complete as seen in Figure 2-14.

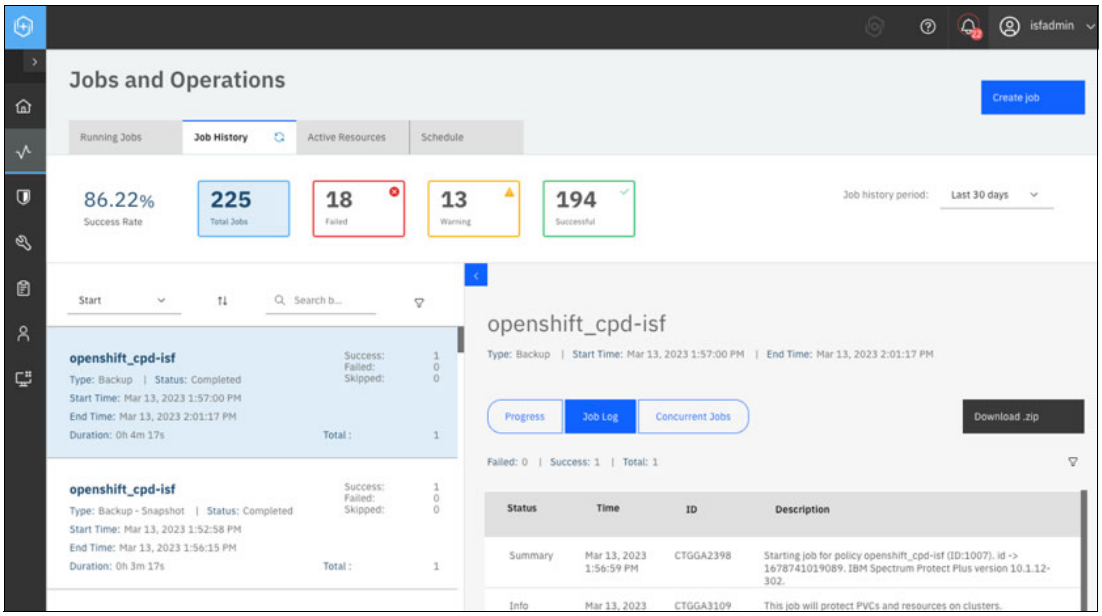

Figure 2-14 The SPP backup jobs for ibm-spectrum-fusion-ns completed successfully

2.4 Backing up the IBM Storage Protect Plus catalog

Now that the Cloud Pak for Data operators, Cloud Pak for Data instance, and IBM Storage Fusion backups are complete, we are ready to back up the actual SPP catalog which is required to restore over to the target cluster. In the SPP UI, go to **Manage Protection** → **IBM Spectrum Protect Plus** → **Backup**. Under SLA Policy Status, click on the **Actions** drop-down menu and select **Start** as seen in Figure 2-15.

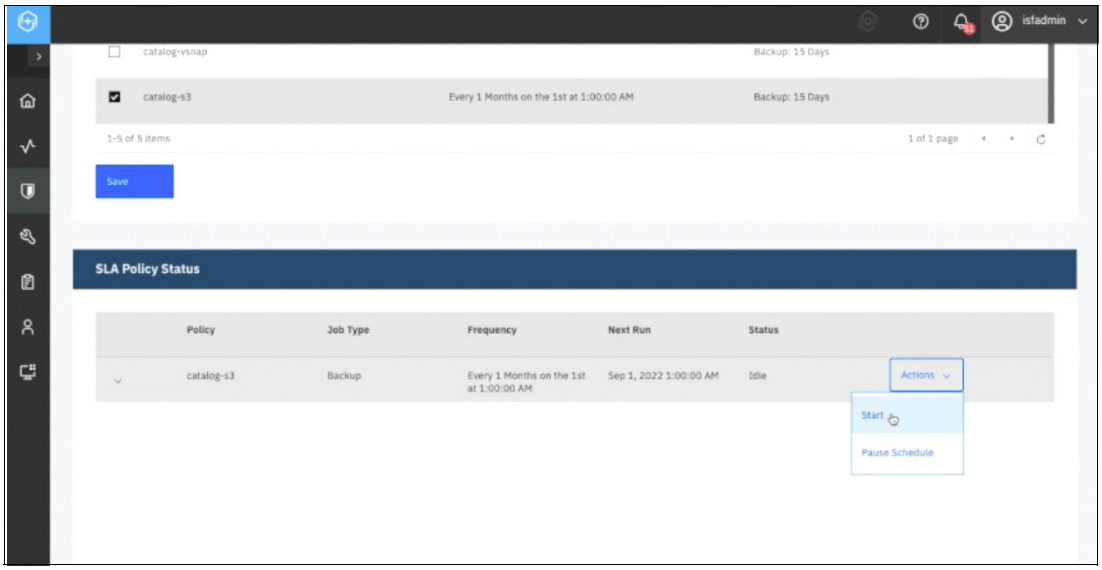

Figure 2-15 Select Start from the Actions drop-down menu to start an SPP catalog backup

Once it starts, we can see the backup is in **Running** State from the **Running Jobs** tab in the **Jobs and Operations** page as seen in Figure 2-16. The SPP catalog backup is meta data only and the backup goes quickly, 1- 5 minutes.

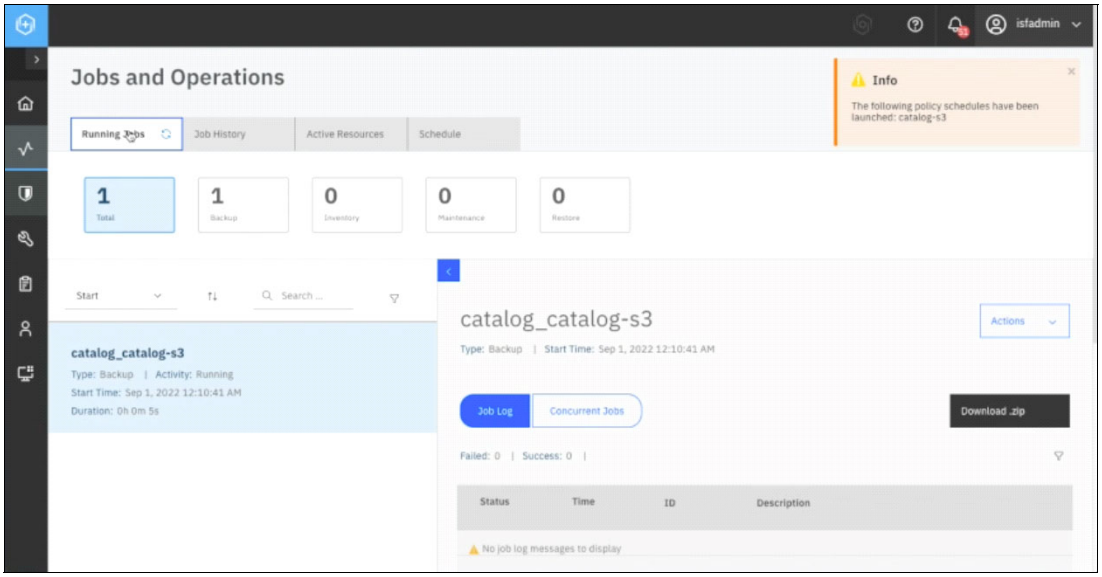

Figure 2-16 Review the progress status of the backup job

3

Restoring to the target cluster

Now that everything on the source cluster is backed up, we can move over to the target cluster to perform the restore. As stated in Chapter 1, “Preparation” on page 1, the source and target clusters need to be at the same OCP major version. This chapter will include restoring the IBM Storage Protect Plus catalog, then restoring the IBM Storage Fusion resources, and then finally restoring the Cloud Pak for Data operators and Cloud Pak for Data instance applications.

3.1 Validating the target cluster is ready for restore

Restore of Cloud Pak for Data requires that the namespaces do not exist prior to restore. On the target cluster, we will check that it does not currently contain **ibm-common-services** and **czen** namespaces. To do this, we issue the **`$ oc get ns |grep -v open`** and the **`$ oc get ns ibm-common-services czen`** commands as seen in Figure 3-1 and the output indicates they do not exist.

```
20220901-00:17:00 [ocp-b:default] ~ $ oc project
Using project "default" on server "https://api.ocp-b.cpst-lab.ibm.com:6443".
20220901-00:17:35 [ocp-b:default] ~ $ oc get ns |grep -v open
NAME                                STATUS   AGE
baas                                Active   10h
csr-auto-approver                   Active   51d
default                             Active   51d
ibm-spectrum-fusion-ns              Active   39h
ibm-spectrum-protect-plus-ns        Active   10h
ibm-spectrum-scale                  Active   37h
ibm-spectrum-scale-csi              Active   39h
ibm-spectrum-scale-operator         Active   39h
kube-node-lease                     Active   51d
kube-public                         Active   51d
kube-system                         Active   51d
rook-ceph                           Active   2d12h
20220901-00:17:46 [ocp-b:default] ~ $ oc get ns ibm-common-services czen
Error from server (NotFound): namespaces "ibm-common-services" not found
Error from server (NotFound): namespaces "czen" not found
20220901-00:17:53 [ocp-b:default] ~ 1 $
```

Figure 3-1 Verify ibm-common-services and czen namespaces do not exist on the target cluster

We also need to verify that none of the Cloud Pak for Data catalog sources exist on the cluster. Issue command **`oc get catsrc -n openshift-marketplace.`** If any Cloud Pak for Data catalogs exist, then it may cause the restore to fail.

We then need to verify this is a fresh installation of IBM Storage Fusion and that it doesn't have any other existing backups. This is because when the SPP catalog is restored from the source cluster, any existing backups will be lost. To check, we need to log into the SPP UI for the target cluster and we need to get the credentials from the target OCP cluster just like we did earlier. To get the credentials, we need to log into the OCP cluster, and from there, go to **Workloads** → **Secrets** → **spp-connection**. This page lists the URL for IBM Storage Protect Plus as well as the username and password. Once we attained our credentials, we logged into the SPP UI and went to the **Jobs and Operations** page. As we can see from the Figure 3-2, we do not have any running jobs.

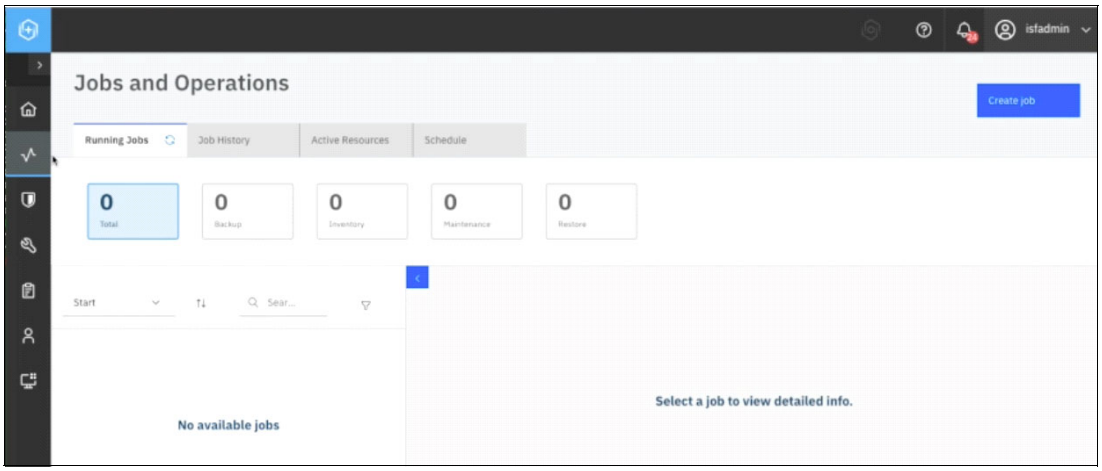

Figure 3-2 Jobs and Operations page of the SSP UI shows no running jobs

From the **Job History** tab, we also checked that there are no previously scheduled backups as seen in Figure 3-3 on page 27.

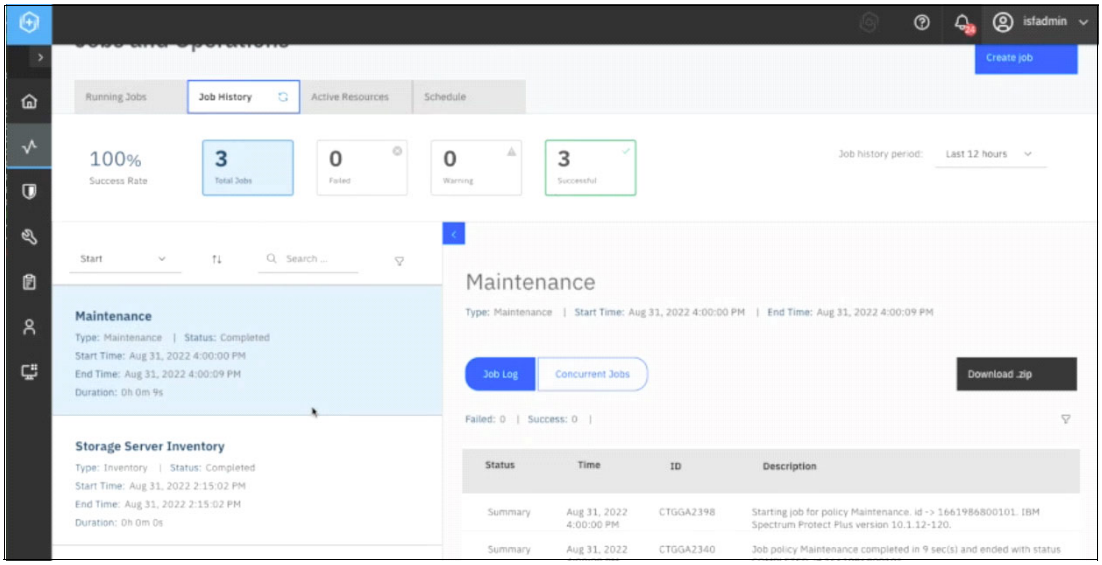

Figure 3-3 No previous scheduled backups

The other thing we need to verify is that we need to have the original storage classes that we had on the source cluster. To check this, we issued **`$ oc get pvc -n czen`** from ocp-a which is our source cluster, as seen in Figure 3-4.

```
20220901-02:22:54 [ocp-a:ibm-spectrum-fusion-ns] ~ $ oc get pvc -n czen
NAME                                    STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS                AGE
activelogs-c-db2oltp-1661967716864728-db2u-0   Bound    pvc-e1962247-71df-4235-8248-b452159bbe1a   100Gi      RWO            ibm-spectrum-scale-rwx      15h
c-db2oltp-1661967716864728-backup       Bound    pvc-56c10e84-bf6f-4aa4-b37e-a4e01adfc540   100Gi      RWX            ibm-spectrum-scale-rwx      15h
c-db2oltp-1661967716864728-meta         Bound    pvc-946e3c96-7c32-4833-ad6f-d4374fa22f4c   100Gi      RWO            ibm-spectrum-scale-rwx      15h
data-c-db2oltp-1661967716864728-db2u-0  Bound    pvc-98aa12ea-5e4d-47c9-97e4-f98a7af51eb3   100Gi      RWO            ibm-spectrum-scale-rwx      15h
data-dsx-influxdb-0                     Bound    pvc-0eacaca4-dd7f-4e94-bbcb-1d6a6c32c3f5   10Gi       RWO            ibm-spectrum-scale-rwo      26h
datadir-zen-metastoredb-0               Bound    pvc-6f45adff-e119-4986-8df4-530411c1eb69   10Gi       RWO            ibm-spectrum-scale-rwo      27h
datadir-zen-metastoredb-1               Bound    pvc-fd964b83-98df-42a0-9007-18e818be9e28   10Gi       RWO            ibm-spectrum-scale-rwo      27h
datadir-zen-metastoredb-2               Bound    pvc-3bb31624-4f7d-4641-b445-9f88cc9b010f   10Gi       RWO            ibm-spectrum-scale-rwo      27h
tempts-c-db2oltp-1661967716864728-db2u-0   Bound    pvc-ff3daf62-dee3-4ef0-96a0-e53931ad0d77   100Gi      RWO            ibm-spectrum-scale-rwx      15h
user-home-pvc                           Bound    pvc-4d64456a-bb42-498b-ac6b-d358ed319a4e   10Gi       RWX            ibm-spectrum-scale-rwx      27h
20220901-02:23:05 [ocp-a:ibm-spectrum-fusion-ns] ~ $
```

Figure 3-4 Verify the original storage classes on the source cluster

From the **STORAGECLASS** column from Figure 3-4, we can see that we have the IBM Storage Scale storage classes for both RWX and RWO: ibm-spectrum-scale-rwo and ibm-spectrum-scale-rwx. Furthermore, we can also see that we have the same storage classes on ocp-b, which is our target cluster by issuing **`$ oc get sc`** as shown in Figure 3-5. If these two storage classes do not exist on both the source and target clusters, the restore will fail.

```
20220901-02:21:51 [ocp-b:default] ~ $ oc get sc
NAME                        PROVISIONER                  RECLAIMPOLICY   VOLUMEBINDINGMODE      ALLOWVOLUMEEXPANSION   AGE
ibm-spectrum-scale-internal   kubernetes.io/no-provisioner   Delete          WaitForFirstConsumer   false                  39h
ibm-spectrum-scale-rwo      spectrumscale.csi.ibm.com    Delete          Immediate              true                   27h
ibm-spectrum-scale-rwx      spectrumscale.csi.ibm.com    Delete          Immediate              true                   27h
ibm-spectrum-scale-sample   spectrumscale.csi.ibm.com    Delete          Immediate              false                  38h
ibm-spectrum-scale-sc       spectrumscale.csi.ibm.com    Delete          Immediate              true                   29h
rook-ceph-block             rook-ceph.rbd.csi.ceph.com   Delete          Immediate              true                   2d14h
20220901-02:22:46 [ocp-b:default] ~ $
```

Figure 3-5 Verify target cluster storage classes are the same as the source cluster

Additionally, the target cluster must be prepared to install Cloud Pak for Data. Change any node settings that are required to match the source cluster. Reference the following Cloud Pak for Data link:
https://www.ibm.com/docs/en/cloud-paks/cp-data/4.5.x?topic=cluster-changing-required-node-settings

3.2 Restore the IBM Storage Protect Plus catalog

Once we verified that we have a fresh installation of Storage Fusion and that there are no existing backups on our target cluster, we are ready to restore the SPP catalog from the IBM Cloud Object Storage.

The first thing that must be setup is access to the same s3 object storage where the SPP catalog was backed up on the source cluster. In the SPP UI, go to the **System Configuration** → **Storage** page as shown in Figure 3-6.

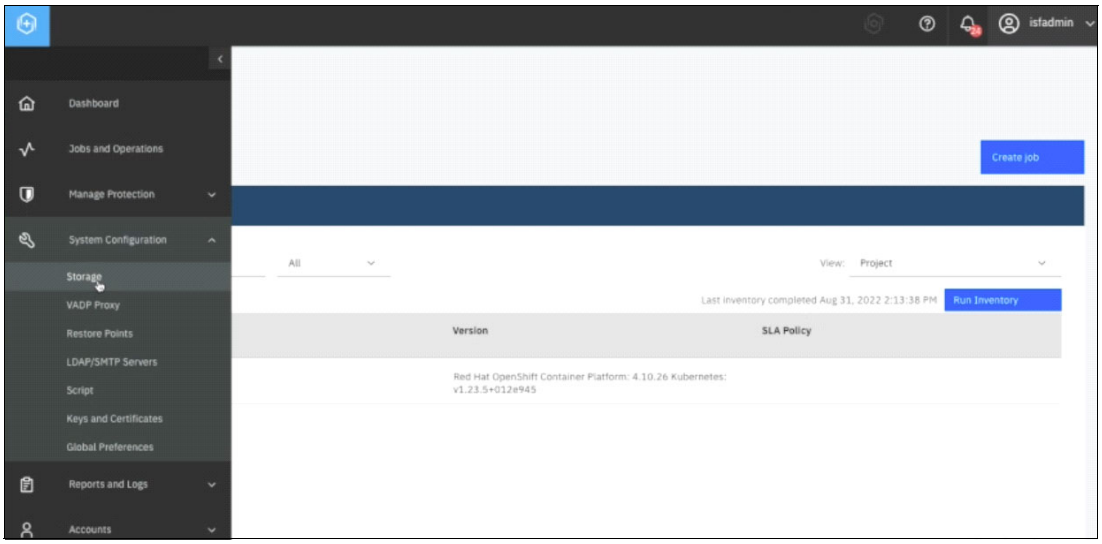

Figure 3-6 Select Storage under System Configuration

The next step is to select **Cloud storage** → **Add cloud storage** and go through the menus as shown in Figure 3-7.

