Redbooks
ibm.com/redbooks

IBM Storage Fusion HCI System
Metro Sync Disaster Recovery Use Case

Pallavi Singh
Prakash Chandraya
Seema Pandhre

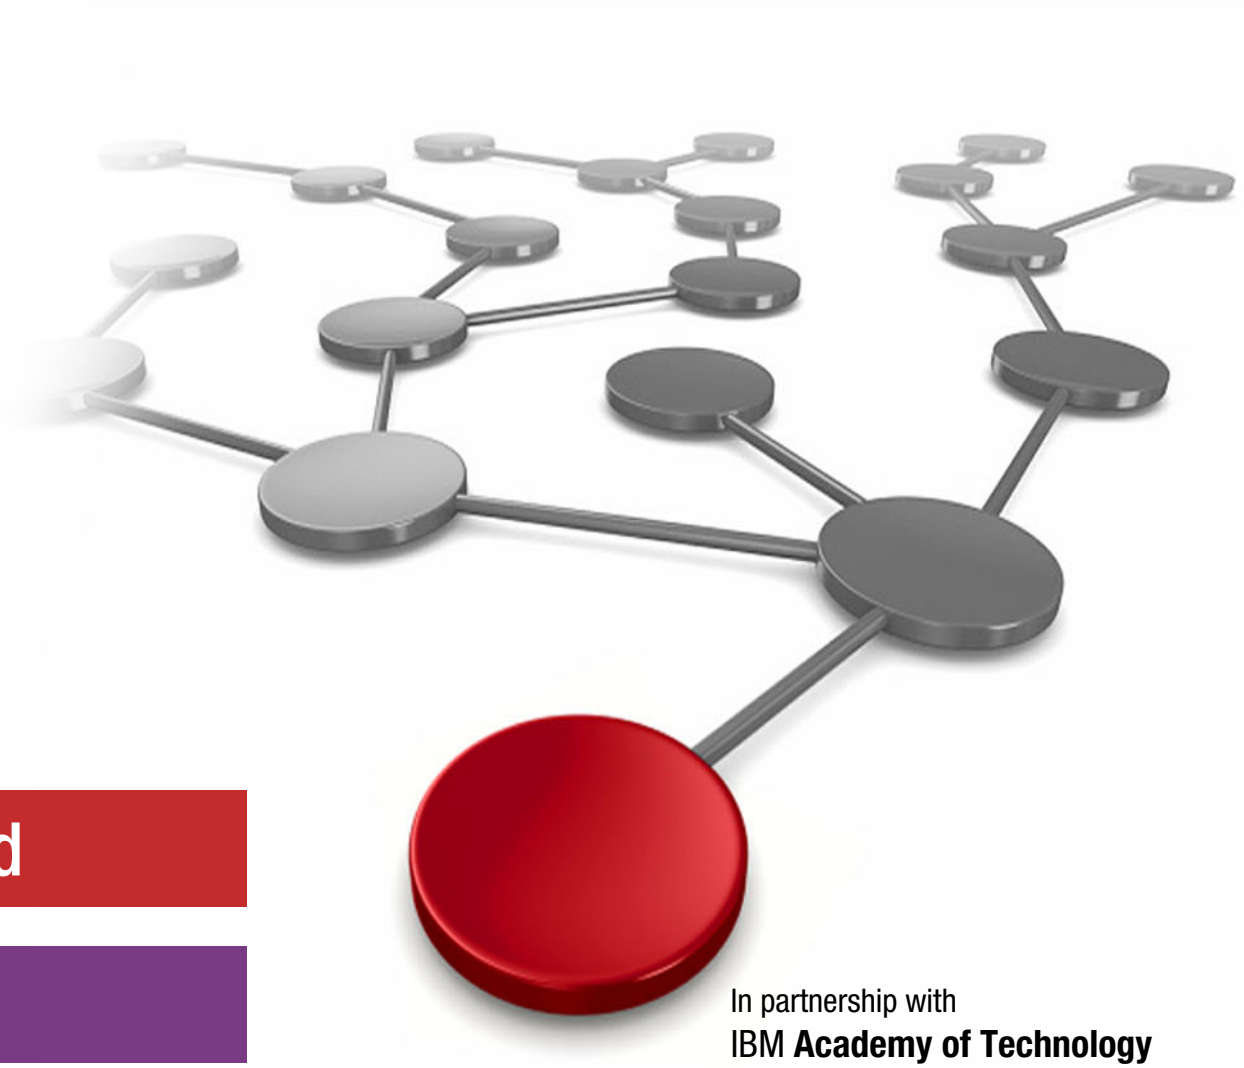

Cloud

Storage

In partnership with
IBM **Academy of Technology**

IBM

Redpaper

IBM Redbooks

IBM Storage Fusion HCI System: Metro Sync Disaster Recovery Use Case

June 2023

REDP-5708-00

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 4, Modification 0 of IBM Spectrum Fusion HCI.

Contents

Notices v
Trademarks vi

Preface vii
Authors vii
Now you can become a published author, too! viii
Comments welcome viii
Stay connected to IBM Redbooks ix

Chapter 1. Metro sync disaster recovery deployment models 1
1.1 Introduction 2
1.2 Deployment Models 3
1.2.1 Deployment Type #1 3
1.2.2 Deployment Type #2 4

Chapter 2. IBM Spectrum Fusion HCI Metro sync disaster recovery installation 5
2.1 Installation 6
2.1.1 Pre-requisites 6
2.1.2 Site1 installation for Deployment type #1 6
2.1.3 Site1 installation for Deployment type #2 6
2.1.4 Site2 Installation 15
2.1.5 Tiebreaker installation 23

Chapter 3. Application failover and failback 35
3.1 Application Failover/Failback between local site and remote site 36
3.1.1 Pre-requisites for application failover/failback 36
3.1.2 Setting up applications for disaster recovery 37
3.1.3 Failover 43
3.1.4 Failback 49

Related publications 55
IBM Redbooks 55
Online resources 55
Help from IBM 55

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.

IBM®	IBM Spectrum Fusion™	Redbooks (logo) ®
IBM Spectrum®	Redbooks®	Spectrum Fusion™

The following terms are trademarks of other companies:

Ansible, OpenShift, Red Hat, 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

Metro sync disaster recovery (DR) provides two-way synchronous data replication between IBM Spectrum Fusion™ HCI clusters installed at two sites. In the event of a site disaster, applications can be failed over to the second site. The replication between the sites is synchronous, hence, the Metro sync DR solution is only available for metropolitan distance data centers with 40 millisecond latency or less.

Note: The procedures described in this paper for IBM Spectrum Fusion HCI 2.4 Metro sync DR are the same for IBM Storage Fusion HCI 2.5.2 Metro-DR.

This IBM Redpaper publication will help you install and configure the new Metro sync DR function). The use case will show the end to end process with the failover and failback of the WordPress application.

Note: IBM Spectrum Fusion HCI and IBM Spectrum Fusion have become IBM Storage Fusion HCI System and IBM Storage Fusion. This edition uses 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.

Authors

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

Pallavi Singh is a Senior Technical Staff Member and works as an architect with IBM® System Labs. She has 15 years of experience in product architecture and design. Pallavi is an IBM Redbooks® Author, a member of the IBM Academy of Technology, and a Corporate Service Corps alumni.

Prakash Chandraya is an Advisory Software Engineer and works as a quality lead with IBM Systems Labs. He has 16 years of experience in product quality assurance. Prakash is a first time Redbooks author and has expertise in leadership, problem solving, mentorship, customer support, and time management.

Seema Pandhre is an Advisory Software Engineer and works as a quality lead with IBM Systems Labs. She has 15 years of industry experience in the software quality assurance field. She has been with IBM since 2007 and her areas of expertise include virtualization, cloud computing, and software engineering.

Special thanks to the following IBM Software Engineers:

Pallavi Joshi, IBM Systems, Pallavi Joshi is a software engineer with IBM Systems Labs. She has experience in software development across various software domains.

Kaustubh I Katruwar, IBM Systems, Kaustubh Indrakant Katruwar is a senior software engineer with the IBM Systems Labs. His experience spans across storage software encompassing clustered file systems, authentication and authorization, file protocols and multi protocol access, encryption, and disaster recovery. He is a storage security enthusiast with inclination towards secure engineering practices and security posture.