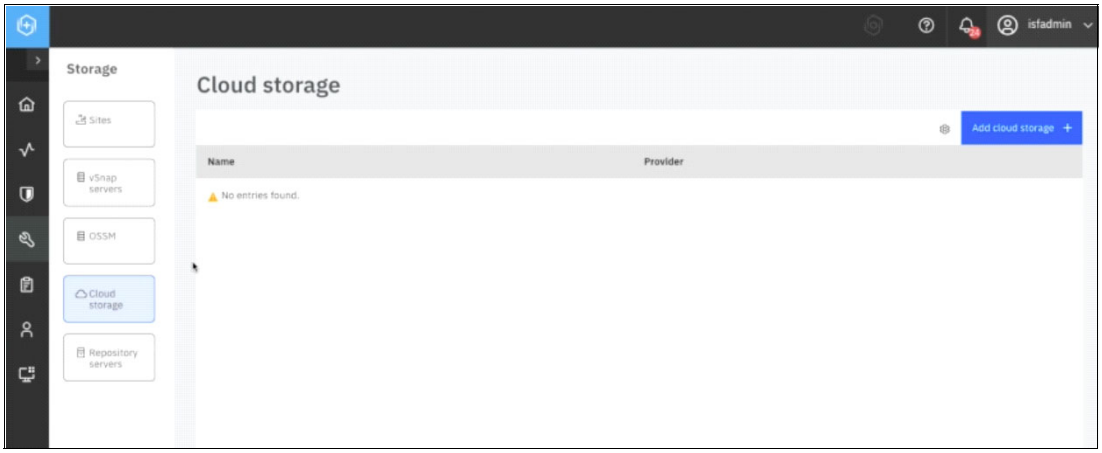

Figure 3-7 Select Add cloud storage

In our example, we used IBM Cloud Object storage as the backup storage location. In the first page, we select **IBM Cloud Object Storage** and in the next page, we need to enter the details for location. This will be the same location we performed the backup to on the source cluster so that the target cluster will be able to look up that catalog backup and be able to restore from it. See Figure 3-8 on page 29.

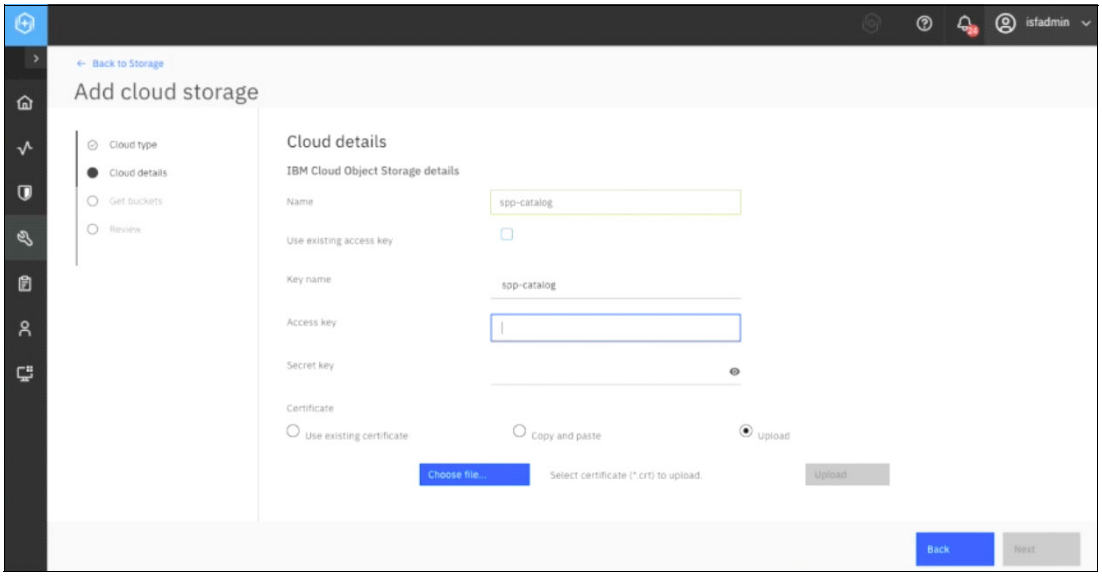

Figure 3-8 Enter the Cloud details

Once we entered the **Cloud details**, the **next** page, prompts us to enter the details of the bucket endpoint. In our example, it was IBM Cloud Object Storage, but it could be set to the location of your choice as seen in Figure 3-9.

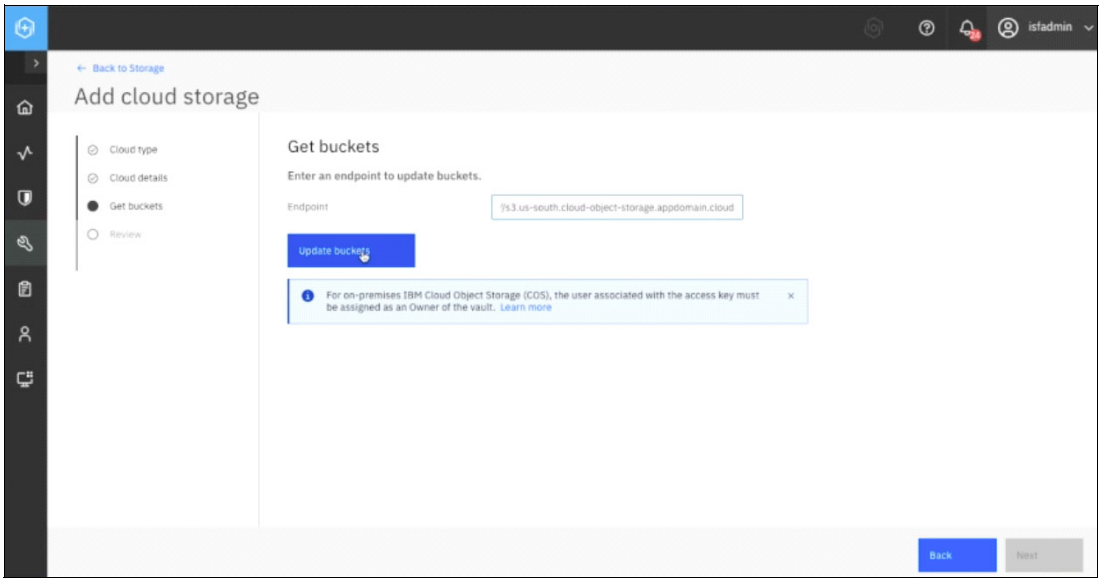

Figure 3-9 Enter the details of the bucket endpoint

In the last page, we reviewed the details we entered, and once we confirmed it was correct, we clicked **Submit** as shown in Figure 3-10.

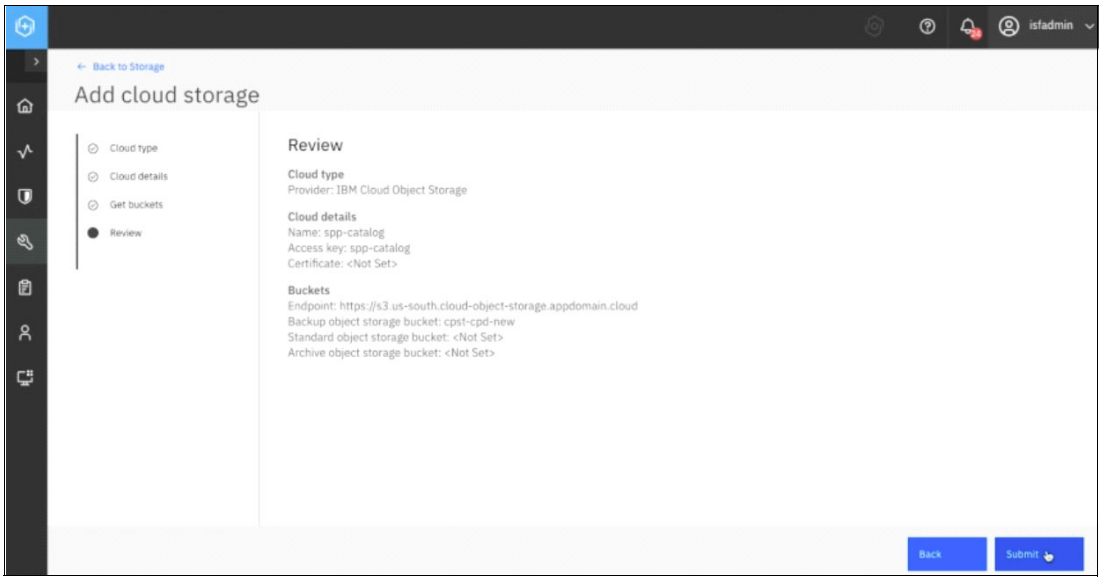

Figure 3-10 Review and submit to add cloud storage

Now we are ready to restore the SPP catalog. Go to the **Manage Protection** → **IBM Spectrum Protect Plus** → **Restore** Page as seen in Figure 3-11.

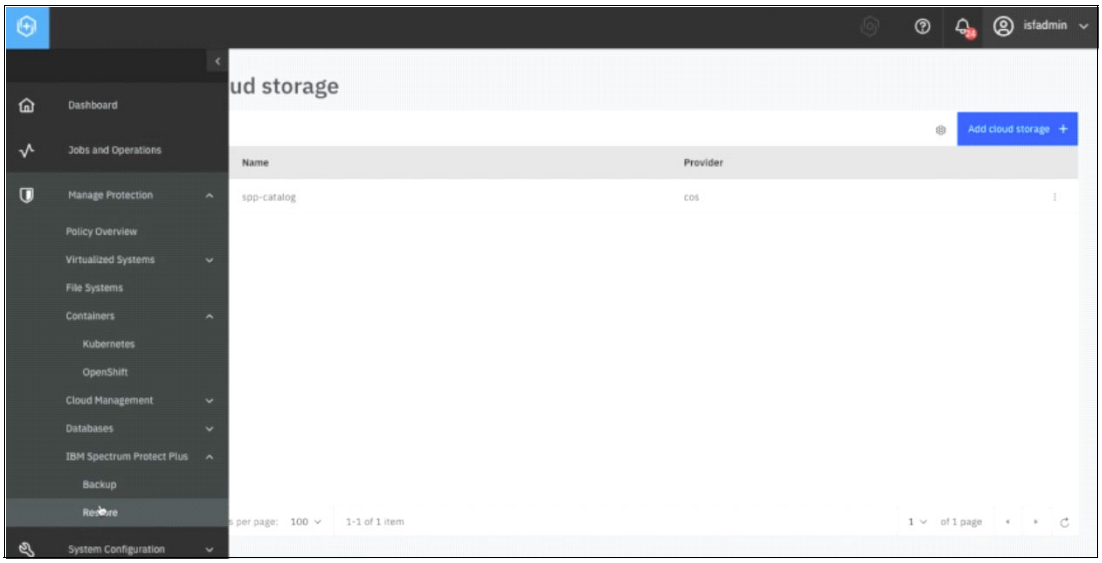

Figure 3-11 Select Restore

In this example, we performed the restore from IBM Cloud Object Storage, so we go to the **From cloud storage** tab at the top and click on the storage location **spp-catalog** as seen in Figure 3-12 on page 31.

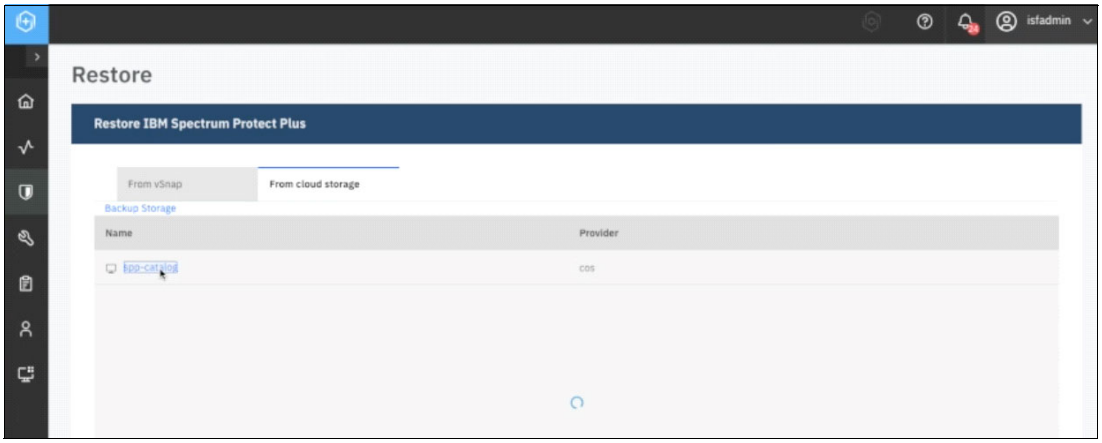

Figure 3-12 Restore SPP catalog from cloud storage

Then we found the most recent backup we just performed and clicked on **Restore** as shown in Figure 3-13.

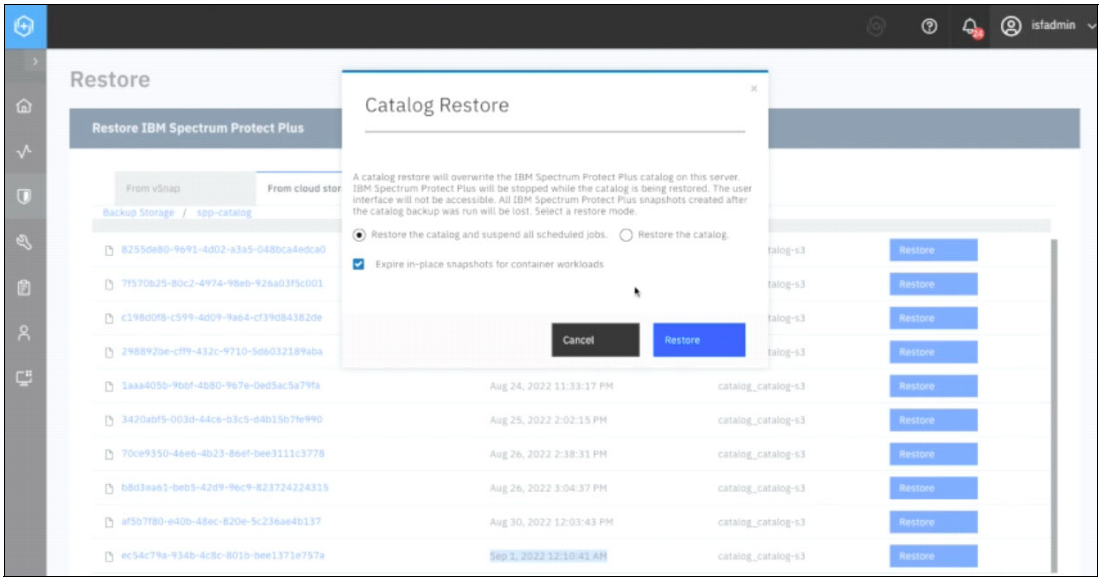

Figure 3-13 Restore the most recent spp-catalog

We then left the default values to **Restore the catalog and suspend all scheduled jobs** and we left the **Expire in-place snapshots for container workloads** checkbox selected and clicked on **Restore**.

A warning box asking if you are sure you want to proceed then appeared and by clicking **Yes**, the restore for SPP is started. See Figure 3-14.

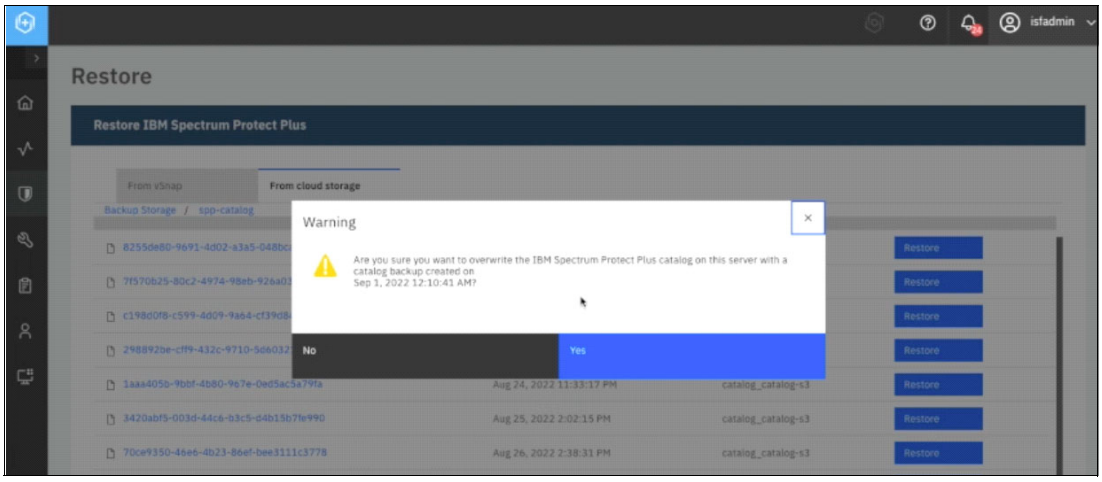

Figure 3-14 Select Yes on message box to start the restore for SPP

The restore for SPP took about 20 minutes in our example and we followed the progress via command line within the SPP namespace by following the **sppvirgo** pod logs. To find the name of the **sppvirgo** pod, we issued the following command (see Figure 3-15).

```
$ oc get pod -n ibm-spectrum-protect-plus-ns
```

```
[ocp-b:default] ~ $ oc get pod -n ibm-spectrum-protect-plus-ns
NAME                                READY   STATUS    RESTARTS      AGE
spp-awsebs-5db99995f-s5jx8          1/1     Running   0             10h
spp-awsec2-7d66b5b55-gsqfd          1/1     Running   0             10h
spp-ingressproxy-54d7dcc67c-vzmks   1/1     Running   0             10h
spp-manager-77b7b9f7c9-slh2b        1/1     Running   0             10h
spp-operator-7fd9f8c748-mvj2g       2/2     Running   0             10h
spp-plugins-mongo-66f478dbd7-7svw4  1/1     Running   0             10h
spp-plugins-redis-6cbb55b57-67jv8   1/1     Running   0             10h
spp-proxy-755d6bf5c9-2hcws          1/1     Running   0             10h
sppdbmongo-5c4bc8bcd5-x72d4         1/1     Running   0             10h
sppdbmongo2-75ff87d96f-zfkv5        1/1     Running   0             10h
sppdbpostgres-7546d9bfc9-lg47b      1/1     Running   0             10h
sppkc-6b457df57c-k57kk              1/1     Running   0             10h
sppnodejs-744fd65b4f-pmddd          1/1     Running   9 (10h ago)   10h
sppui-7bc6dd7898-blx5n              1/1     Running   0             10h
sppvadp-59db565b9-qmwkm             1/1     Running   0             10h
sppvirgo-58f764f974-5sr9q           1/1     Running   0             10h
```

Figure 3-15 Find the sppvirgo pod name

The command to follow the **sppvirgo** pod logs is as follows:

```
$ oc logs <sppvirgo-pod-name> -n ibm-spectrum-protect-plus-ns -f --since=1m
```

In our example the name of our **sppvirgo** pod is **sppvirgo-58f764f974-5sr9q** so we issued the following command as shown in Figure 3-16 on page 33:

```
$ oc logs sppvirgo-58f764f974-5sr9q -n ibm-spectrum-protect-plus-ns -f --since=1m
```

```
20220901-00:20:28 [ocp-b:default] ~ $ oc logs sppvirgo-58f764f974-5sr9q -n ibm-spectrum-protect-plus-ns -f --since=1m
[2022-09-01T07:20:15.729Z] System Bundle Shutdown       <KE0010I> Shutdown initiated.
[2022-09-01T07:20:22.614Z] System Bundle Shutdown       <TC0002I> Stopping Tomcat.
[2022-09-01T07:20:24.515Z] System Bundle Shutdown       <TC0003I> Stopped Tomcat.
[2022-09-01T07:20:26.325Z] INFO  iLogServiceListener@42b4c7d6 org.osgi.service.log.LogService                    Bundle org.eclipse.virgo.medic.core_3.7.2.RELEASE, Service 94
, ServiceEvent UNREGISTERING
[2022-09-01T07:20:26.325Z] INFO  iLogServiceListener@42b4c7d6 org.osgi.service.log.LogService                    Bundle org.eclipse.virgo.medic.core_3.7.2.RELEASE, Service 34
, ServiceEvent UNREGISTERING
[2022-09-01T07:20:26.325Z] INFO  iLogServiceListener@42b4c7d6 org.osgi.service.log.LogService                    Bundle org.eclipse.virgo.medic.core_3.7.2.RELEASE, Service 32
, ServiceEvent UNREGISTERING
[2022-09-01T07:20:26.325Z] INFO  iLogServiceListener@42b4c7d6 org.osgi.service.log.LogService                    Bundle org.eclipse.virgo.medic.core_3.7.2.RELEASE, Service 33
, ServiceEvent UNREGISTERING
[2022-09-01T07:20:26.325Z] INFO  iLogServiceListener@42b4c7d6 org.osgi.service.log.LogService                    Bundle org.eclipse.virgo.medic.core_3.7.2.RELEASE, Service 36
, ServiceEvent UNREGISTERING
[2022-09-01T07:20:26.325Z] INFO  iLogServiceListener@42b4c7d6 org.osgi.service.log.LogService                    Bundle org.eclipse.virgo.medic.core_3.7.2.RELEASE, Service 96
, ServiceEvent UNREGISTERING
[2022-09-01T07:20:26.326Z] INFO  iLogServiceListener@42b4c7d6 org.osgi.service.log.LogService                    Bundle org.eclipse.virgo.medic.core_3.7.2.RELEASE, Service 37
, ServiceEvent UNREGISTERING
[2022-09-01T07:20:26.326Z] INFO  iLogServiceListener@42b4c7d6 org.osgi.service.log.LogService                    Bundle org.eclipse.virgo.medic.core_3.7.2.RELEASE, Service 40
, ServiceEvent UNREGISTERING
[2022-09-01T07:20:26.326Z] INFO  iLogServiceListener@42b4c7d6 org.osgi.service.log.LogService                    Bundle org.eclipse.virgo.medic.core_3.7.2.RELEASE, Service 39
, ServiceEvent UNREGISTERING
[2022-09-01T07:20:26.326Z] INFO  iLogServiceListener@42b4c7d6 org.osgi.service.log.LogService                    Bundle org.eclipse.virgo.medic.core_3.7.2.RELEASE, Service 41
, ServiceEvent UNREGISTERING
```

Figure 3-16 Follow the sppvirgo-58f764f974-5sr9q pod logs to check restore progress

From the logs we can see it unregistered everything from the original instance and it is restoring the instance from the source cluster. Once that process completes, we will see the SPP pods, including the sppnodejs pod, in **Running** state as seen in Figure 3-17.

```
20220901-00:41:02 [ocp-b:default] ~ 1 $ oc get pod -n ibm-spectrum-protect-plus-ns
NAME                                READY   STATUS    RESTARTS      AGE
spp-awsebs-5db99995f-s5jx8          1/1     Running   0             10h
spp-awsec2-7d66b5b55-gsqfd          1/1     Running   0             10h
spp-ingressproxy-54d7dcc67c-vzmks   1/1     Running   0             10h
spp-manager-77b7b9f7c9-slhzb        1/1     Running   0             10h
spp-operator-7fd9f8c748-mvjzg       2/2     Running   0             10h
spp-plugins-mongo-66f478dbd7-7svw4  1/1     Running   0             10h
spp-plugins-redis-6cbb55b57-67jv8   1/1     Running   0             10h
spp-proxy-755d6bf5c9-2hcws          1/1     Running   0             10h
sppdbmongo-5c4bc8bcd5-x72d4         1/1     Running   0             10h
sppdbmongo2-75ff87d96f-zfkv5        1/1     Running   0             10h
sppdbpostgres-7546d9bfc9-lg47b      1/1     Running   0             10h
sppkc-6b457df57c-k57kk              1/1     Running   0             10h
sppnodejs-744fd65b4f-l8lrj          1/1     Running   2 (83s ago)   4m38s
sppui-7bc6dd7898-blx5n              1/1     Running   0             10h
sppvadp-59db565b9-qmwkm             1/1     Running   0             10h
sppvirgo-58f764f974-5sr9q           1/1     Running   0             10h
20220901-00:41:06 [ocp-b:default] ~ $
```

Figure 3-17 SPP pods show running

The next step is to log back into the SPP UI from ocp-b, which is our target cluster, but this time, our old credentials will not work because the credentials changed after the restore completed and now the credentials are from the original source cluster. Therefore, we need to get the credentials from the source cluster by following the same procedure as before and logging into the OpenShift cluster and from there, go to **Workloads** → **Secrets** → **spp-connection**. Once we have the credentials, we can log into the SPP UI now that the restore is completed.

Once we log in, the first step is to go into the **Accounts** → **User** page and change the **isfadmin** user that IBM Storage Fusion uses and change the password back to what IBM Storage Fusion on the target cluster has recorded for it.

To do this, from the **Users** page, we click on the three dots within the **isfadmin** box and this will bring up the **Change password** window as shown in Figure 3-18.

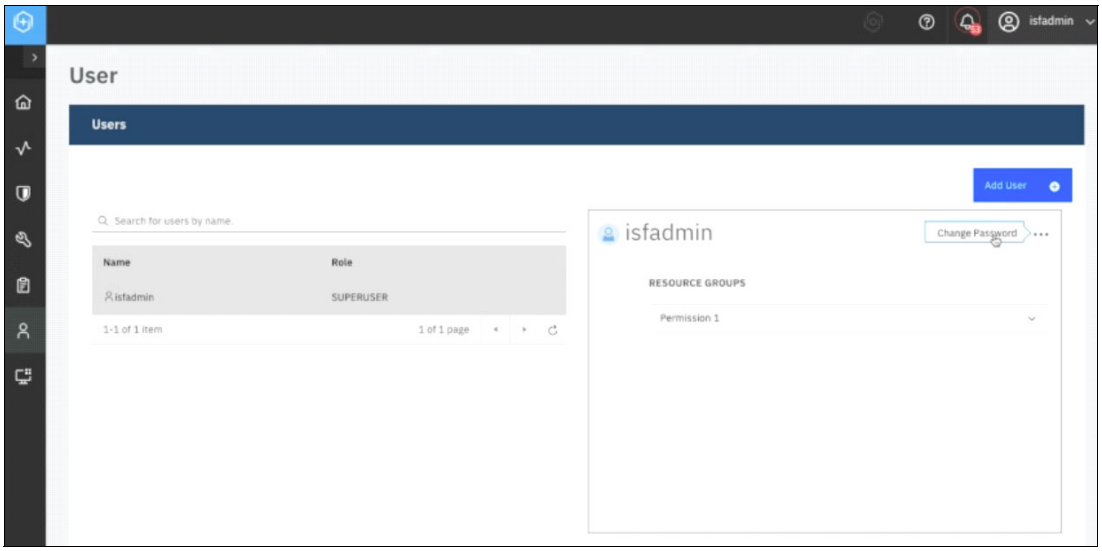

Figure 3-18 Change the isfadmin user password to what was recorded on the target cluster

Once we enter the password, we click on **Update user** and then a window appears saying we must log in again because the password changed as shown in Figure 3-19. Therefore, we log back in with the credentials from ocp-b, which is our target cluster.

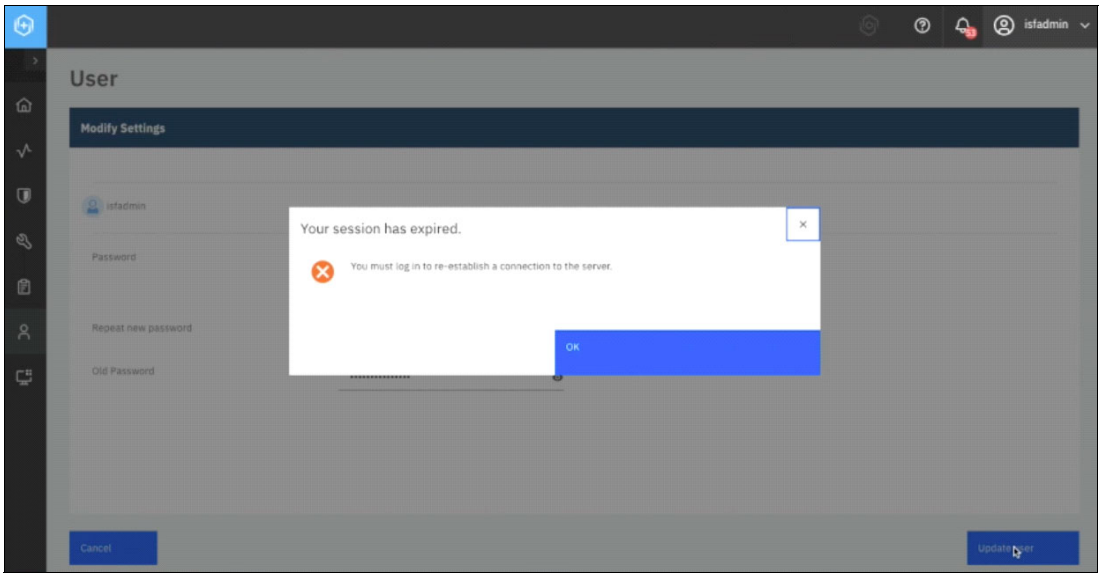

Figure 3-19 Log in again after the password was updated

Once we log into the SPP UI for our target cluster (ocp-b), we go to the **Jobs and Operations** → **Job History** tab and we can see that the restore contains all the previous backups that were performed on the source cluster and now they are available on the target cluster as shown in Figure 3-20 on page 35.

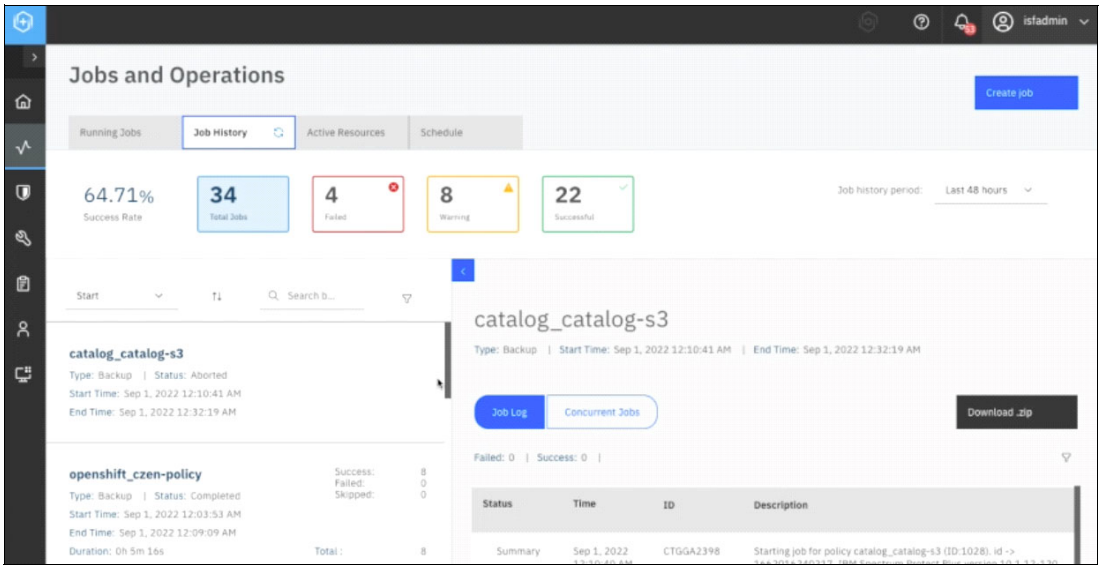

Figure 3-20 Verify all the backups from the source are restored on the target cluster

Then we go to **System Configuration** → **Storage** → **Cloud storage** to confirm the storage locations are also restored from the source cluster as shown in Figure 3-21.

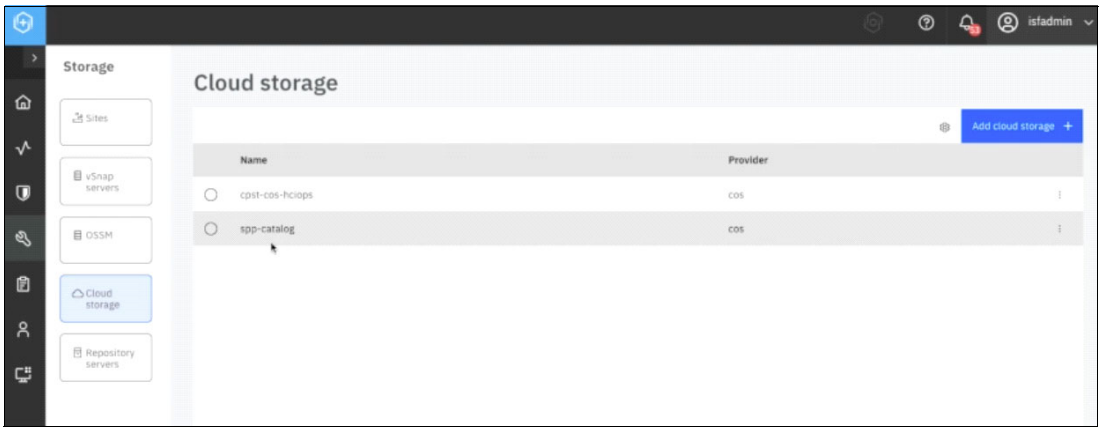

Figure 3-21 Confirm the storage locations are also restored from the source cluster

The restore brings over the cluster registration from the source cluster. So first, we need to delete the managed cluster registration from the OpenShift containers section and delete that registration. To do this, from within the SPP UI for our target cluster, we go to the **Manage Protection** → **Containers** → **OpenShift** page and click on **Manage clusters** as shown in Figure 3-22.

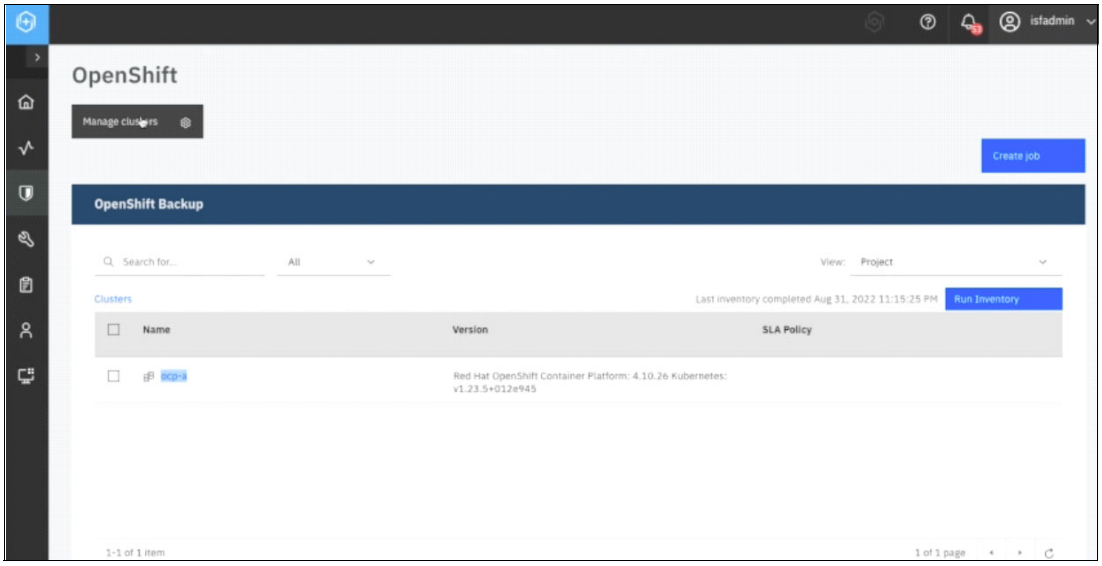

Figure 3-22 Go to Manage clusters

Then we clicked on the **trash bin** icon to delete the prior host address, entered the code to confirm deletion and then clicked on **UNREGISTER** as shown in Figure 3-23.

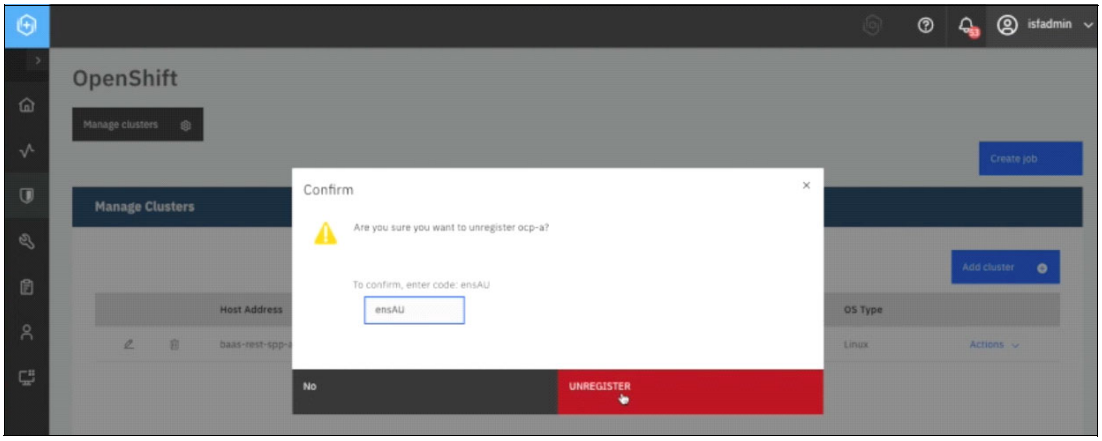

Figure 3-23 Unregister ocp-a

After successful deletion, a message stating **The provider was successfully unregistered** appeared, and we clicked **Ok** as seen in Figure 3-24 on page 37.

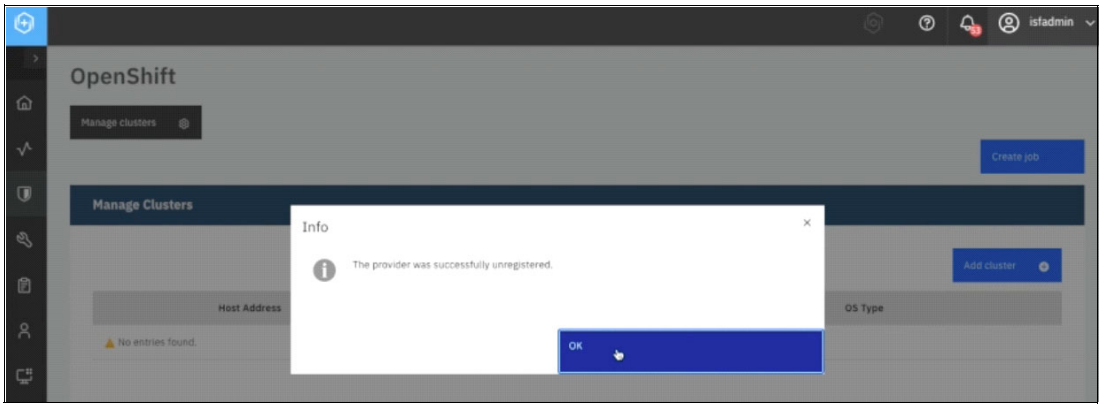

Figure 3-24 Successfully unregistered message

The next step is to delete one of the baas transaction manager pods and by doing this, it will cause an automatic re-registration of the local cluster. This will register the baas agent locally and it will go back to being the target cluster (ocp-b), instead of the source cluster (ocp-a). This step is very important to ensure everything is set up and works correctly.