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

Larry Coyne
IBM Redbooks, Tucson Center

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

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

Stay connected to IBM Redbooks

- Find us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

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

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

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

1

Metro sync disaster recovery deployment models

IBM Spectrum Fusion HCI provides a feature, Metro sync DR, to achieve a highly available storage infrastructure. This chapter describes the Metro sync DR use case for disaster recovery.

1.1 Introduction

Metro sync DR (Disaster Recovery) provides two-way synchronous data replication between IBM Spectrum Fusion HCI clusters installed at two sites. In the event of a site disaster, applications can be failed over to the second site. The replication between the sites is synchronous, hence, the Metro sync DR solution is only available for metropolitan distance data centers with 40 millisecond latency or less.

The Metro sync DR architecture consists of the following components as shown in Figure 1-1.

1. Home Site - This is the first site in the Metro sync DR configuration, also referred to by Local site or Home Site or Site1 in this guide.
2. Remote Site - This is the second site in the Metro sync DR configuration, also referred to as Remote Site or Site2 in this guide.
3. Tiebreaker

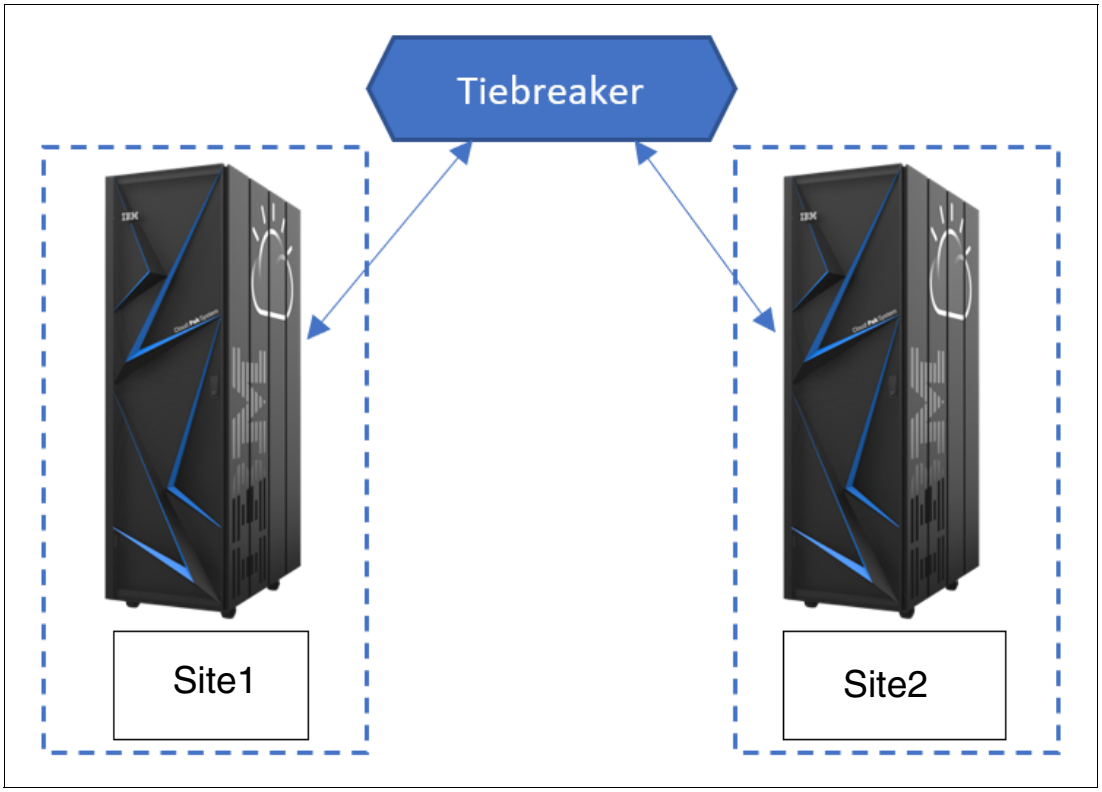

Figure 1-1 Metro sync DR configuration

IBM Spectrum Fusion HCI is based on IBM Spectrum Scale and uses its “stretch cluster” feature to provide a unique active-active resiliency across data centers that are up to 150km distant. This is achieved via spanning an IBM Spectrum Scale file system across two IBM Spectrum Fusion and Red Hat OpenShift Container Platform (OCP) clusters and synchronously replicating the data between both availability zones as shown in Figure 1-2 on page 3. The goal is to allow for an IBM Spectrum Fusion to synchronously replicate data to another nearby IBM Spectrum Fusion, each with their own OpenShift clusters.

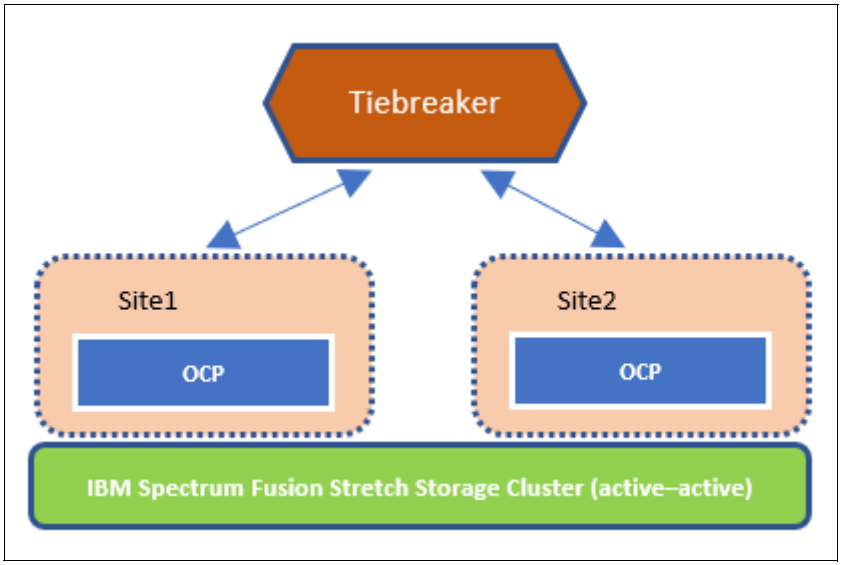

Figure 1-2 Metro sync DR component architecture

1.2 Deployment Models

There are two deployment models for Metro sync DR configuration.

1.2.1 Deployment Type #1

As a client, you buy a single IBM Spectrum Fusion appliance. Later, as the business requirements increase, you need another IBM Spectrum Fusion appliance as well if you want to achieve disaster recovery. In this case as shown in Figure 1-3, you convert the first appliance into Site1 (Refer to 2.1.2, “Site1 installation for Deployment type #1” on page 6 for the steps) and then continue with installation of the second appliance as Site2.

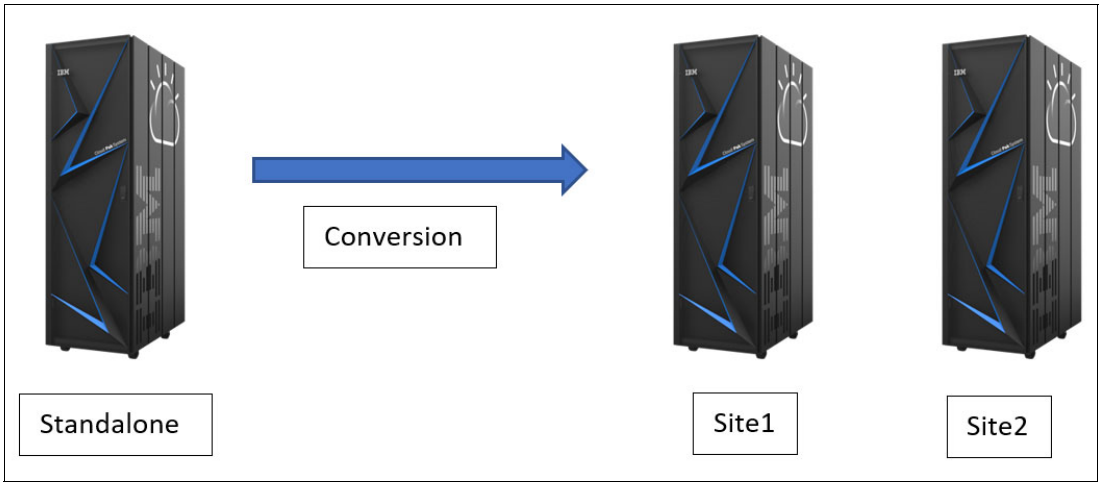

Figure 1-3 Deployment model Type #1 for setting up disaster recovery

1.2.2 Deployment Type #2

As a client, based on business requirements, you plan to have disaster recovery. In this case, you would have to purchase two IBM Spectrum Fusion appliances, see Figure 1-4. The installation of the two appliances will be done in sequence. During IBM Spectrum Fusion appliance setup, on the Disaster Recovery step, ensure to select the appropriate options. For the first rack install, select Site1 and for the second appliance install, select Site2. Refer to 2.1, “Installation” on page 6.

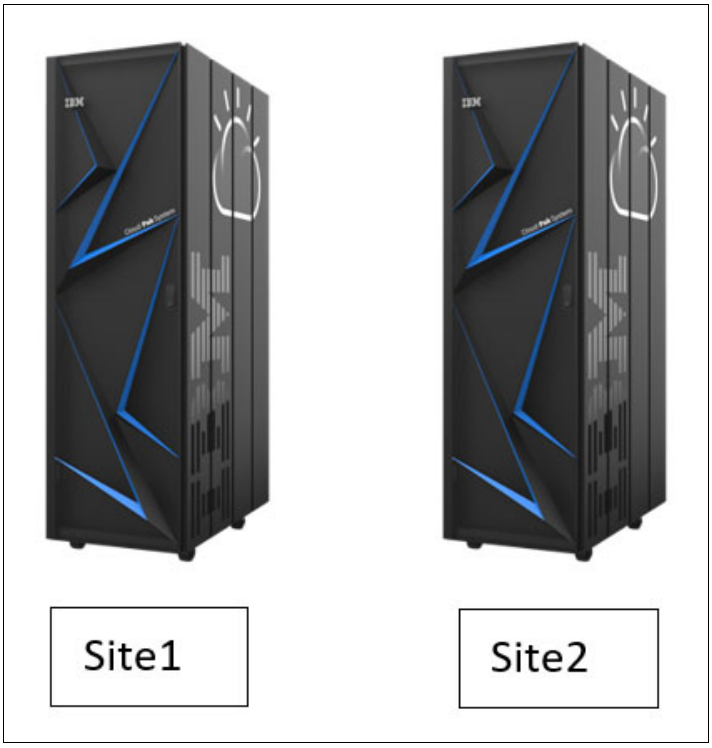

Figure 1-4 Deployment model Type #2 for setting up disaster recovery

2

IBM Spectrum Fusion HCI Metro sync disaster recovery installation

IBM Spectrum Fusion HCI provides a feature, Metro sync DR, to achieve a highly available storage infrastructure. This chapter describes the Metro sync DR installation steps based on the selected deployment model.

2.1 Installation

This section will guide you through the installation of the IBM Spectrum Fusion appliances in a Metro sync DR configuration. It will also guide to install the Tiebreaker and configure it.

2.1.1 Pre-requisites

Before you begin installation, you need to ensure that the IBM Spectrum Fusion HCI appliance is setup in datacenter as per guidance from IBM. Refer to the following links in IBM Documentation to setup the appliance:

1. Instructions to setup the appliance in Data Center

   https://www.ibm.com/docs/en/spectrum-fusion/2.4?topic=hci-planning-prerequisites

2. Instructions to setup network and connectivity between Metro sync DR clusters

   https://www.ibm.com/docs/en/spectrum-fusion/2.4?topic=recovery-general-metro-sync-dr-prerequisites

2.1.2 Site1 installation for Deployment type #1

If you have considered Deployment type #1 from 1.2.1, “Deployment Type #1” on page 3, you would need to convert the configuration from standalone to Site1. Use the following steps to complete the Site1 configuration:

1. Log in to OpenShift Container Platform web console of Standalone site.
2. Go to ibm-spectrum-fusion-ns namespace.
3. Go to **Workload** → **Secret**, then, search and open secret: userconfig-secret.
4. Change the value of metrodrsite to Site1 i.e. “metrodrsite”: “site1” as shown in Figure 2-1 and save.

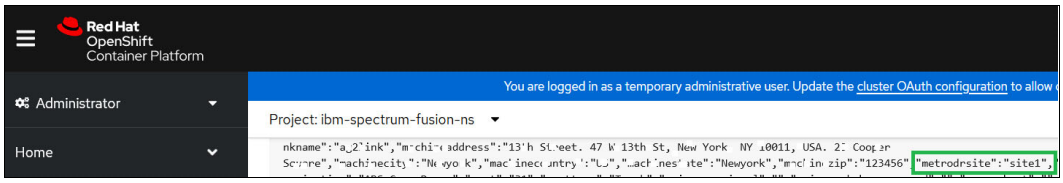

Figure 2-1 Update metrodrsite variable to Site1

The Disaster recovery section in the **Dashboard** page is enabled and the **Disaster recovery** page is available in the menu.

2.1.3 Site1 installation for Deployment type #2

If you have considered deployment type #2 from 1.2.2, “Deployment Type #2” on page 4, follow these steps to start the Site1 install.

Network configuration

Follow these steps to set up the network configuration:

1. Enter the details of the Network configuration as show in Figure 2-2 on page 7.
2. Click on **Validate Network** button.

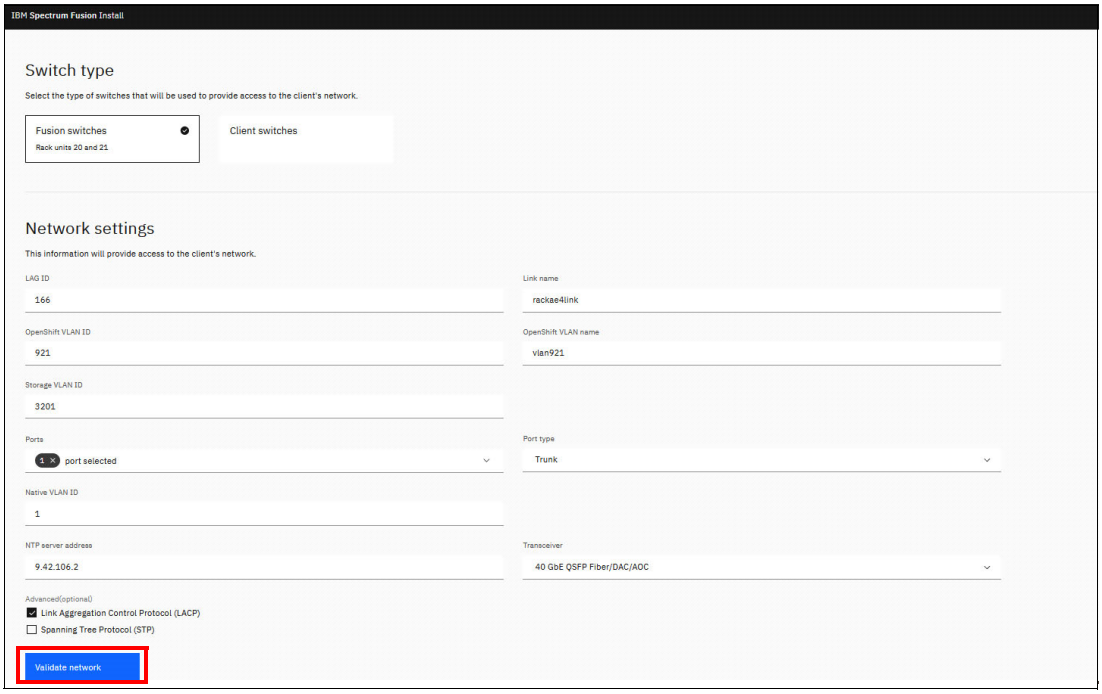

Figure 2-2 Network Configuration

The Network configuration provisions Node IP, configures DHCP, and NTP as shown in Figure 2-3, Figure 2-4, and Figure 2-5 on page 8.