To do this, we first list the pods in the **baas** namespace as shown in Figure 3-25 by issuing `$ oc get pod -n baas and then issued $ oc delete pod -n baas-transaction-manager-5bf458648-6xk2z`

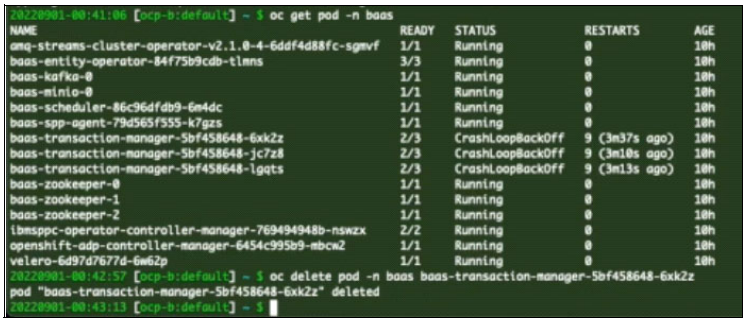

Figure 3-25 Delete one of the baas transaction manager pods to cause local cluster re-registration

Then, we confirmed the **baas transaction manager** pods are recreated and are in **Running** state as seen in Figure 3-26.

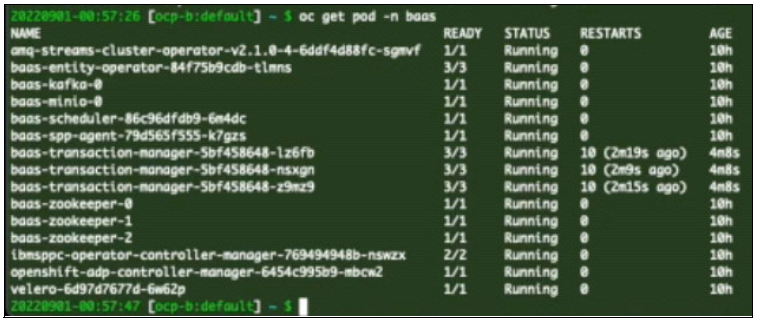

Figure 3-26 baas transaction manager pods are recreated and running

The next step is to go back to the SPP UI from ocp-b, which is our target cluster, and go to the **Manage Protection** → **Containers** → **OpenShift** page to confirm the creation of the new registration and to perform a test and inventory on it to ensure everything looks as expected.

To perform the test, we clicked on the **Actions** drop-down arrow and clicked on **Test** as shown in Figure 3-27.

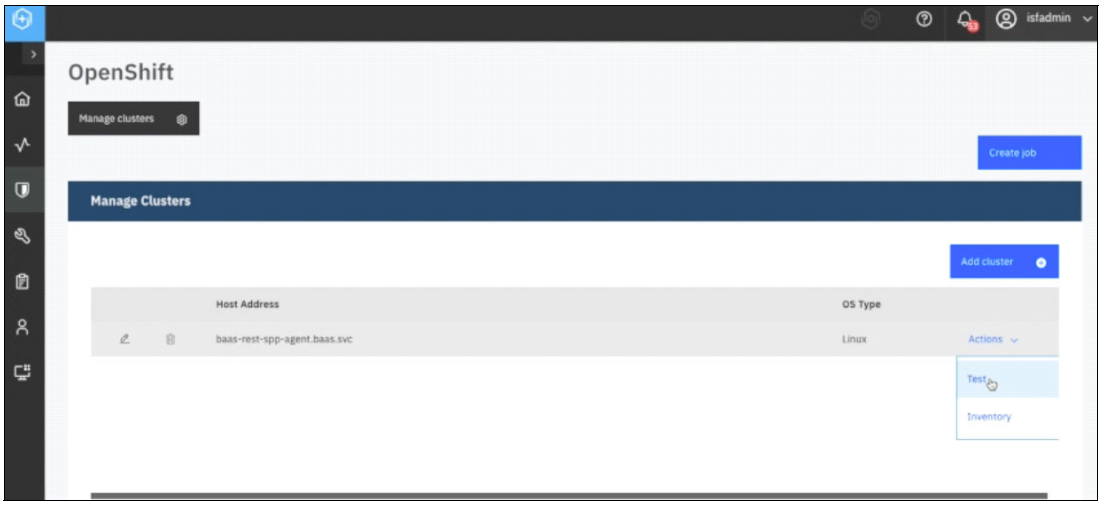

Figure 3-27 Perform a test

Once the test completes, the results of the test are displayed as seen in Figure 3-28.

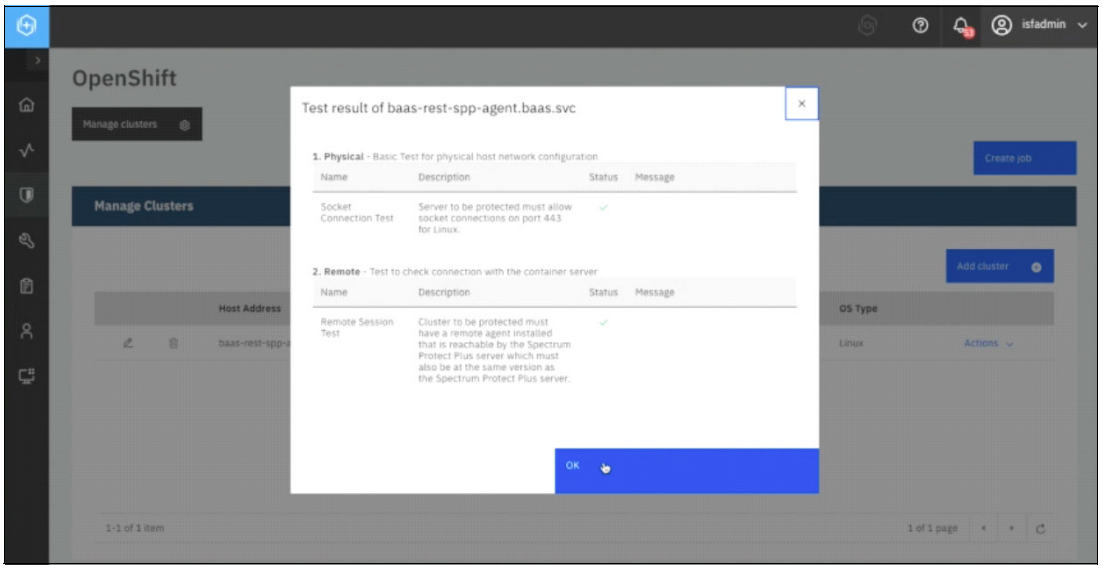

Figure 3-28 Results of the test are displayed in the window

The next step is to run the inventory by clicking on the **Actions** drop-down arrow and selecting **Inventory** as seen in Figure 3-29 on page 39.

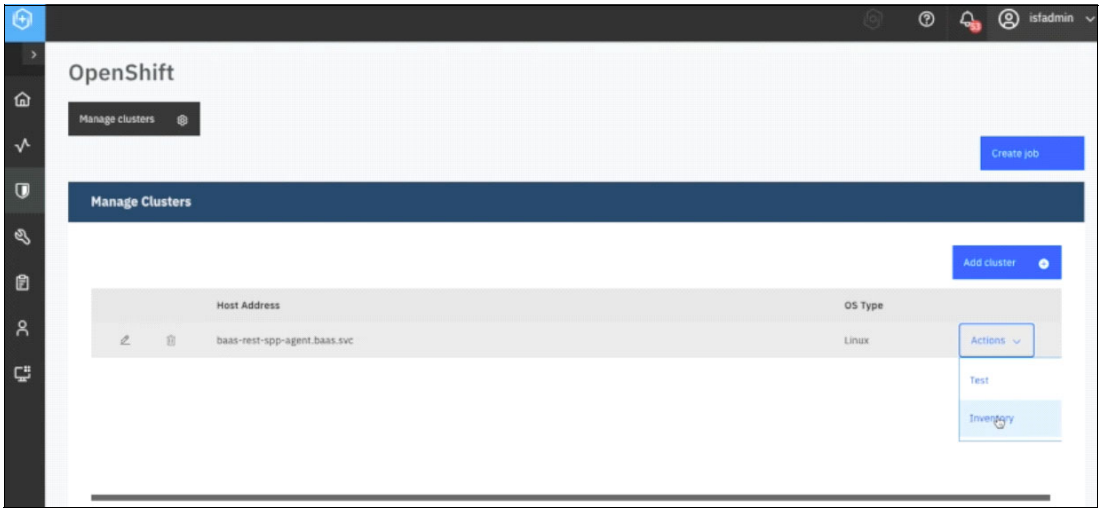

Figure 3-29 Select Inventory from the actions drop-down menu

After the inventory job is created, we go to the **Jobs and Operations** → **Job History** page and we confirmed that the Inventory completed as seen in Figure 3-30.

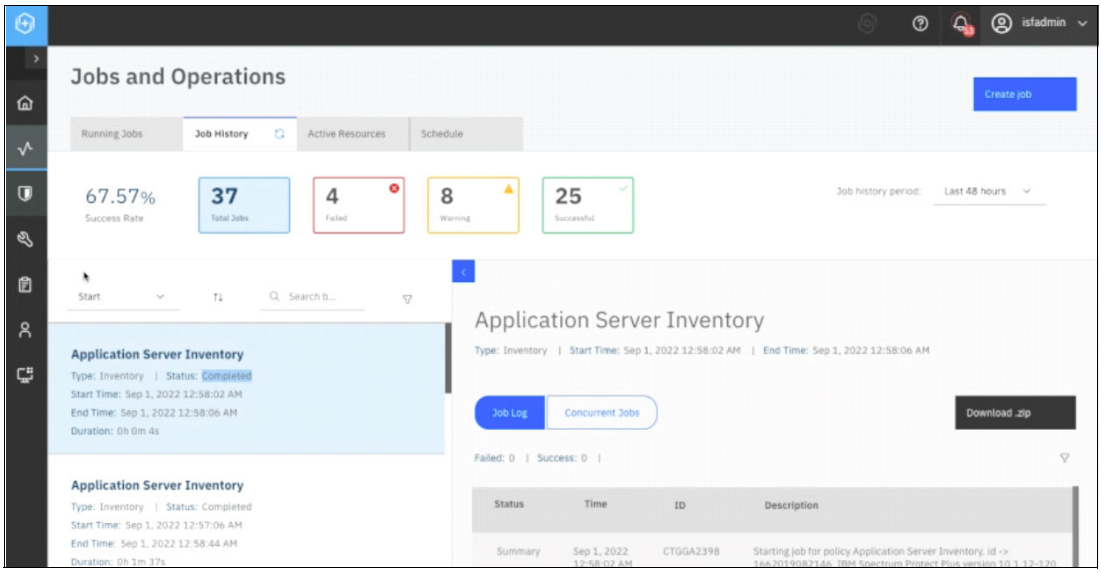

Figure 3-30 Confirm the inventory completed

The next step is to go to the **Manage Protection** → **Containers** → **OpenShift** page and we confirmed that the name of the cluster is no longer **ocp-a**; it is **opc-b** now, which is our target cluster as show in Figure 3-31.

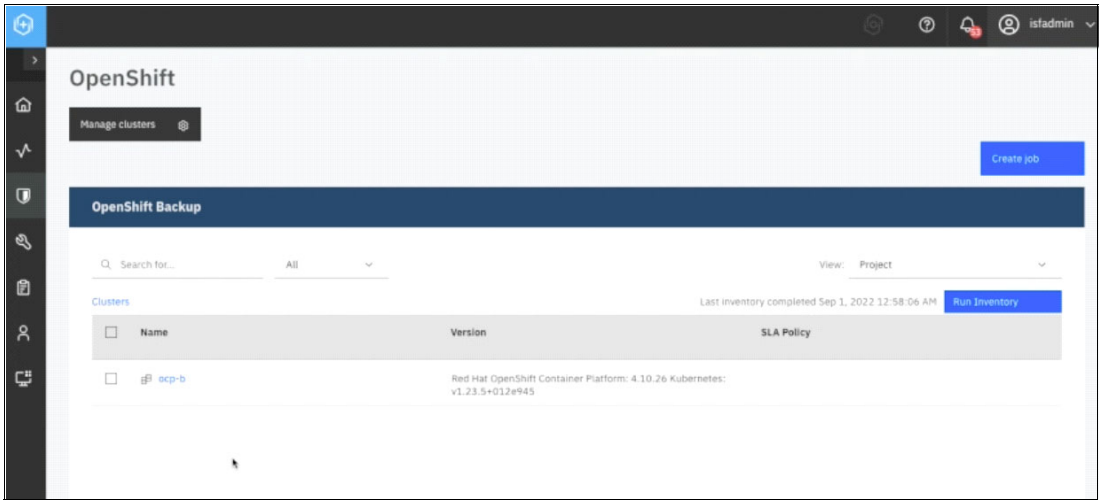

Figure 3-31 Confirm the cluster is opc-b our target cluster

3.3 Restore the IBM Storage Fusion application

To restore the IBM Storage Fusion application resources, from the IBM Storage Protect Plus UI, go to **Jobs and Operations** and click **Create job**. Then select **Restore** as seen in Figure 3-32.

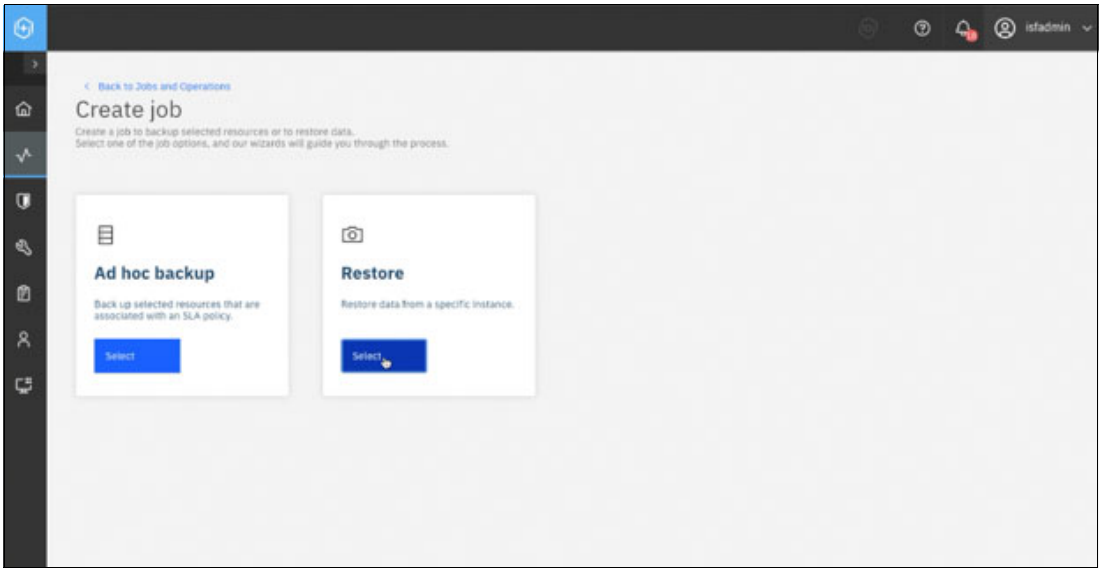

Figure 3-32 In SPP, create a restore job for the ibm-spectrum-fusion-ns application

In the restore job view, select **Containers** → **OpenShift** and click **Next**, as seen in Figure 3-33 on page 41.

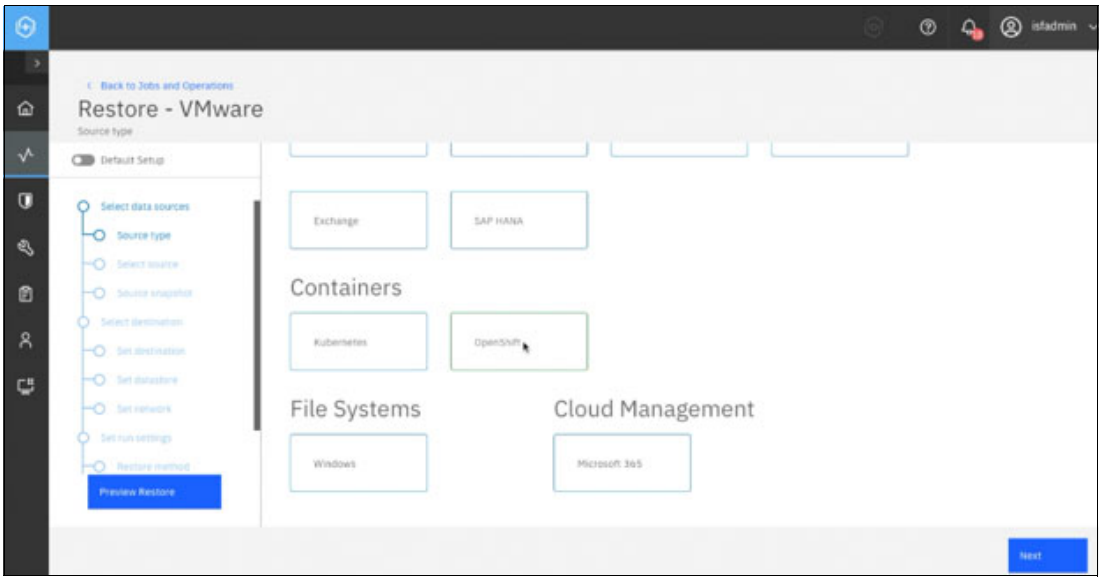

Figure 3-33 Restore job will be OpenShift restore type

For the restore source, select the source cluster and click the **plus** icon next to **ibm-spectrum-fusion-ns**, as seen in Figure 3-34.

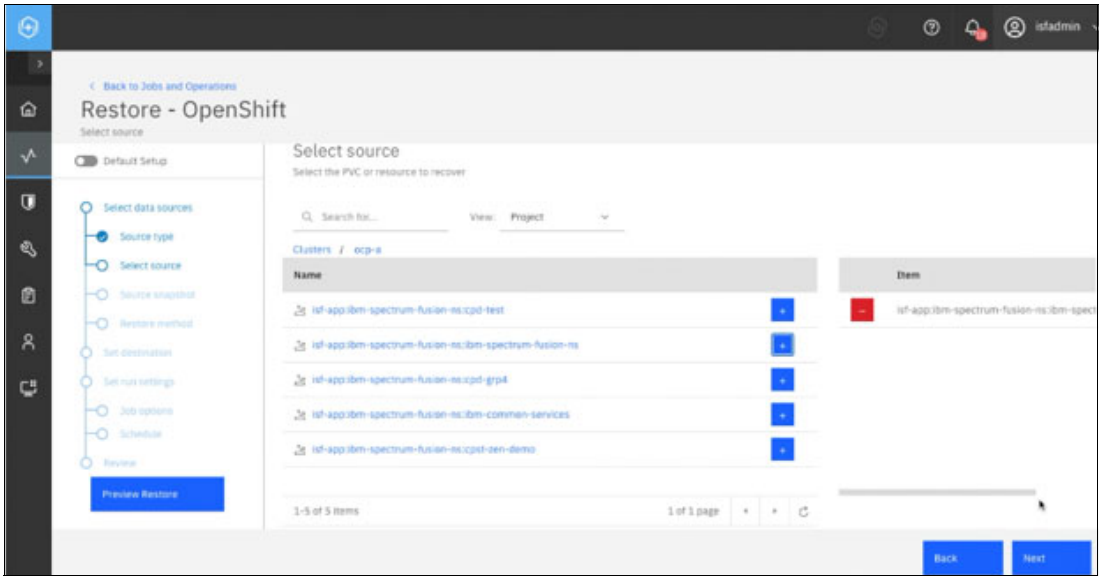

Figure 3-34 The restore resource will be the isf application for ibm-spectrum-fusion-ns

For the source snapshot, select **From Copy** and **On-Demand**. Then select the desired restore point, in this case the most recent backup. See Figure 3-35.

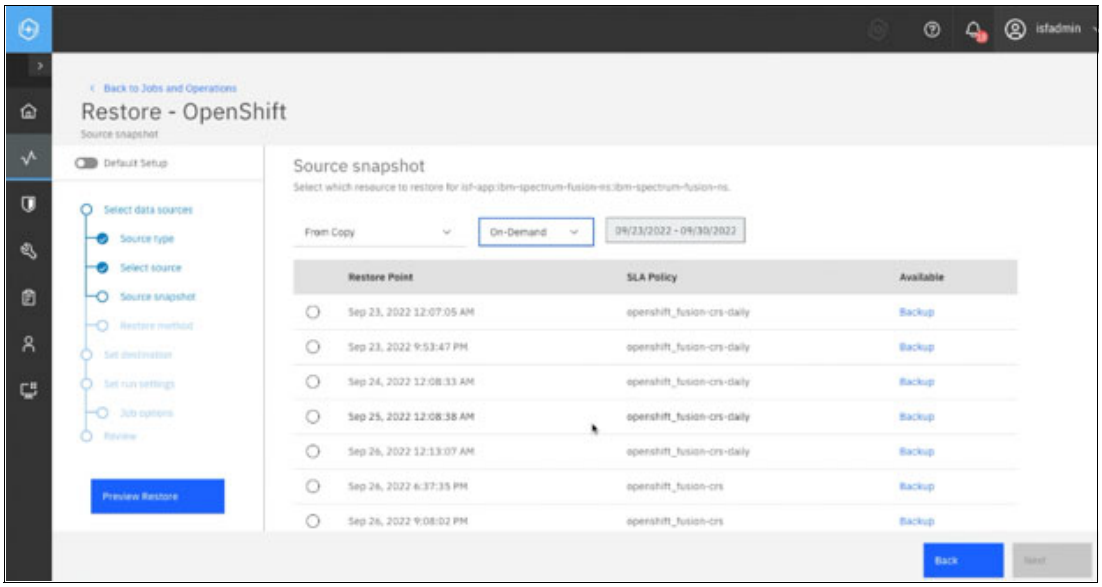

Figure 3-35 Restore type of From Copy and On-Demand

For destination, select **Restore to alternate cluster**, then select the local cluster (ocp-b in this example), as seen in Figure 3-36.

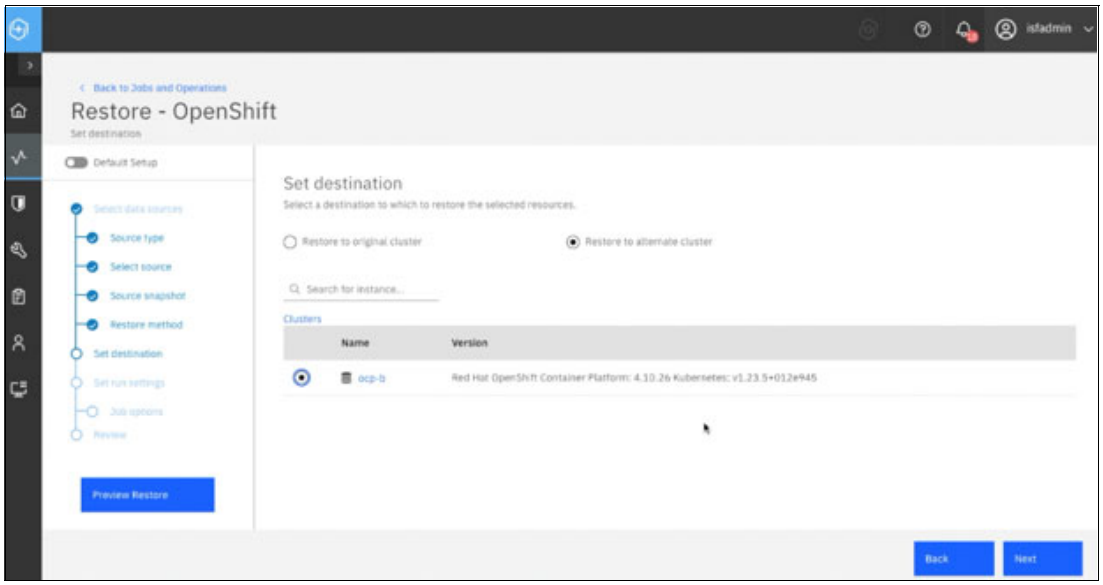

Figure 3-36 Restore to alternate cluster

Use the defaults for the remaining selections and submit the restore job.

From **Jobs and Operations** panel, watch the job progress. For this example, the IBM Storage Fusion namespace restore took five minutes and six seconds, as seen in Figure 3-37 on page 43.

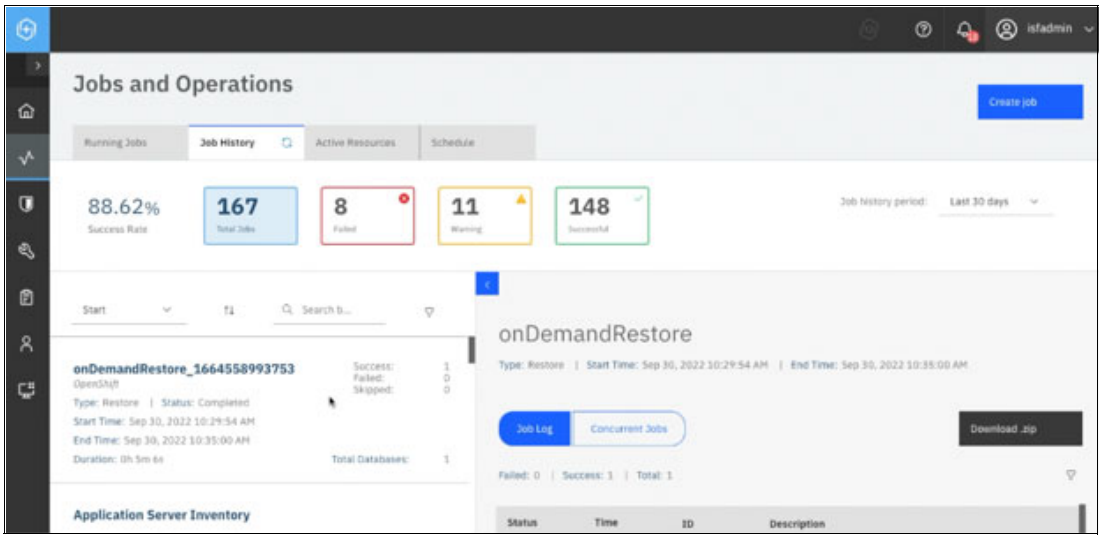

Figure 3-37 Monitor the SPP jobs for the ibm-spectrum-fusion-ns application restore

3.4 Restore Cloud Pak for Data

Before restoring Cloud Pak for Data operators and instance applications, we must install the `cpdbr-oadp` service in the Cloud Pak for Data operators and instance namespaces on the restore cluster. This will setup the required cluster roles, cluster rolebindings, and permissions for the `cpdbr` service to perform the restore operations. Reference the Cloud Pak for Data link:

https://www.ibm.com/docs/en/cloud-paks/cp-data/4.7.x?topic=utilities-installing-cpdbr-service-storage-fusion-integration

Using cpd-cli, install the `cpdbr oadp` service for Cloud Pak for Data operators. In this example, it is an Express install, where both the foundation and operators namespaces are ibm-common-services.

Note: The **ibm-common-services** namespace will not yet exist on the restore cluster, this step will only setup cluster permissions.

```
./cpd-cli oadp install --foundation-namespace=ibm-common-services
--operators-namespace=ibm-common-services --component=cpdbr-ops-hooks
--cpdbr-hooks-image-prefix=icr.io/cpopen/cpd --log-level=debug -verbose
```

Then install the `cpdbr oadp` service for Cloud Pak for Data instance. In this example, the instance is in the **czen** namespace.

Note: The **czen** namespace will not yet exist on the restore cluster, this step will only setup cluster permissions.

```
./cpd-cli oadp install --cpd-namespace=czen --component=cpdbr-hooks
--cpdbr-hooks-image-prefix=icr.io/cpopen/cpd --log-level=debug -verbose
```

3.4.1 Restore the Cloud Pak for Data operators

Now, we are ready to prepare for the restore of the Cloud Pak for Data operators. The first step is to click on **Create job** from the **Jobs and Operations** page as seen in Figure 3-38.

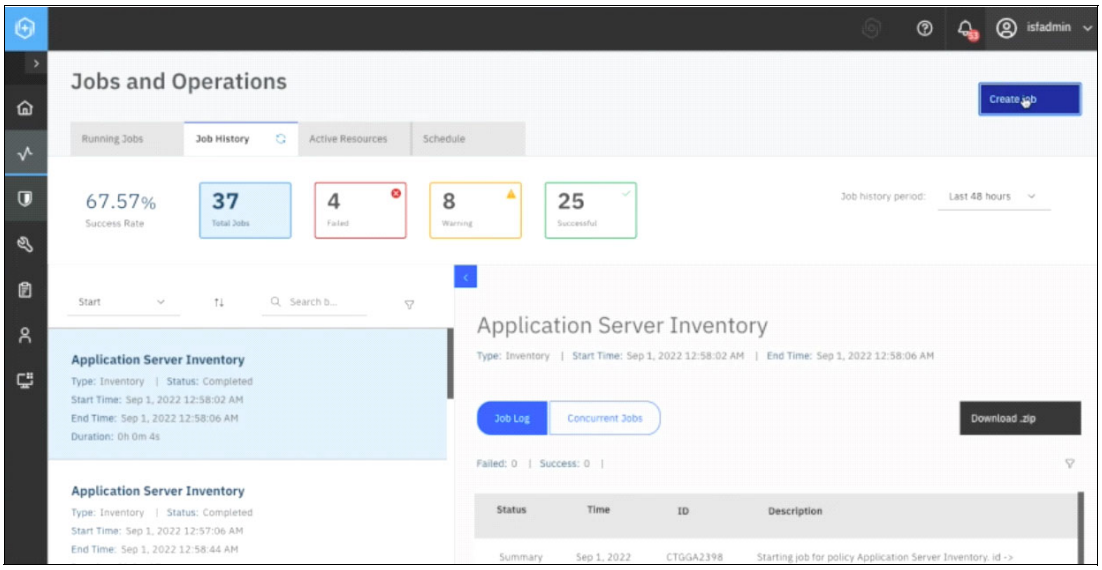

Figure 3-38 Select the Create job button

The next page is the **Create job** page and here, we selected **Restore** which brings up the **Restore - VMWare** page and we selected **OpenShift** under the **Containers** section as seen in the Figure 3-39.

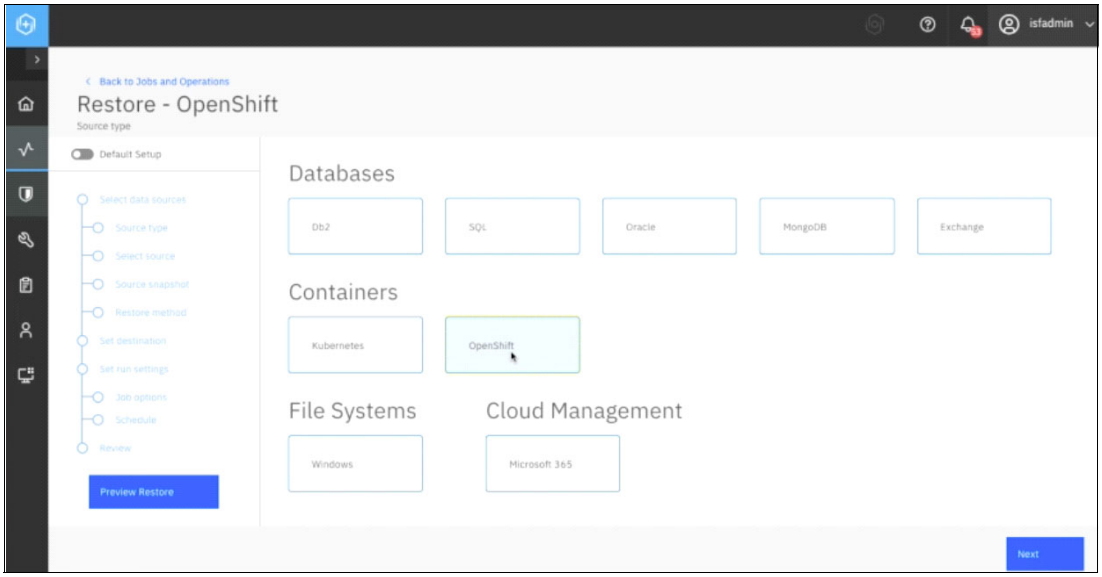

Figure 3-39 Select OpenShift

On the **Select source** menu, we clicked on the **ocp-a** cluster and we clicked on the **plus** icon next to the pvc for **ibm-common-services** and after doing this, it appears under the **Item** list on the right hand side. Afterwards, we clicked on the **Next** button as seen in the Figure 3-40 on page 45.

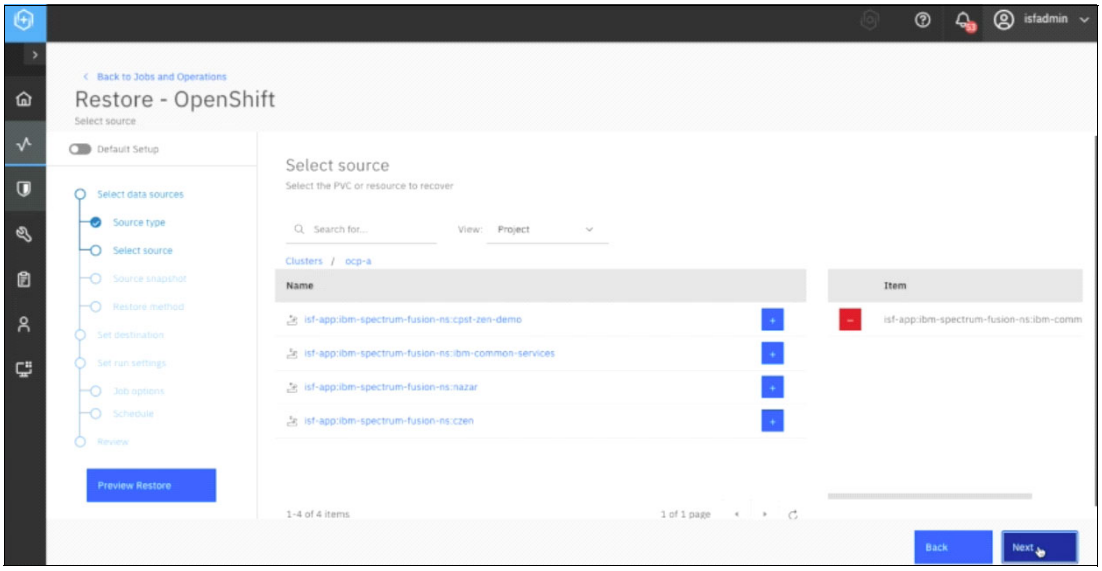

Figure 3-40 Select the ocp-a cluster

On the **Source snapshot** page, we expanded the **Origin** drop-down menu and selected **From Copy** as the source of the snapshot and selected **On-Demand** as the **Type of Restore** as shown in Figure 3-41.

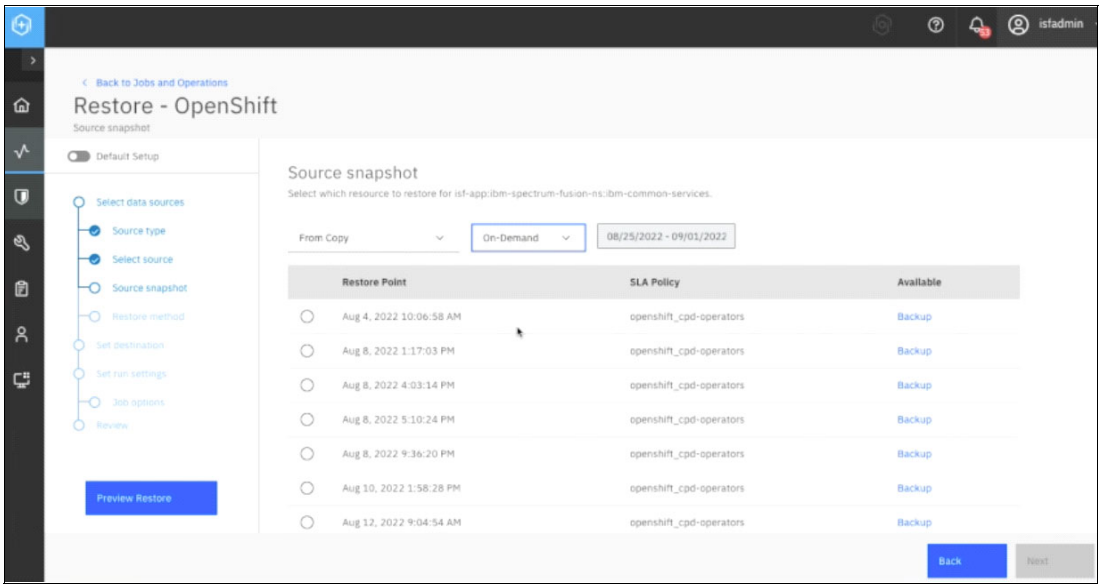

Figure 3-41 Selected From Copy from Origin menu and On-Demand from Type of restore

Then we scrolled down the page to find and select the most recent snapshot, and then clicked **Next** as can be seen in Figure 3-42.

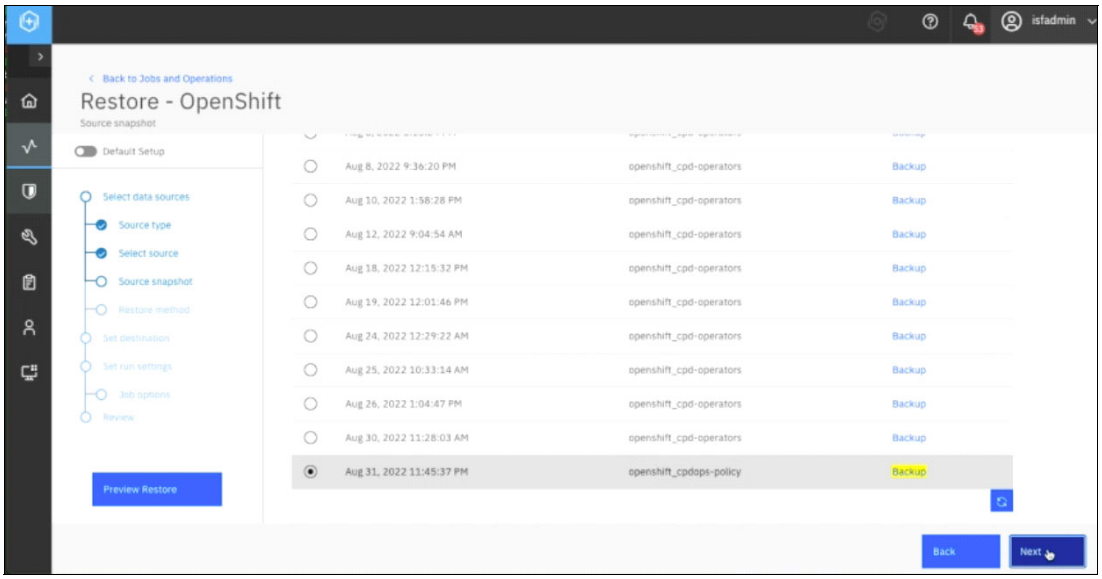

Figure 3-42 Scrolled to find the most recent snapshot then selected Next

On the **Restore method** page, we left the defaults and just clicked **Next** as shown in Figure 3-43.

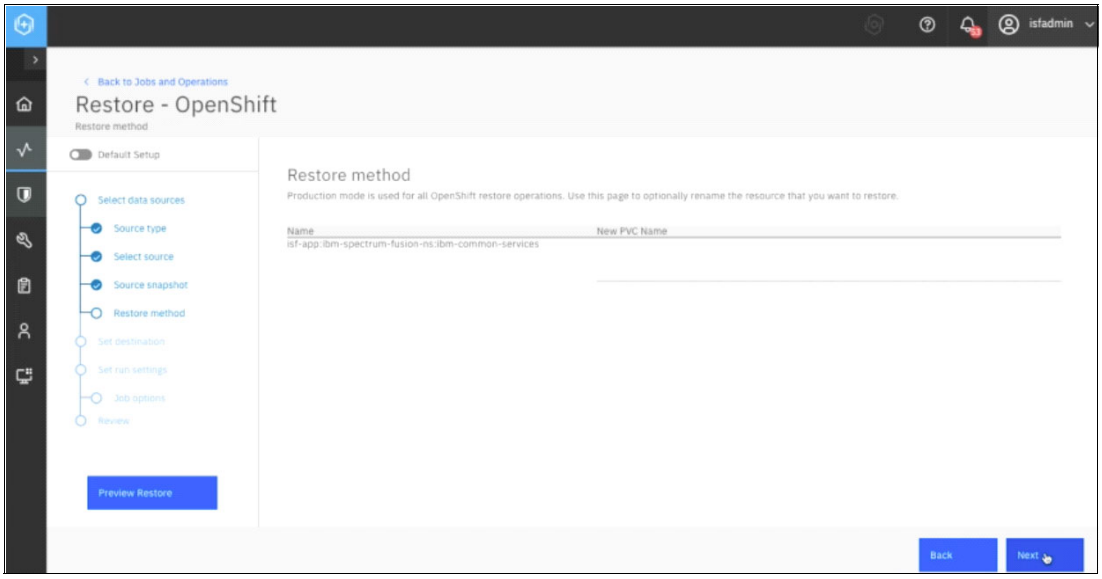

Figure 3-43 Accepted the defaults and selected Next

On the **Set destination** page, we selected the radio button next to **Restore to alternate cluster** and also selected the radio button next to our target cluster, ocp-b, and then clicked **Next** as depicted in Figure 3-44 on page 47.