Figure 2-3 Network configuration - Provisioning node IP

Figure 2-4 Network configuration - DHCP configuration

Figure 2-5 Network configuration - NTP configuration

3. The network configuration is completed successfully as shown in Figure 2-6.

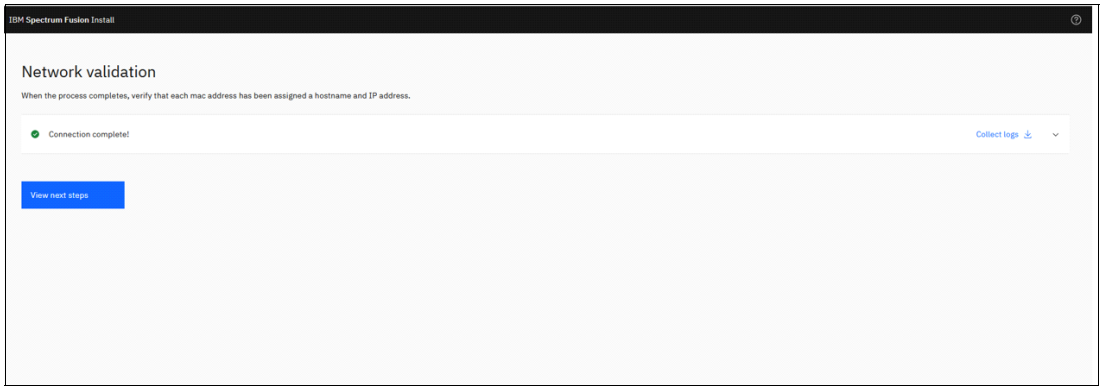

Figure 2-6 Network configuration - Successful completion

4. Click the **View next steps** button. This page will continue with the installation as shown in Figure 2-7.

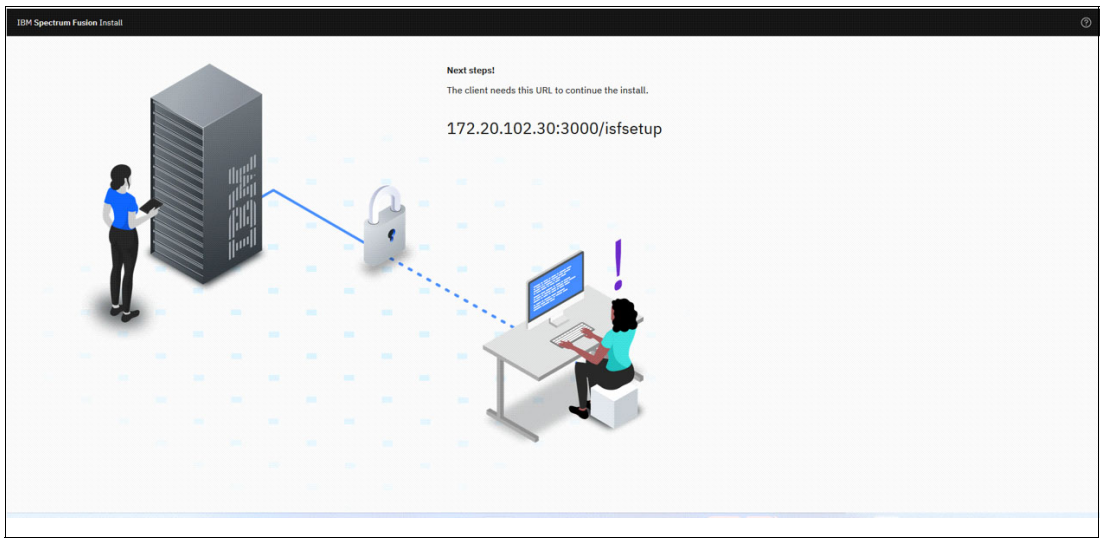

Figure 2-7 Next Steps

IBM Spectrum Fusion installer

The following steps describe the installation process for IBM Spectrum Fusion:

1. To proceed further with IBM Spectrum Fusion install, use the URL as shown in Figure 2-7 for your system.

2. You are presented with a **License agreement** page. Read the License agreement and Privacy policy. Then, accept the license and click on **Continue** button as shown in Figure 2-8.

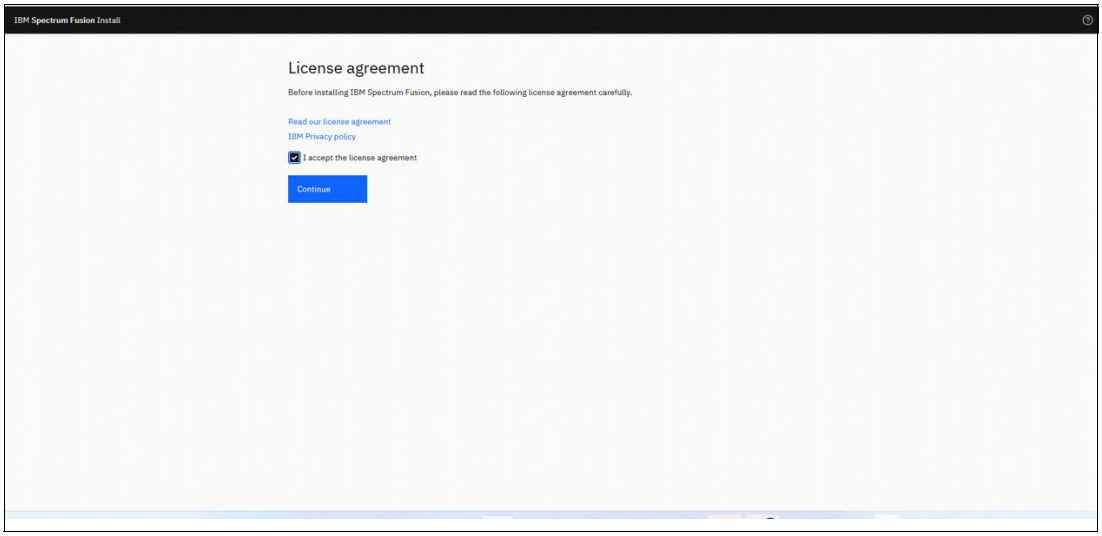

Figure 2-8 License agreement

3. The **Getting Started** page displays the procedure for Install process as shown in Figure 2-9.

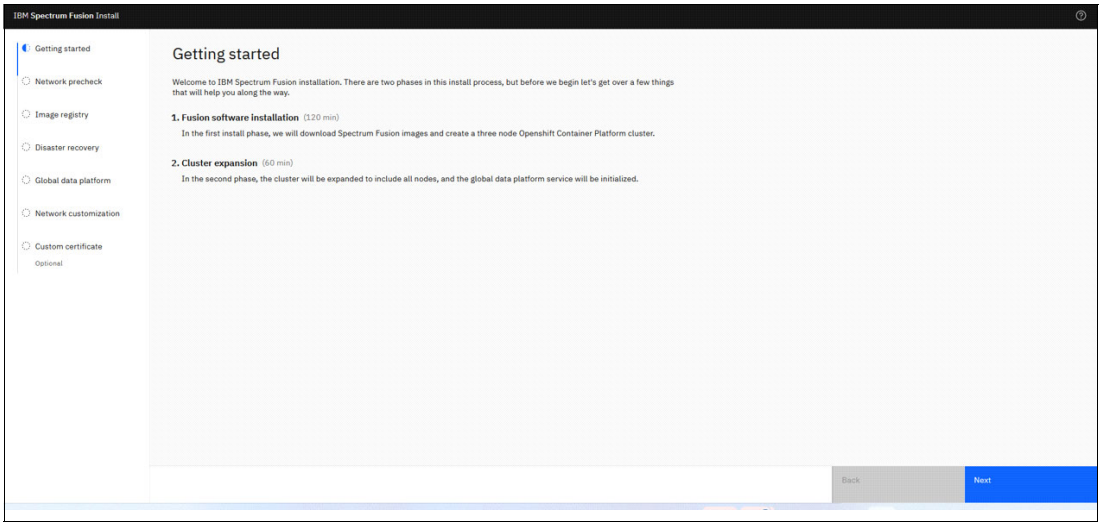

Figure 2-9 Install procedure

4. The **Network precheck** page displays all the nodes of the appliance along with the other details like MAC address, status, location, hostname, and IP address as shown in Figure 2-10.

.