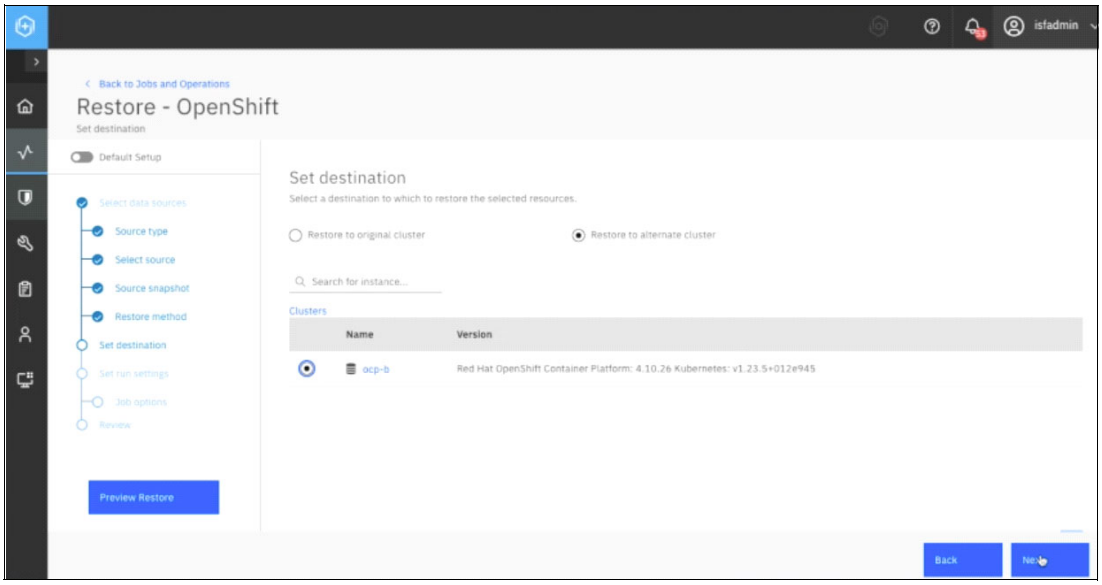

Figure 3-44 Select Restore to alternate cluster ocp-b

On the **Job options** page, we left the defaults and clicked **Next** as shown in Figure 3-45.

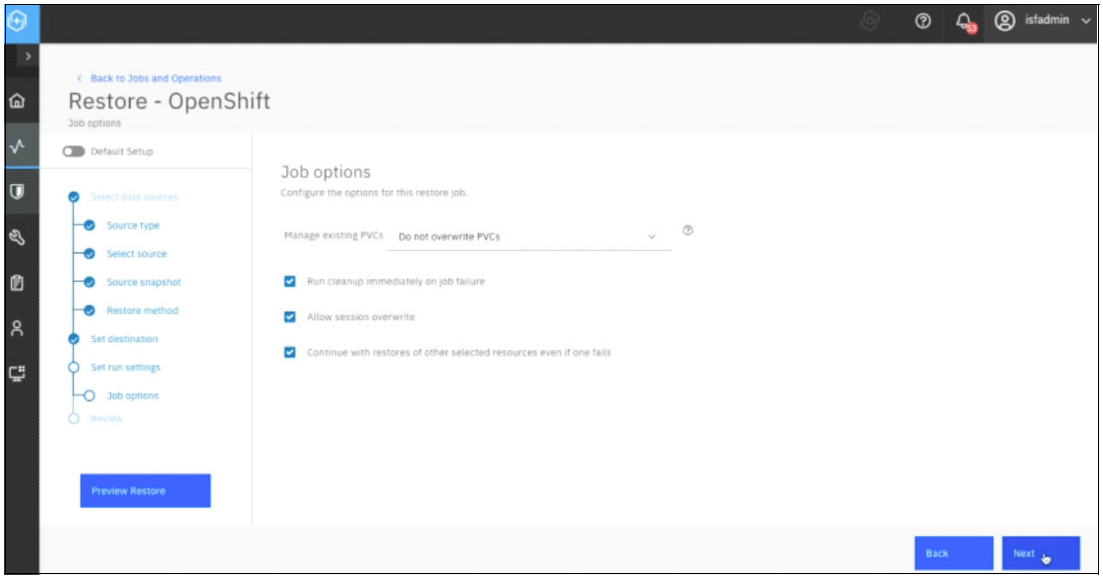

Figure 3-45 Accepted the defaults and selected Next

In the last page, we reviewed all our selections and clicked **Submit** as shown in Figure 3-46.

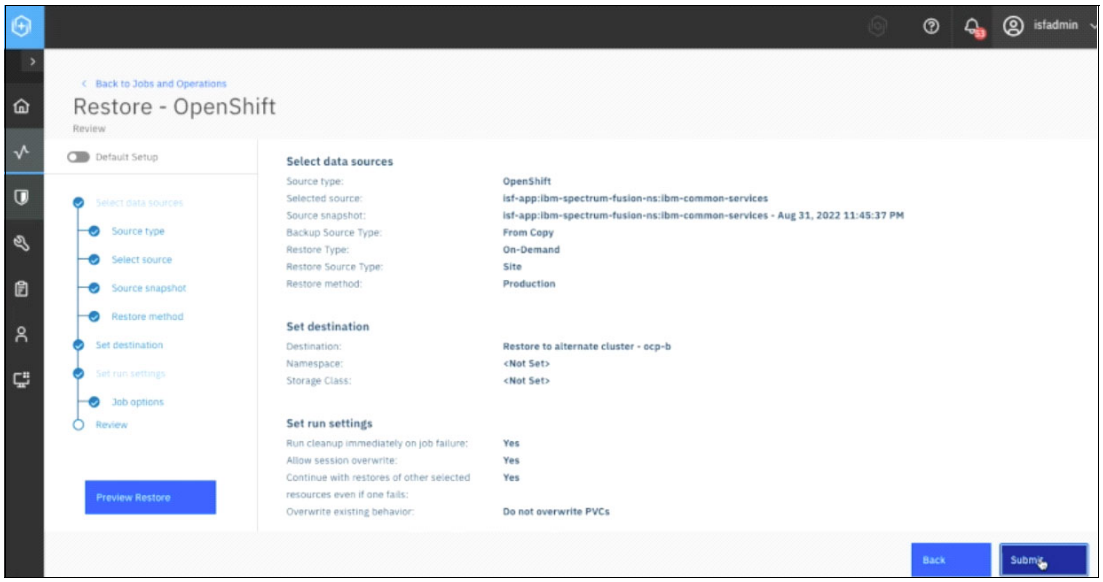

Figure 3-46 Review the selections and submit

The restore job is then created and the confirmation message appears as seen in Figure 3-47.

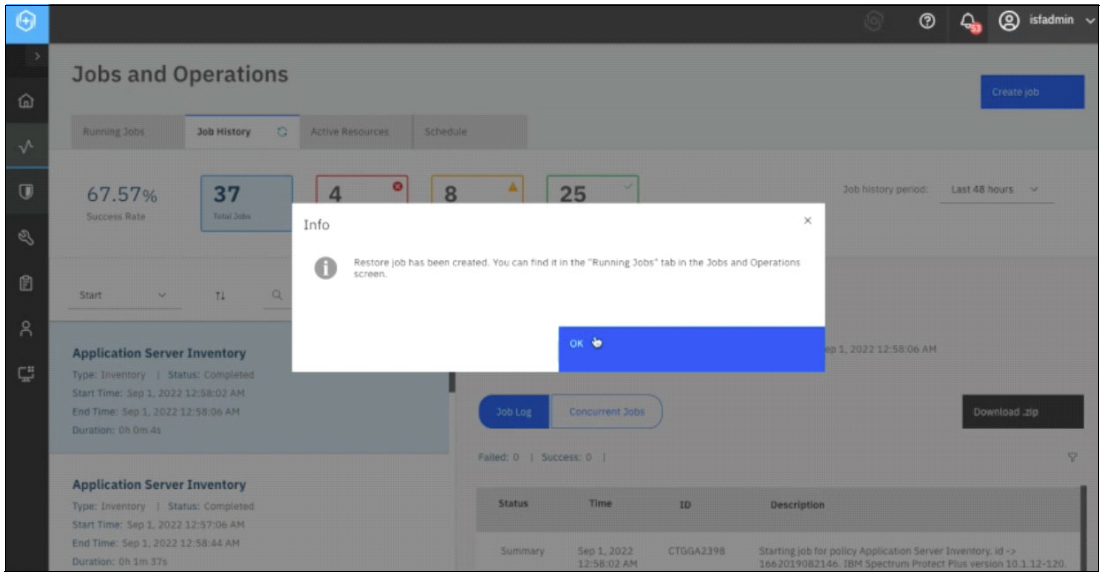

Figure 3-47 Message stating the Restore job was created and can be viewed in the Running Jobs tab

To view the progress of our restore job, we went to the **Running Jobs** tab and we can see the job is running in Figure 3-48 on page 49. This restore process involves restoring the catalog sources, the subscriptions, the csv installations and preparing all of the IBM Cloud Pak for Data operators.

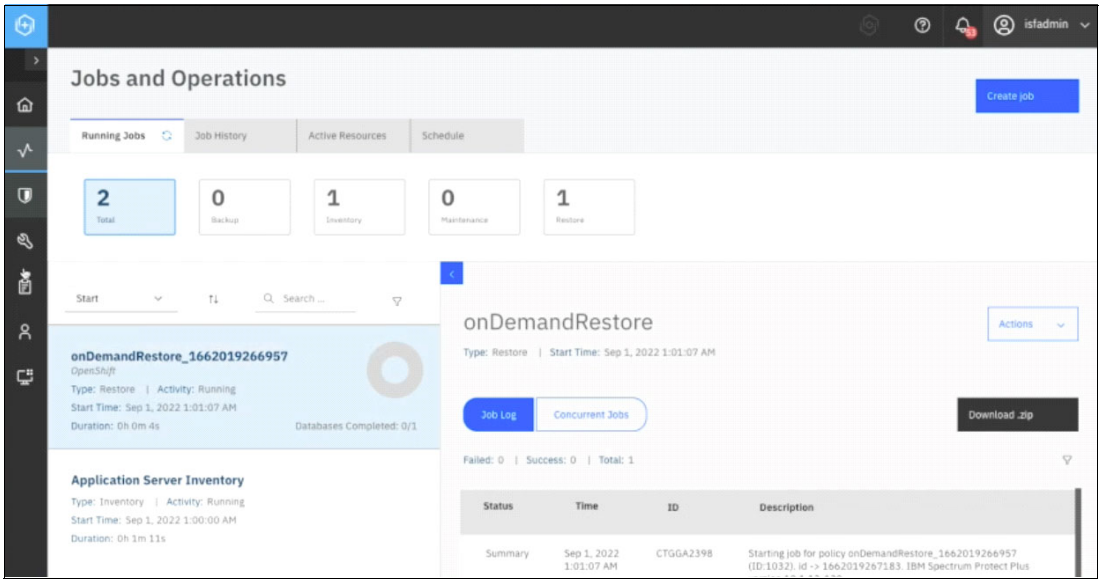

Figure 3-48 View the Running Jobs

The Cloud Pak for Data operators are installed in **ibm-common-services** and after a few minutes, we can see the pod in *Running* state by running the following command:

```
$ oc get pod -n ibm-common-services
```

We can also see that the recipe that was part of the backup is also restored by issuing the following command (see Figure 3-49):

```
$ oc get recipes.spp-data-protection.isf.ibm.com -n ibm-common-services
```

```
20220901-01:05:44 [ocp-b:default] ~ $ oc get pod -n ibm-common-services
NAME                                  READY   STATUS    RESTARTS   AGE
cpdbr-ops-service-6bcbc98fb7-kkmcw    1/1     Running   0          2m1s
20220901-01:05:51 [ocp-b:default] ~ $ oc get recipes.spp-data-protection.isf.ibm.com -n ibm-common-services
NAME                AGE
ibmcpd-operators    2m9s
```

Figure 3-49 Check that the pod is running and the recipe is restored

After about an hour and 17 minutes, we can see that the restore job for the Cloud Pak for Data operators completed successfully from the **Jobs and Operations** page in the SPP UI on ocp-b, which is our storage cluster as shown in Figure 3-50. In our source cluster we had several catalog sources and operators that were running on our source cluster and were restored. As such, the amount of time it takes for the restore to complete depends on the number of services and catalogs that are on the Cloud Pak for Data instance.

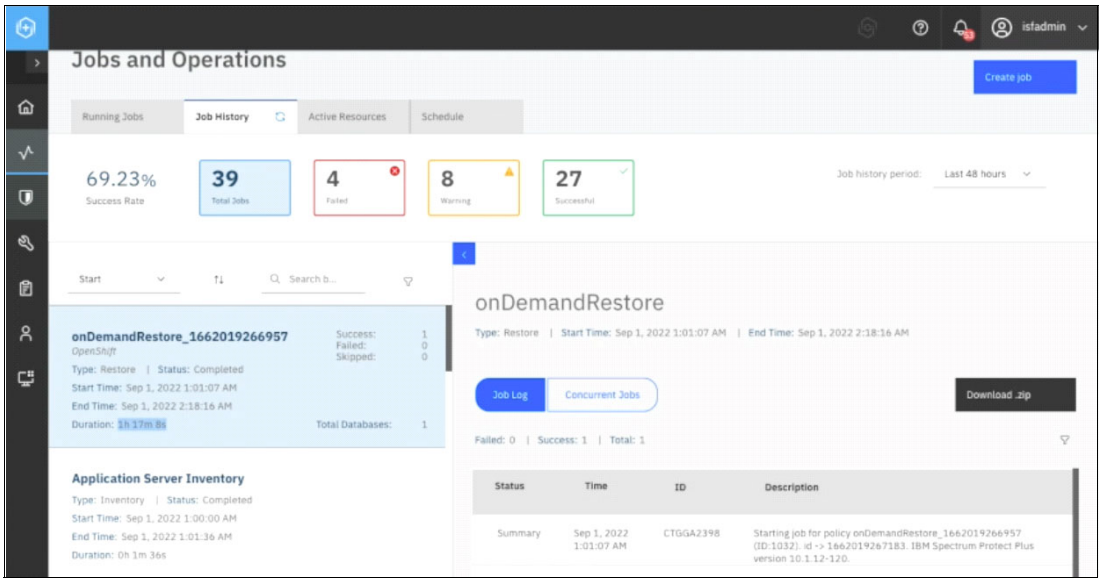

Figure 3-50 On-Demand Restore completed

We verified the creation of the catalog sources on our cluster (ocp-b) by running **`$ oc get catsrc -n openshift-marketplace`** as seen in Figure 3-51.

```
20220901-02:18:46 [ocp-b:default] ~ $ oc get catsrc -n openshift-marketplace
NAME                                  DISPLAY                                                TYPE   PUBLISHER          AGE
certified-operators                   Certified Operators                                    grpc   Red Hat            51d
community-operators                   Community Operators                                    grpc   Red Hat            51d
cpd-platform                          Cloud Pak for Data                                     grpc   IBM                74m
ibm-cloud-databases-redis-operator-catalog  ibm-cloud-databases-redis-operator-catalog        grpc   IBM                73m
ibm-cpd-ae-operator-catalog           Cloud Pak for Data IBM Analytics Engine powered by Apache Spark  grpc  IBM  72m
ibm-cpd-ccs-operator-catalog          CPD Common Core Services                               grpc   IBM                70m
ibm-cpd-datarefinery-operator-catalog Cloud Pak for Data IBM DataRefinery                    grpc   IBM                69m
ibm-cpd-datastage-operator-catalog    IBM CPD DataStage                                      grpc   IBM                67m
ibm-cpd-iis-operator-catalog          CPD IBM Information Server                             grpc   IBM                54m
ibm-cpd-scheduling-catalog            IBM Cloud Pak for Data Scheduler Catalog               grpc   IBM                53m
ibm-cpd-wkc-operator-catalog          CPD WKC                                                grpc   IBM                52m
ibm-cpd-wml-operator-catalog          Cloud Pak for Data Watson Machine Learning             grpc   IBM                50m
ibm-cpd-ws-operator-catalog           CPD IBM Watson Studio                                  grpc   IBM                49m
ibm-cpd-ws-runtimes-operator-catalog  CPD Watson Studio Runtimes                             grpc   IBM                48m
ibm-db2aaservice-cp4d-operator-catalog IBM Db2aaservice CP4D Catalog                         grpc   IBM                46m
ibm-db2oltp-cp4d-operator-catalog     IBM Db2oltp CP4D Catalog                               grpc   IBM                45m
ibm-db2uoperator-catalog              IBM Db2U Catalog                                       grpc   IBM                43m
ibm-db2wh-cp4d-operator-catalog       IBM Db2wh CP4D Catalog                                 grpc   IBM                42m
ibm-dmc-operator-catalog              ibm-dmc-operator-catalog                               grpc   IBM                41m
ibm-fdb-operator-catalog              IBM FoundationDB                                       grpc   IBM                39m
ibm-sppc-operator                     IBM SPPC Operator                                      grpc   IBM                12h
isf-catalog                           ISF Catalog                                            grpc   IBM                41h
manta-adl-operator-catalog            MANTA Operator Catalog                                 grpc   MANTA Software     38m
opencloud-operators                   IBMCS Operators                                        grpc   IBM                21d
redhat-marketplace                    Red Hat Marketplace                                    grpc   Red Hat            51d
redhat-operators                      Red Hat Operators                                      grpc   Red Hat            51d
20220901-02:18:58 [ocp-b:default] ~ $
```

Figure 3-51 Verify the creation of the catalog sources on ocp-b cluster

Additionally, we also verified the subscriptions on **ibm-common-services** by issuing **$ oc get sub -n ibm-common-services** as shown in Figure 3-52 on page 51.

Figure 3-52 Verify the subscriptions on ibm-common-services

3.4.2 Restoring the Cloud Pak for Data instance

Now we are ready to move on to the restore of the Cloud Pak for Data instance. First, we checked that the **czen** namespace does not exist by issuing **`$ oc get ns czen`**. Next, we installed `cpdbr-oadp` service. Note, it does not create that namespace but it will set up the cluster role-bindings and permissions required for **cpdbr-oadp** tool to be restored during the restore phase. To install, we issued the following command as seen in Figure 3-53.

```
./cpd-cli oadp install --cpd-namespace=czen --component=cpdbr-hooks
--cpdbr-hooks-image-prefix=icr.io/cpopen/cpd --log-level=debug -verbose
```

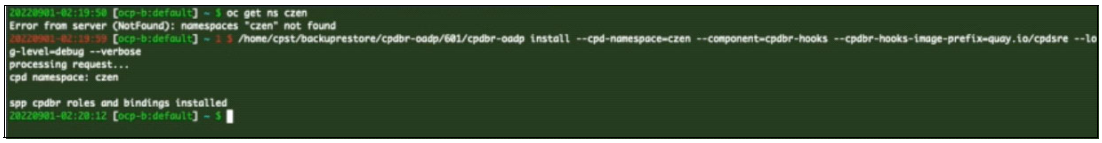

Figure 3-53 Check czen namespace does not exist and install cpdbr-oadp service

Now we are ready to go through the same restore process we previously ran. On the target cluster, which is ocp-b in our example, we go to the SPP UI **Jobs and Operations** page and click on **Create job** as seen in Figure 3-54.