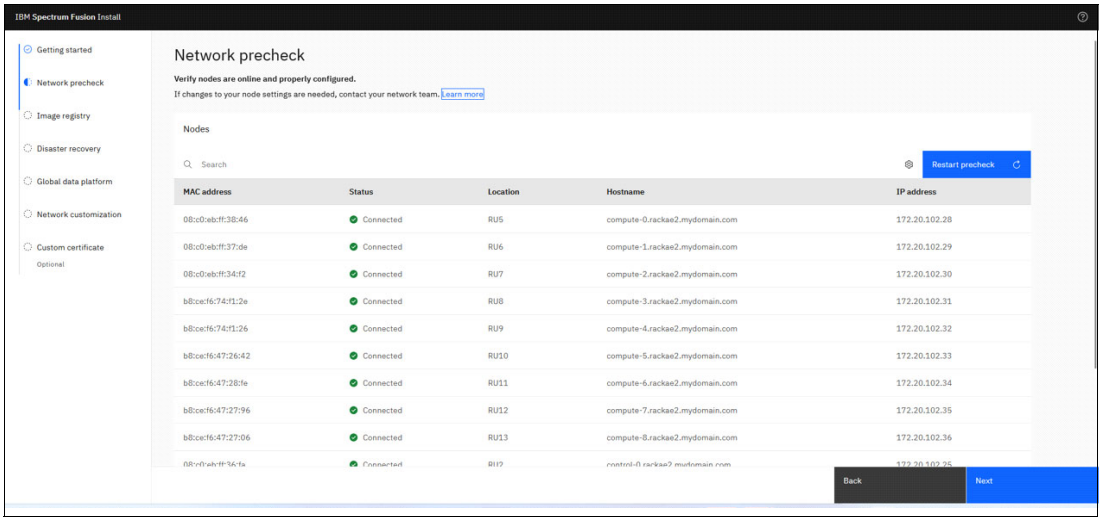

Figure 2-10 Network precheck

5. Select the image registry as per your requirements as shown in Figure 2-11. enter the details and click **Next.**

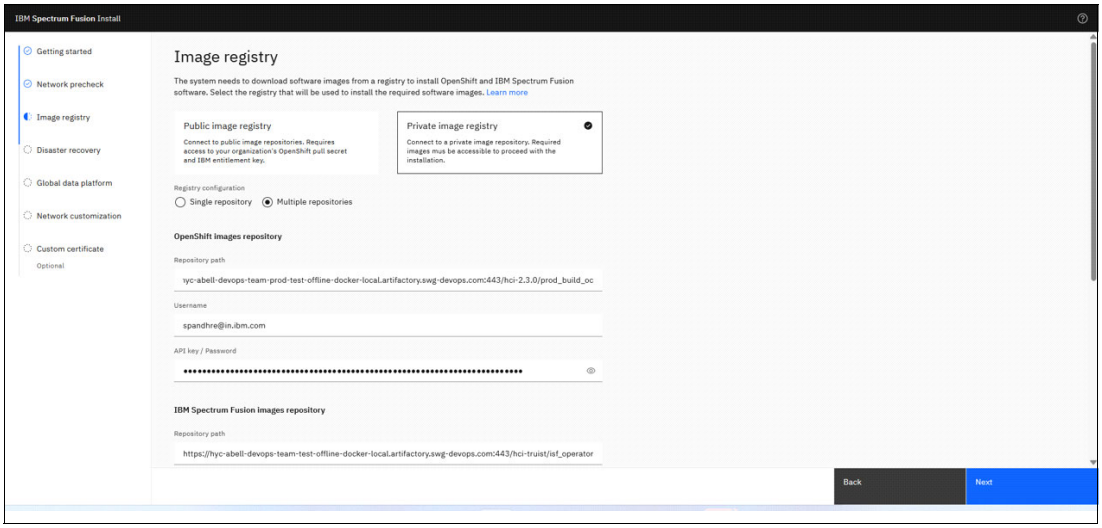

Figure 2-11 Image registry

6. In the **Disaster recovery** page as shown in Figure 2-12 on page 11, select the 2nd tile which is the first site of the Disaster Recovery pair.

   Click **Next** to continue.

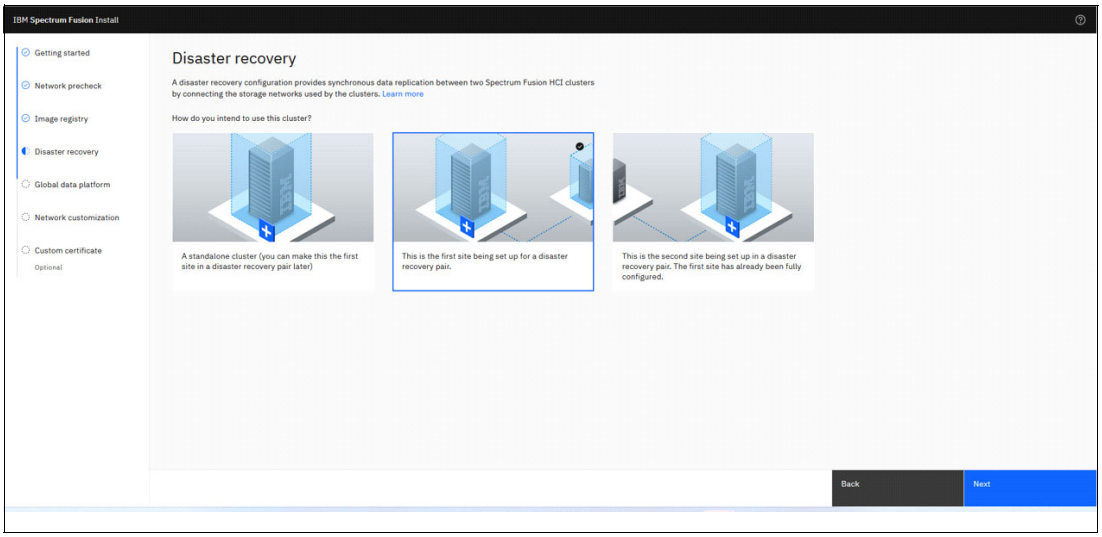

Figure 2-12 Disaster Recovery

7. On the **Global data platform** page as shown in Figure 2-13, select the appropriate building block. Click **Next**.

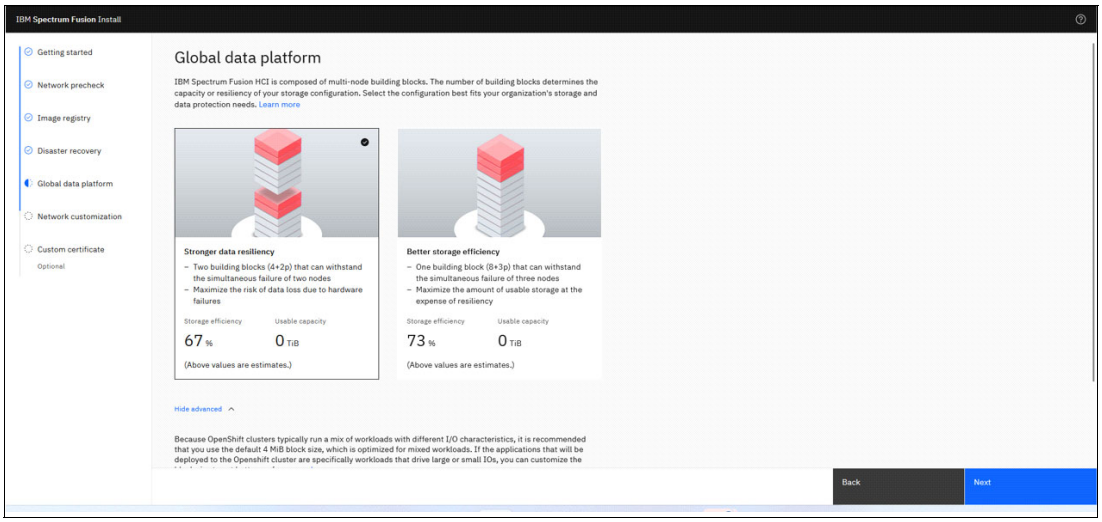

Figure 2-13 Global data platform

8. On the **Network customization** page as shown in Figure 2-14, enter the details of the OpenShift network and storage network.

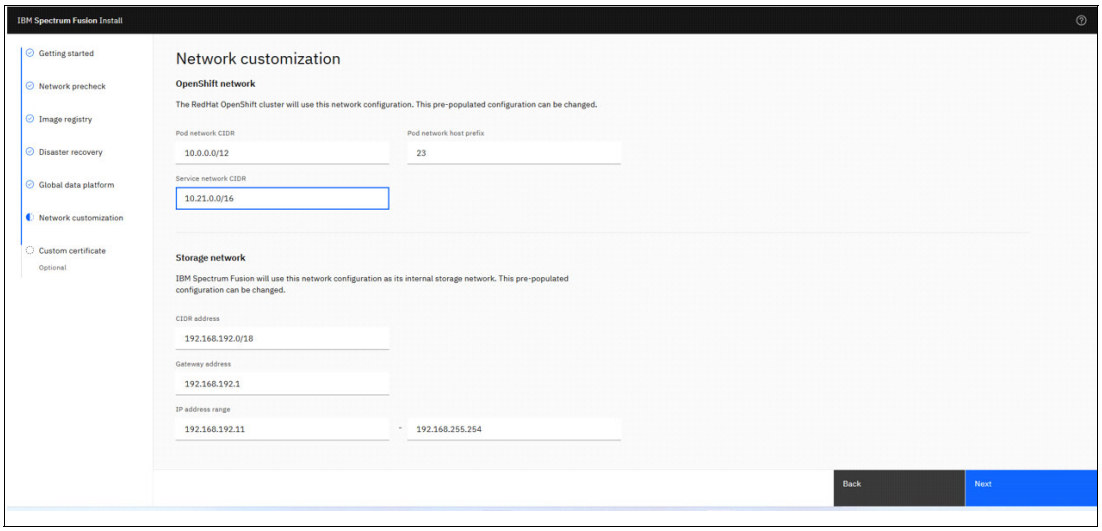

Figure 2-14 Network customization

9. On the **Custom certificate** page as shown in Figure 2-15, provide the details of your organization certificate, if any.

   Click **Finish** to start the installation.

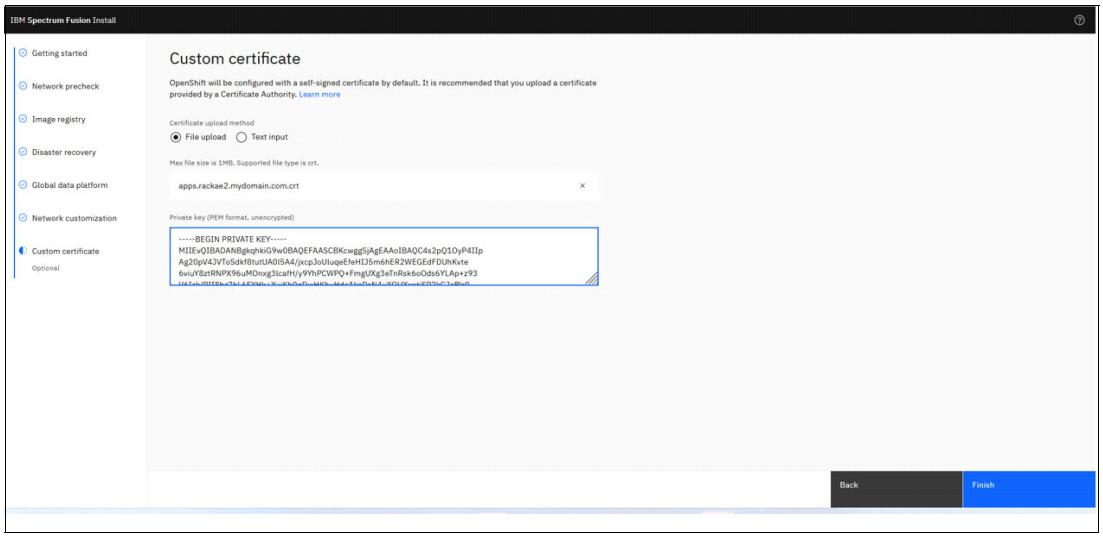

Figure 2-15 Custom certificate

10. Once the Install completes, the message is displayed as shown in Figure 2-16 on page 13.

> **Note:** Ensure to download the OpenShift Password and CoreOS Key and/or copy the credentials from the page.

11. Once the OpenShift credentials are downloaded, confirm by clicking on the checkbox.
12. Now, click on the **IBM Spectrum Fusion** button to proceed with the install.

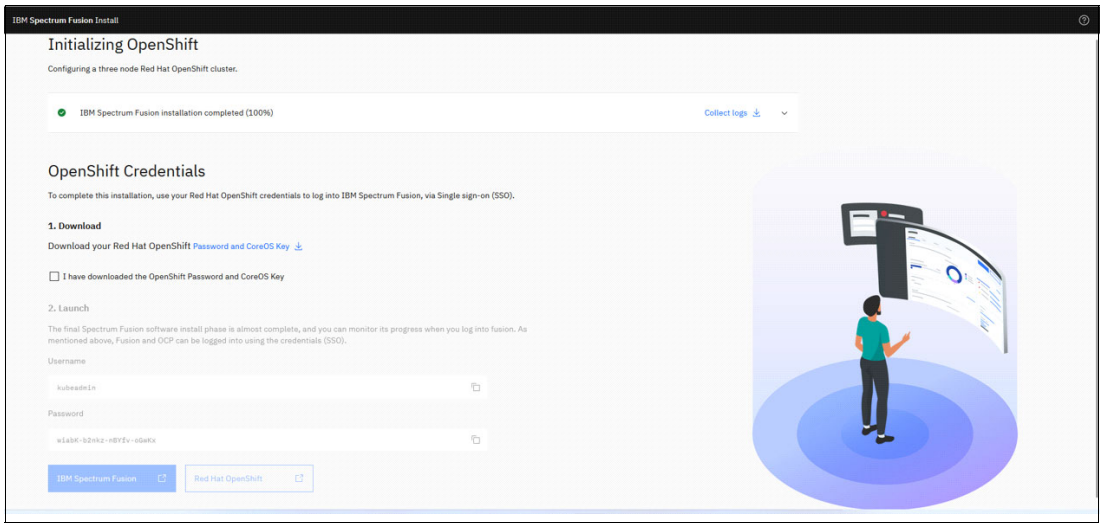

Figure 2-16 Initializing OpenShift

OpenShift configuration and Global data platform installation

This section describes the steps for the OpenShift configuration and Global data platform installation:

1. Wait for the OpenShift configuration and global data platform installation to complete successfully as shown in Figure 2-17.

Figure 2-17 Final installation process of IBM Spectrum Fusion

2. Click on the **Launch Fusion** button in Figure 2-17 to go to the IBM Spectrum Fusion **Quick start** page.

3. Now, we are ready to use IBM Spectrum Fusion as shown in Figure 2-18.

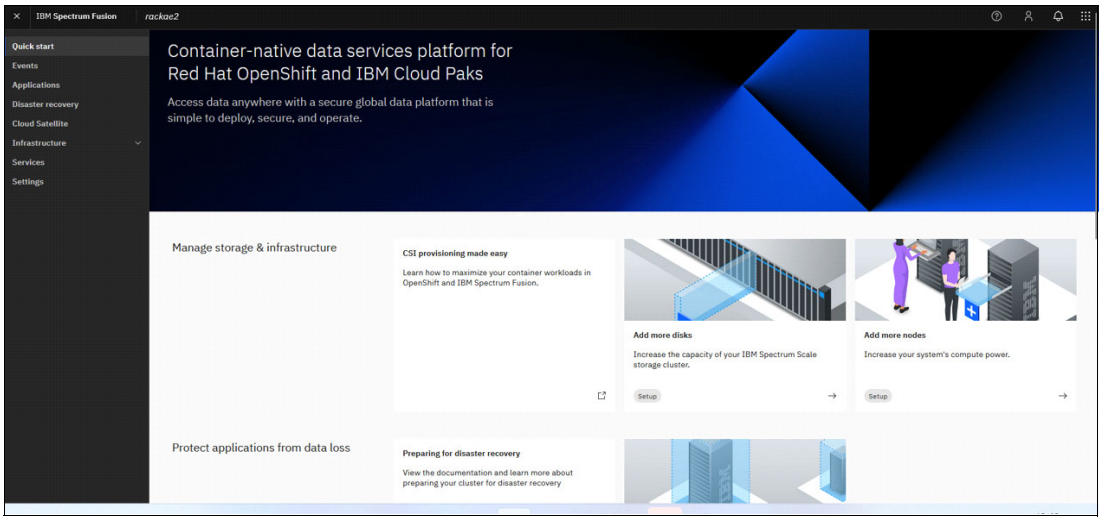

Figure 2-18 IBM Spectrum Fusion Quick start page

Disaster Recovery

On IBM Spectrum Fusion, in the left pane menu, click on **Disaster recovery** option. The **Disaster recovery** page is shown in Figure 2-19.