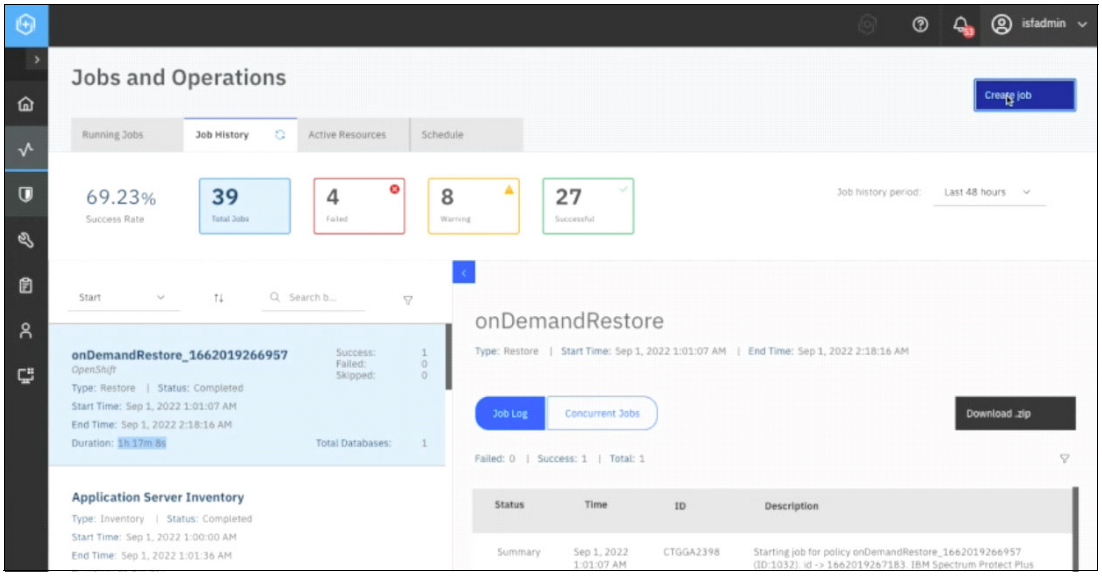

Figure 3-54 Select create job

On the **Create job** page, we select **Restore** as seen below in Figure 3-55.

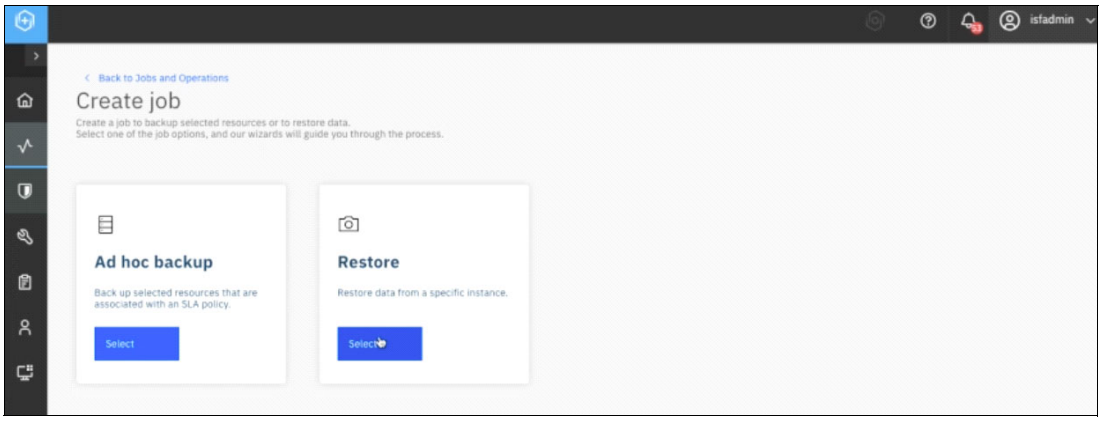

Figure 3-55 Select Restore

On the **Restore - OpenShift** page, we select **OpenShift** from the **Containers** section as seen in Figure 3-56.

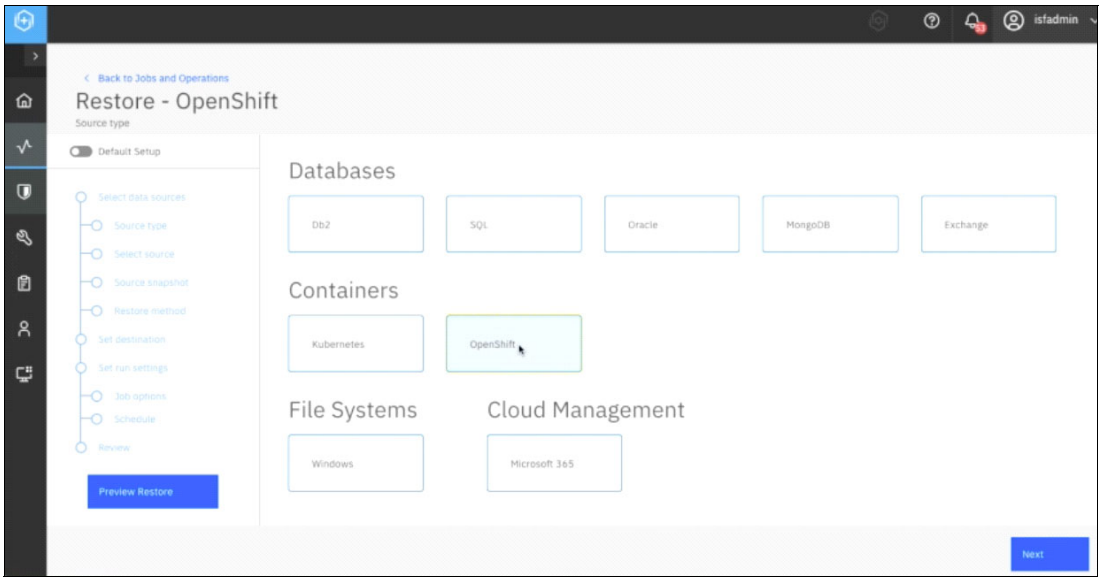

Figure 3-56 Restore OpenShift

On the **Select source** page, we selected the PVC to recover by clicking on the ocp-a cluster and that listed the PVCs within it. Then, we clicked on the **plus sign** next to the isf-app:ibm-spectrum-fusion-ns:czen PVC. This added the PVC to the **Item** list on the right hand side and then we clicked **Next** as shown in Figure 3-57 on page 53.

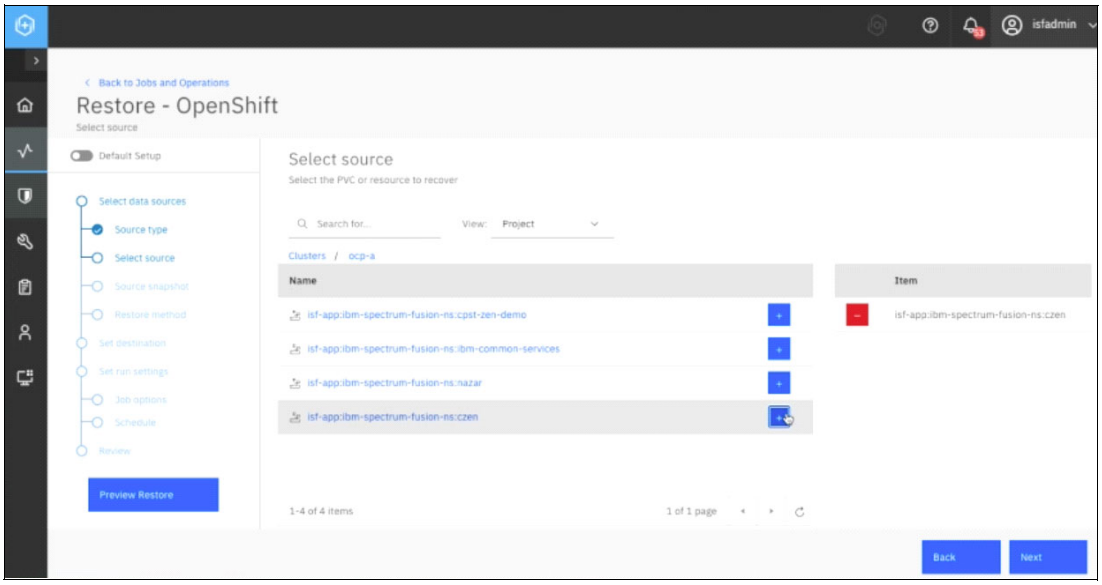

Figure 3-57 Select the PVC to recover

On the **Source snapshot** page, we selected the resource to restore by clicking on the drop-down arrow next to **Origin** and we selected **From Copy** and on the **Type of Restore** drop-down arrow, we selected **On-Demand**. This then listed the restore points, and there is only one in our example below. We selected it and then clicked **Next** as shown in Figure 3-58.

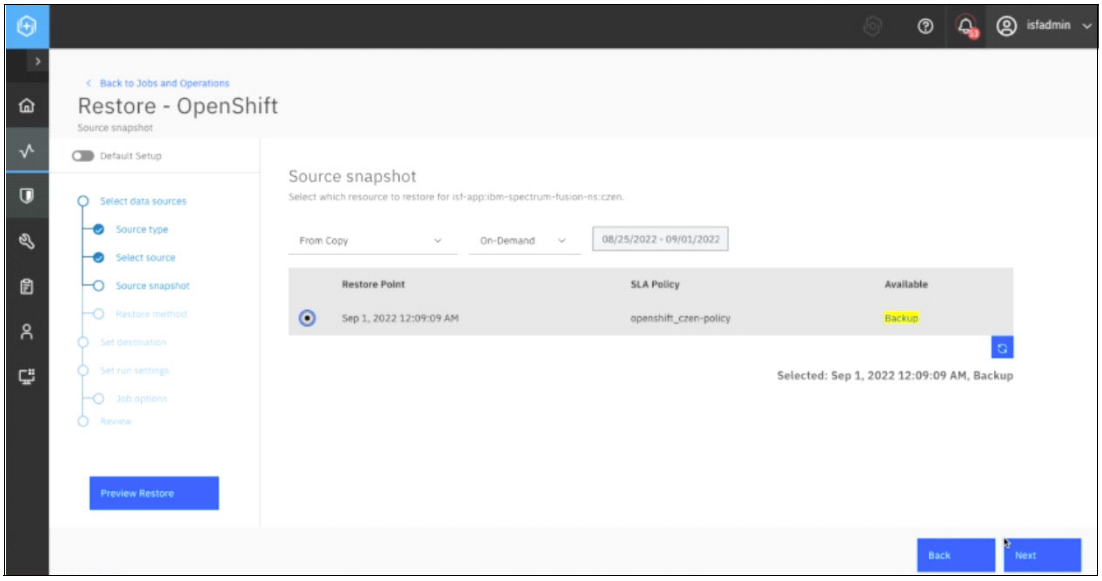

Figure 3-58 Selected From Copy, On-Demand, and Restore Point

On the **Restore method** page, we left the defaults and clicked **Next** as seen on Figure 3-59.

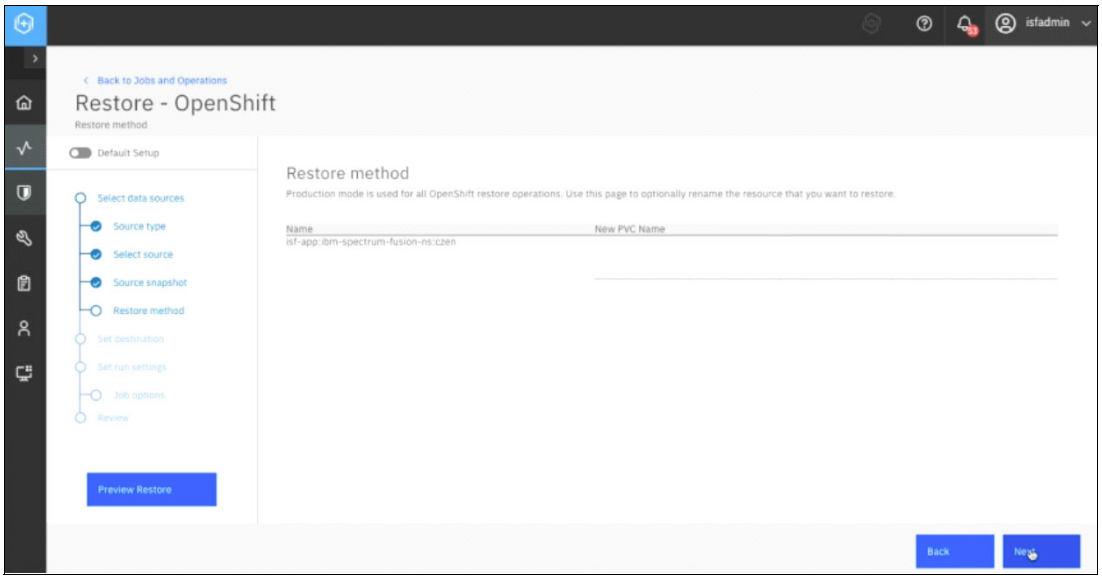

Figure 3-59 Accepted defaults for Restore method

On the **Set destination** page, we selected **Restore to alternate cluster** and ocp-b as depicted in Figure 3-60.

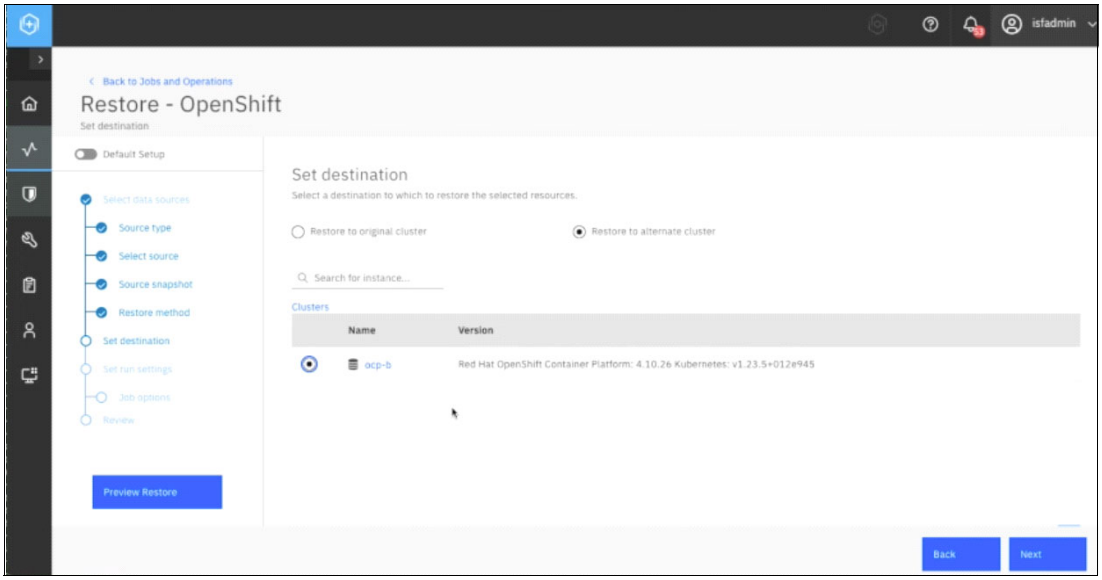

Figure 3-60 Selected Restore to alternate cluster

Also in the **Set destination** page, ensure that **Original Storage class** is selected for the **Storage class for restoring PVCs** and **Original Namespace** is selected for the **Namespace Destination**, as depicted in Figure 3-61 on page 55. Then we clicked **Next**.

Note: We must have the same storage classes on the target cluster that were used on the source cluster, otherwise the restore will fail.

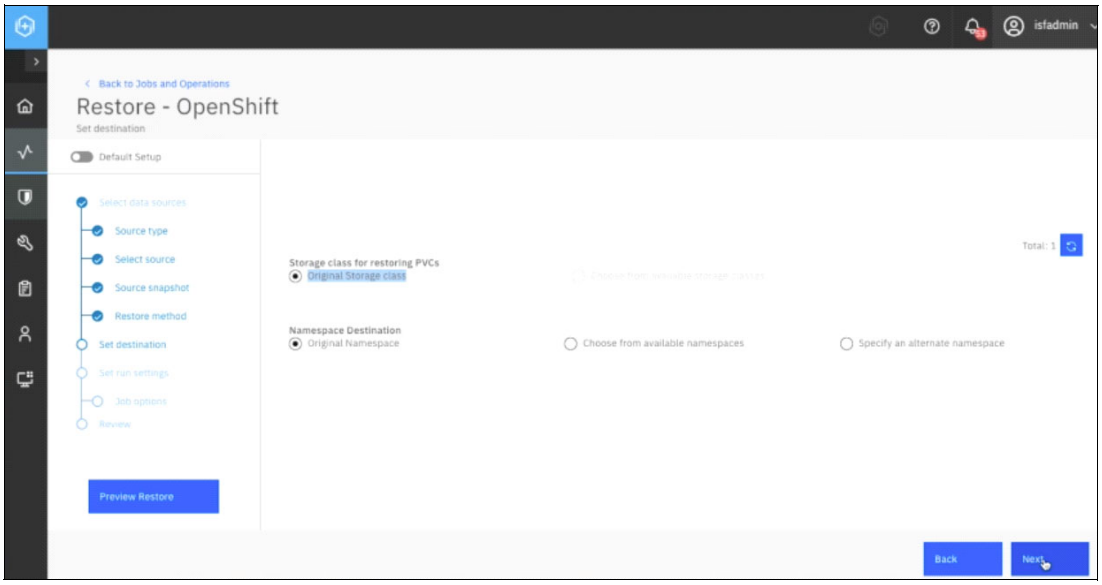

Figure 3-61 Selected Original Storage Class and Original Namespace

On the **Job options** page, we left the defaults and ensured the following options were selected before clicking **Next**: **Do not overwrite PVCs**, **Run cleanup immediately on job failure**, **Allow session overwrite**, **Continue with restores of other selected resources even if one fails** as shown in Figure 3-62.

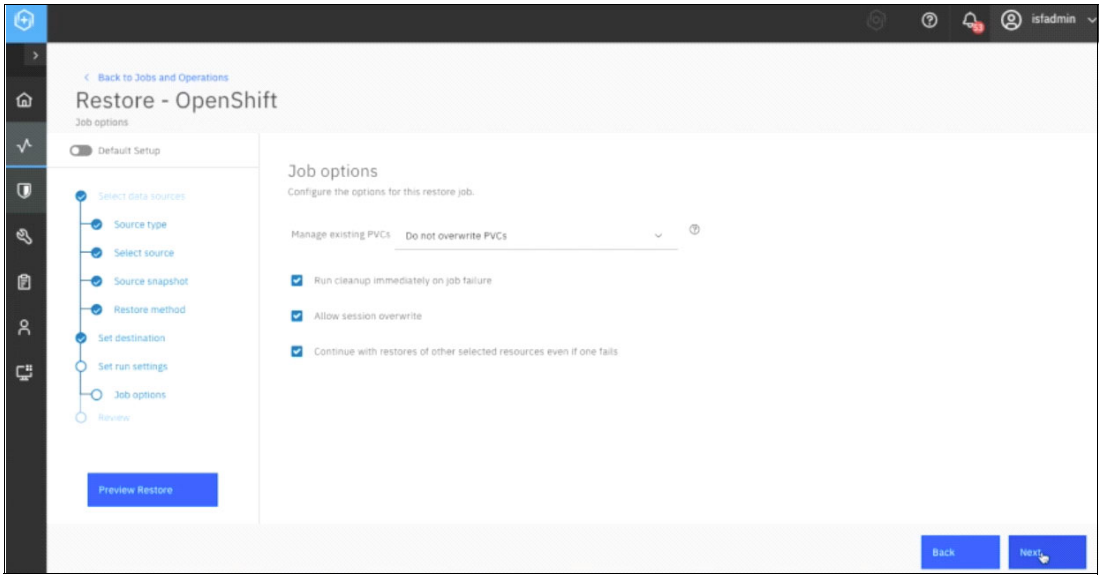

Figure 3-62 Accepted defaults and selected Do not overwrite PVCs and checked the boxes

On the **Review** page, we verified our selections were accurate and then clicked **Submit** as seen in Figure 3-63 to perform the restore.