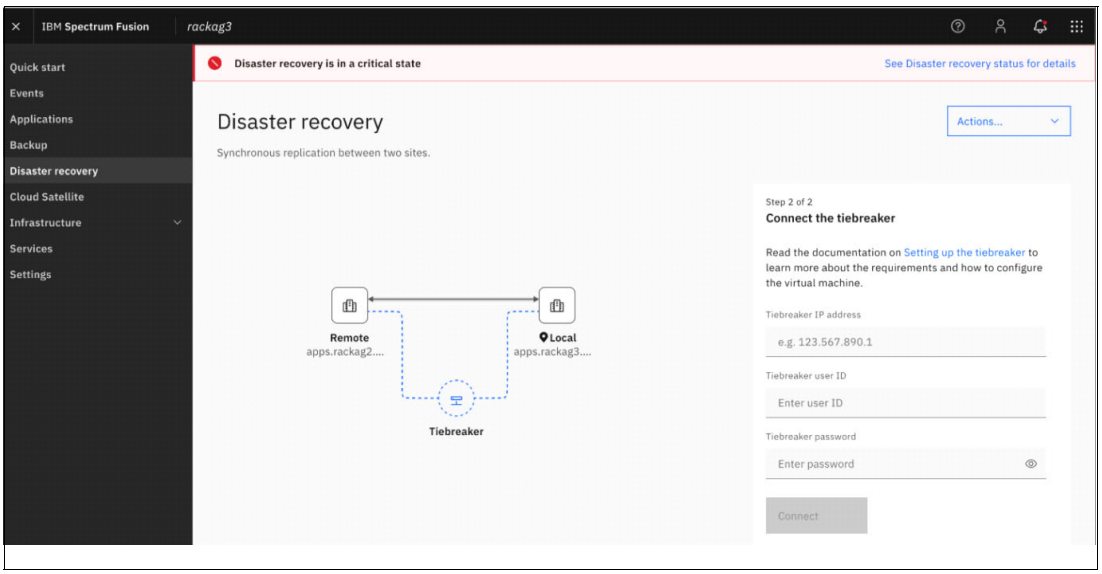

Figure 2-19 Disaster recovery option

2.1.4 Site2 Installation

This section describes the steps to install Site2, irrespective of whatever deployment type is chosen from 1.2, “Deployment Models” on page 3.

Network configuration

1. Enter the details of the Network configuration as show in Figure 2-20.
2. Click on the **Validate Network** button.

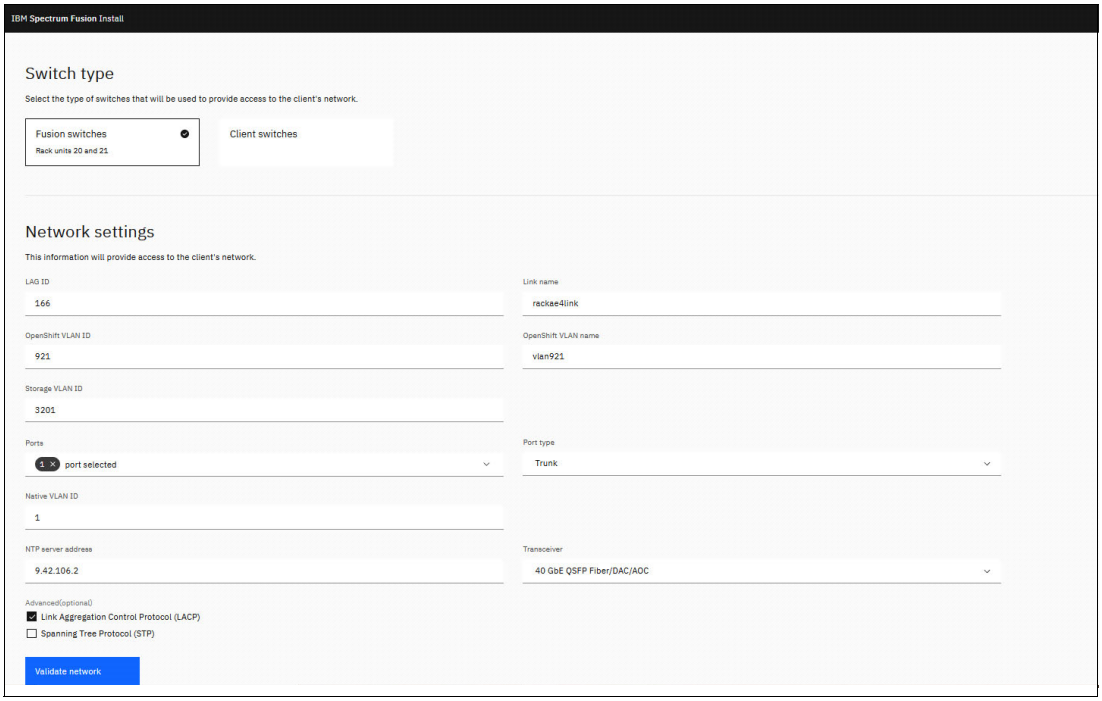

Figure 2-20 Network configuration

The Network configuration provisions Node IP, configures DHCP, and NTP as shown in Figure 2-21, Figure 2-22 on page 16, and Figure 2-23 on page 16.

Figure 2-21 Network configuration - Provisioning node IP

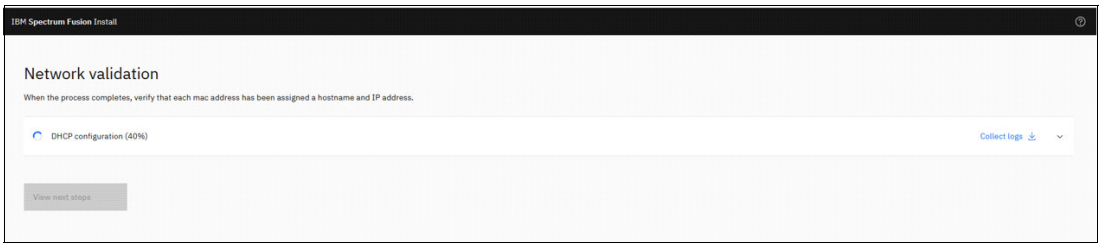

Figure 2-22 Network configuration - DHCP configuration

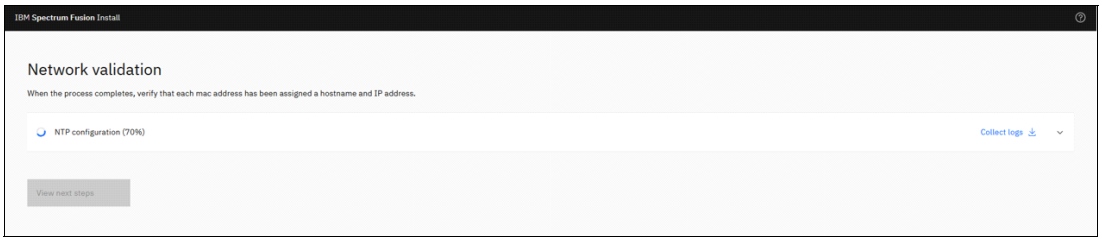

Figure 2-23 Network configuration - NTP configuration

3. The network configuration has completed successfully as shown in Figure 2-24.

Figure 2-24 Network configuration - Successful completion

4. Click the **View next steps** button. This page will further guide to proceed with installation as shown in Figure 2-25.