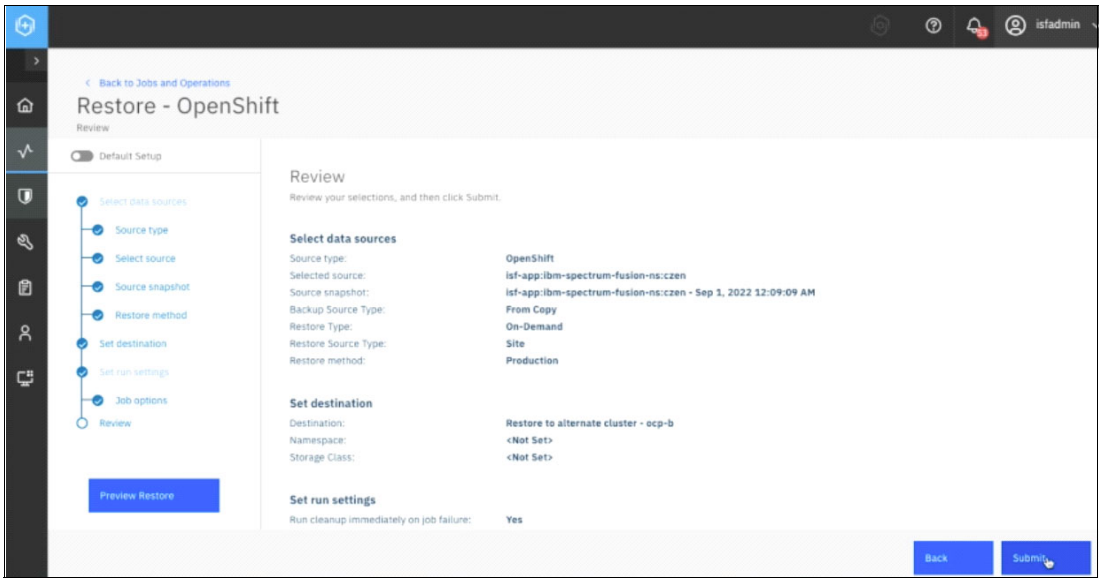

Figure 3-63 Verified that the selections were accurate

A confirmation then appeared informing us the restore job had been created and we clicked **OK** as seen in Figure 3-64.

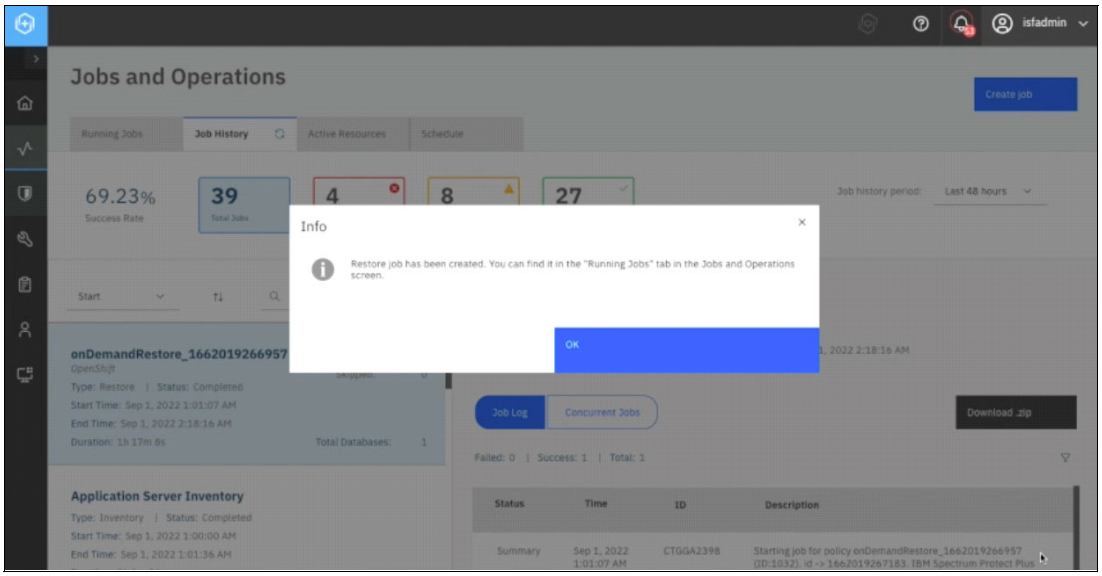

Figure 3-64 Confirmation that the restore job was created

3.5 Verification of the Cloud Pak for Data restore

After the job was completed, we went to the **Job History** tab within the **Jobs and Operations** and saw that this restore took 12 minutes and 6 seconds to complete. The restore included the restore of all of the services, the Cloud Pak for Data platform, and in our example, we had an IBM Db2 instance that was deployed on the original source cluster (ocp-a) and it was also restored and deployed on the target cluster (ocp-b). We verified the creation of all of the pods in the **czen** namespace by issuing **`$ oc get pod -n czen`** on our target cluster (ocp-b) as seen in Figure 3-65.

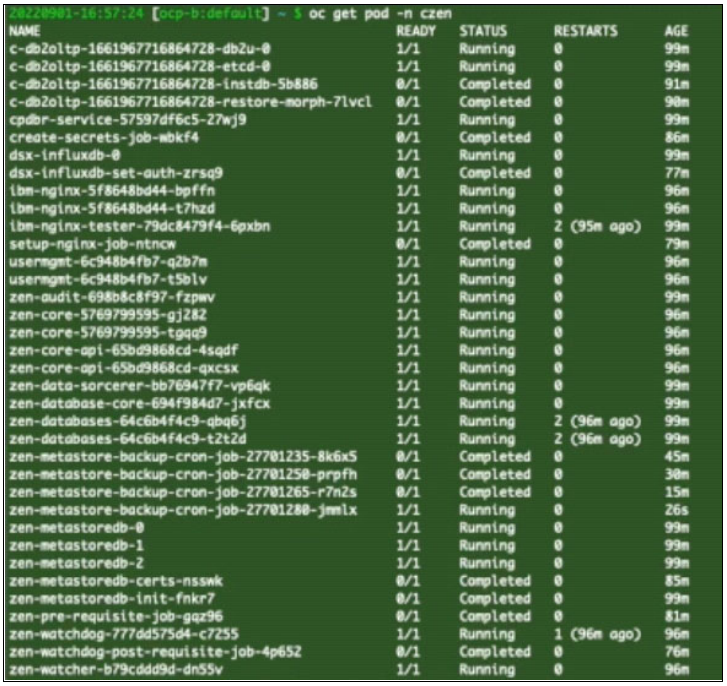

```
[20220901-16:57:24 [ocp-b:default] ~ $ oc get pod -n czen
NAME                                              READY   STATUS      RESTARTS      AGE
c-db2oltp-1661967716864728-db2u-0                 1/1     Running     0             99m
c-db2oltp-1661967716864728-etcd-0                 1/1     Running     0             99m
c-db2oltp-1661967716864728-instdb-5b886           0/1     Completed   0             91m
c-db2oltp-1661967716864728-restore-morph-7lvcl    0/1     Completed   0             90m
cpdbr-service-57597df6c5-27wj9                    1/1     Running     0             99m
create-secrets-job-wbkf4                          0/1     Completed   0             86m
dsx-influxdb-0                                    1/1     Running     0             99m
dsx-influxdb-set-auth-zrsq9                       0/1     Completed   0             77m
ibm-nginx-5f8648bd44-bpffn                        1/1     Running     0             96m
ibm-nginx-5f8648bd44-t7hzd                        1/1     Running     0             96m
ibm-nginx-tester-79dc8479f4-6pxbn                 1/1     Running     2 (95m ago)   99m
setup-nginx-job-ntncw                             0/1     Completed   0             79m
usermgmt-6c948b4fb7-q2b7m                         1/1     Running     0             96m
usermgmt-6c948b4fb7-t5blv                         1/1     Running     0             96m
zen-audit-698b8c8f97-fzpwv                        1/1     Running     0             99m
zen-core-5769799595-gjz82                         1/1     Running     0             96m
zen-core-5769799595-tgqq9                         1/1     Running     0             96m
zen-core-api-65bd9868cd-4sqdf                     1/1     Running     0             96m
zen-core-api-65bd9868cd-qxcsx                     1/1     Running     0             96m
zen-data-sorcerer-bb76947f7-vp6qk                 1/1     Running     0             99m
zen-database-core-694f984d7-jxfcx                 1/1     Running     0             99m
zen-databases-64c6b4f4c9-qbq6j                    1/1     Running     2 (96m ago)   99m
zen-databases-64c6b4f4c9-t2t2d                    1/1     Running     2 (96m ago)   99m
zen-metastore-backup-cron-job-27701235-8k6x5      0/1     Completed   0             45m
zen-metastore-backup-cron-job-27701250-prpfh      0/1     Completed   0             30m
zen-metastore-backup-cron-job-27701265-r7n2s      0/1     Completed   0             15m
zen-metastore-backup-cron-job-27701280-jmmlx      1/1     Running     0             26s
zen-metastoredb-0                                 1/1     Running     0             99m
zen-metastoredb-1                                 1/1     Running     0             99m
zen-metastoredb-2                                 1/1     Running     0             99m
zen-metastoredb-certs-nsswk                       0/1     Completed   0             85m
zen-metastoredb-init-fnkr7                        0/1     Completed   0             99m
zen-pre-requisite-job-gqz96                       0/1     Completed   0             81m
zen-watchdog-777dd575d4-c7255                     1/1     Running     1 (96m ago)   96m
zen-watchdog-post-requisite-job-4p652             0/1     Completed   0             76m
zen-watcher-b79cddd9d-dn55v                       1/1     Running     0             96m
```

Figure 3-65 Verified the creation of all the pods in the czen namespace on target cluster ocp-b

Then we verified the IBM Db2 instance deployment by issuing **`$ oc get db2ucluster -n czen`** as seen in Figure 3-66.

```
[20220901-17:00:34 [ocp-b:default] ~ $ oc get db2ucluster -n czen
NAME                        STATE   MAINTENANCESTATE   AGE
db2oltp-1661967716864728    Ready   None               100m
```

Figure 3-66 Verified the IBM Db2 instance deployment

We then described **zenservice** by issuing **`$ oc describe zenservice lite-cr -n czen`** and verified it is 100% complete as seen in Figure 3-67 on page 58.

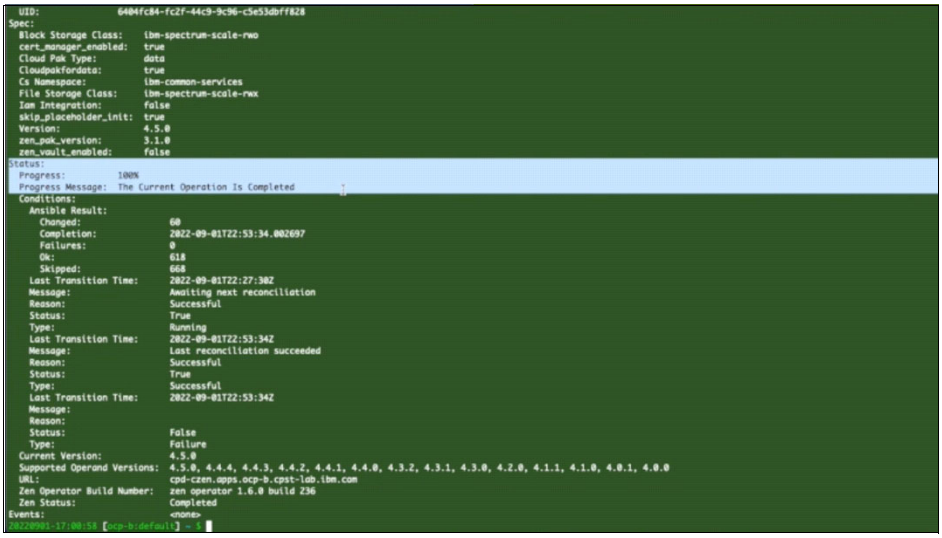

Figure 3-67 Described the zenservice and verified it is 100% complete

Now we can access the instance console for Cloud Pak for Data by first getting the link from our target cluster (ocp-b) by issuing **`$ oc get route -n czen`** as seen in Figure 3-68.

Figure 3-68 Get the link to access the instance console for Cloud Pak for Data from ocp-b

Once we log in, we clicked on **View all** from within the **My instances** section in the main page for the IBM Cloud Pak for Data console as shown in Figure 3-69.

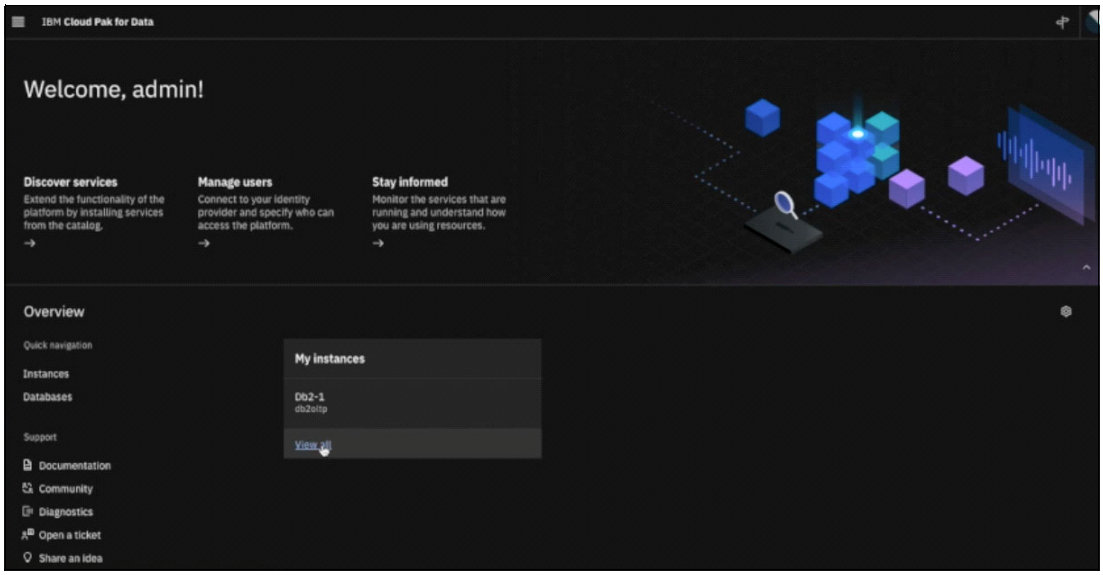

Figure 3-69 Selected View all from My instances for the IBM Cloud Pak for Data console

From within the **Instances** page, we confirmed that the instance name is the same as our source cluster (Db2-1) and the **Status** is green and **Running** as shown in Figure 3-70 on page 59.

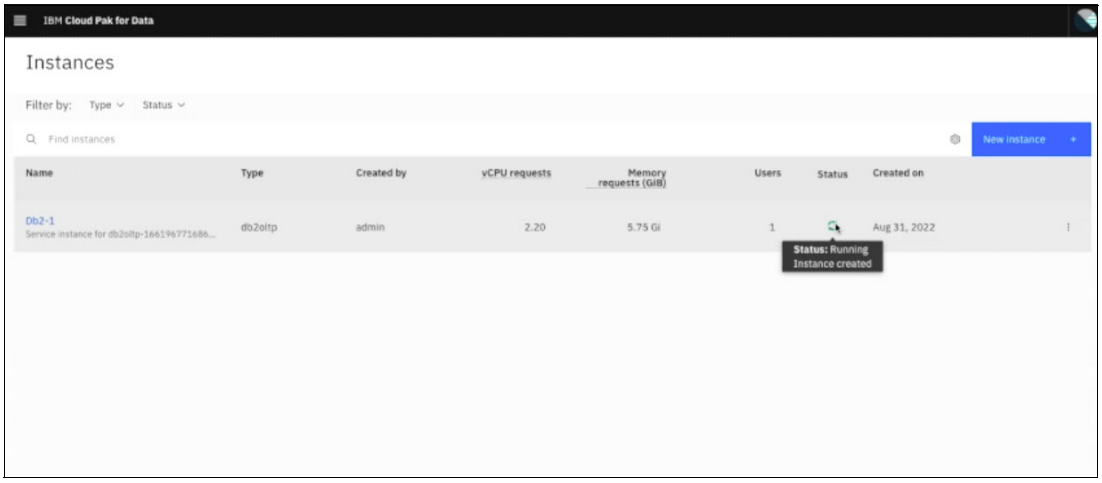

Figure 3-70 Confirmed instance name is the same as source cluster (Db2-1)

We then verified that all the PVCs for the **czen** namespace were restored on the target cluster (ocp-b) by issuing **`$ oc get pvc -n czen`** as seen in Figure 3-71.

```
20220901-17:01:08 [ocp-b:default] ~ $ oc get pvc -n czen
NAME                                        STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS                AGE
activelogs-c-db2oltp-1661967716864728-db2u-0   Bound    pvc-3b546421-a4cb-4110-9150-8ec98071e2ae   100Gi      RWO            ibm-spectrum-scale-rwx      105m
c-db2oltp-1661967716864728-backup           Bound    pvc-7a065a18-e65f-4d89-b24d-f38cb9c0f7fa   100Gi      RWX            ibm-spectrum-scale-rwx      105m
c-db2oltp-1661967716864728-meta             Bound    pvc-cb848300-1990-45dc-8733-4848ab250045   100Gi      RWX            ibm-spectrum-scale-rwx      105m
data-c-db2oltp-1661967716864728-db2u-0      Bound    pvc-ef1c5057-9095-4d0b-96ec-3ac9134044ee   100Gi      RWO            ibm-spectrum-scale-rwx      105m
data-dsx-influxdb-0                         Bound    pvc-108c2851-b9d5-49b9-957b-d3cfc8c58e2b   10Gi       RWO            ibm-spectrum-scale-rwo      105m
datadir-zen-metastoredb-0                   Bound    pvc-1950e7c8-a3b6-422c-9d0f-bbfba81ded80   10Gi       RWO            ibm-spectrum-scale-rwo      101m
datadir-zen-metastoredb-1                   Bound    pvc-96084c6c-c7f7-49eb-9b33-b7c0e8bcac5c   10Gi       RWO            ibm-spectrum-scale-rwo      101m
datadir-zen-metastoredb-2                   Bound    pvc-ef3ab757-a7a9-4da7-bce9-b5e53c1c8a71   10Gi       RWO            ibm-spectrum-scale-rwo      101m
tempts-c-db2oltp-1661967716864728-db2u-0    Bound    pvc-c8d8420e-8b5f-471b-8995-7bce64296e0c   100Gi      RWO            ibm-spectrum-scale-rwx      105m
user-home-pvc                               Bound    pvc-6f2a5280-1bbe-41a5-9ae7-d1bc422015de   10Gi       RWX            ibm-spectrum-scale-rwx      105m
20220901-17:02:05 [ocp-b:default] ~ $
```

Figure 3-71 Verified all the PVCs for the czen namespace were restored on the ocp-b

By running all the steps described in this paper, we have completed a successful online backup of the Cloud Pak for Data application and restored it to an alternate cluster using IBM Storage Fusion Data Protection.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *Accelerating IBM watsonx.data with IBM Storage Fusion HCI System,* REDP-5720
- *IBM Storage Fusion HCI System Metro Sync Dr Use Case*, REDP-5708
- *IBM Storage Fusion Product Guide*, REDP-5688
- *Multi-Factor Authentication Using IBM Security Verify for IBM Spectrum Fusion*, REDP-5662

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- IBM Storage Fusion

  https://www.ibm.com/products/storage-fusion
- IBM Documentation - IBM Storage Fusion

  https://www.ibm.com/docs/en/storage-fusion
- IBM Storage Fusion HCI announcements

  https://www.ibm.com/docs/en/search/storage%20fusion?type=announcement
- IBM Support - IBM Storage Fusion Support Reference Guide

  https://www.ibm.com/support/pages/ibm-storage-fusion-support-reference-guide
- IBM Documentation - IBM Cloud Pak for Data documentation

  https://www.ibm.com/docs/en/cloud-paks/cp-data
- IBM Documentation - IBM Spectrum Protect Plus documentation

  https://www.ibm.com/docs/en/spp/10.1.15?topic=product-support-documentation
- Evolving the IBM Storage Portfolio Brand Identity and Strategy

  https://www.ibm.com/cloud/blog/evolving-the-ibm-storage-portfolio-brand-identity-and-strategy
- IBM Storage

  https://www.ibm.com/storage

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

REDP-5706-00

ISBN 0738461156

Printed in U.S.A.