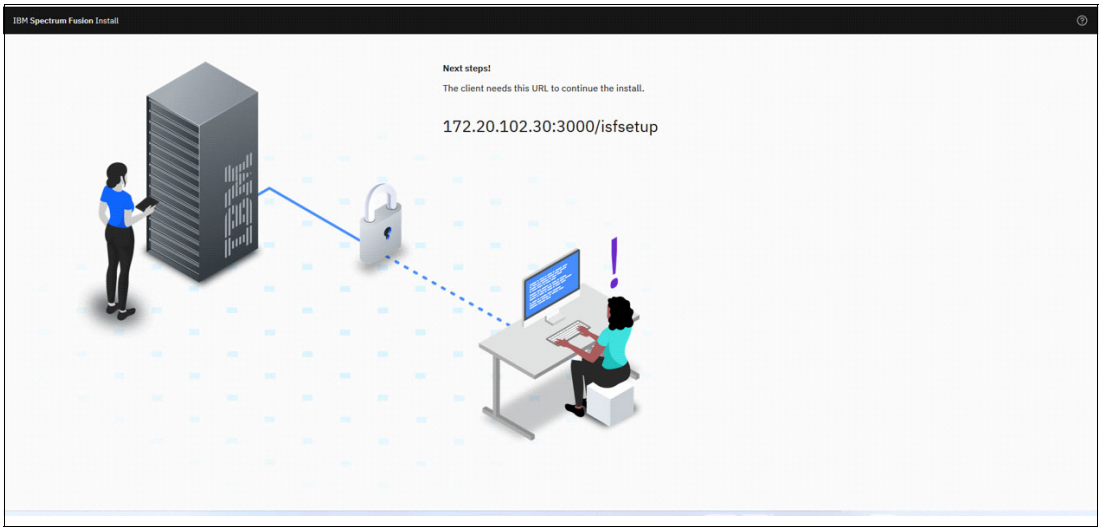

Figure 2-25 Next Steps

IBM Spectrum Fusion installer

The following steps describe the installation process for IBM Spectrum Fusion:

1. To proceed with IBM Spectrum Fusion install, use the URL as shown in Figure 2-25 on page 16 for your system.
2. You are presented with a **License agreement** page. Read the License agreement and Privacy policy. Then, accept the license and click the **Continue** button as shown in Figure 2-26.

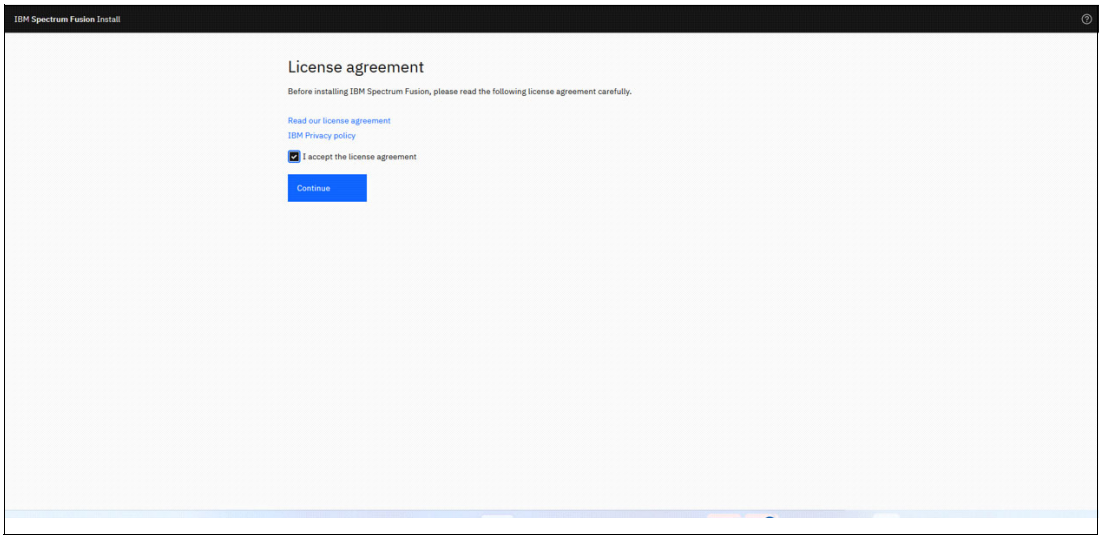

Figure 2-26 License agreement

3. The **Getting Started** page displays the procedure for Install process as shown in Figure 2-27.

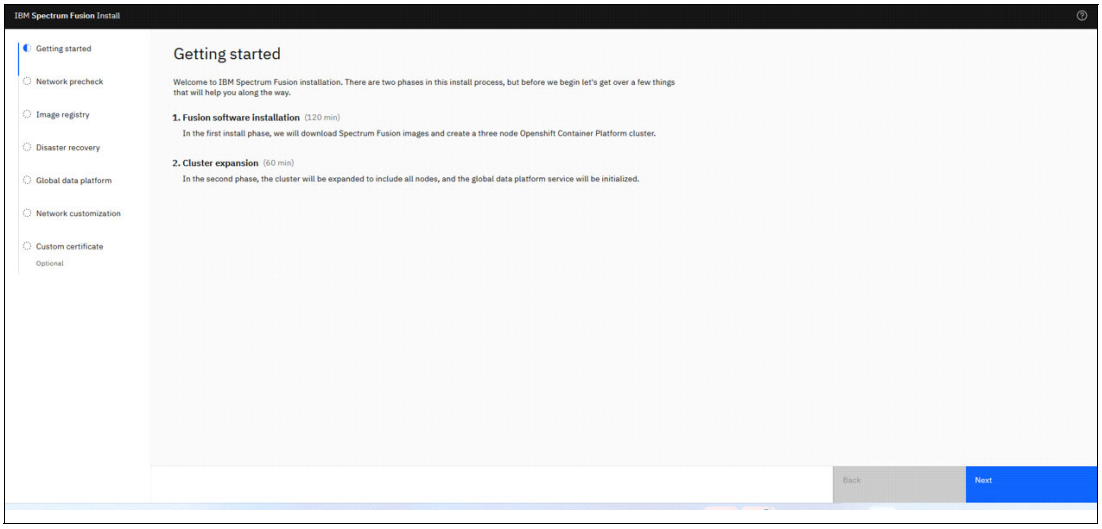

Figure 2-27 Install procedure

4. The **Network precheck** page displays all the nodes of the appliance along with the other details like MAC address, status, location, hostname, and IP address as shown Figure 2-28.

> **Note:** If any issues are indicated in the page, the recommendation is to connect with IBM to resolve the issue before proceeding further.

Click the **Next** button.

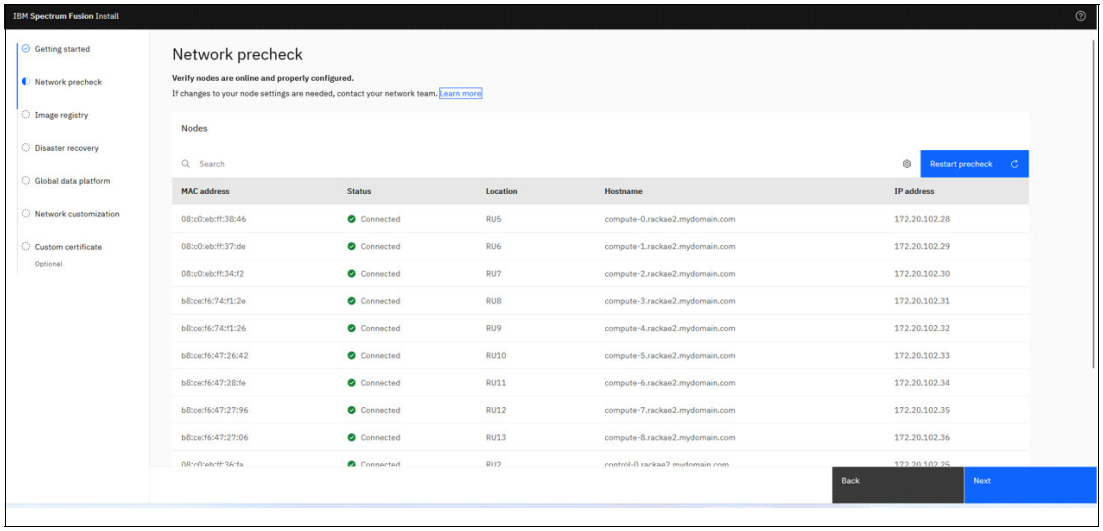

Figure 2-28 Network precheck

5. Select the image registry as per your requirements as shown in Figure 2-29. Enter the details and click the **Next** button.

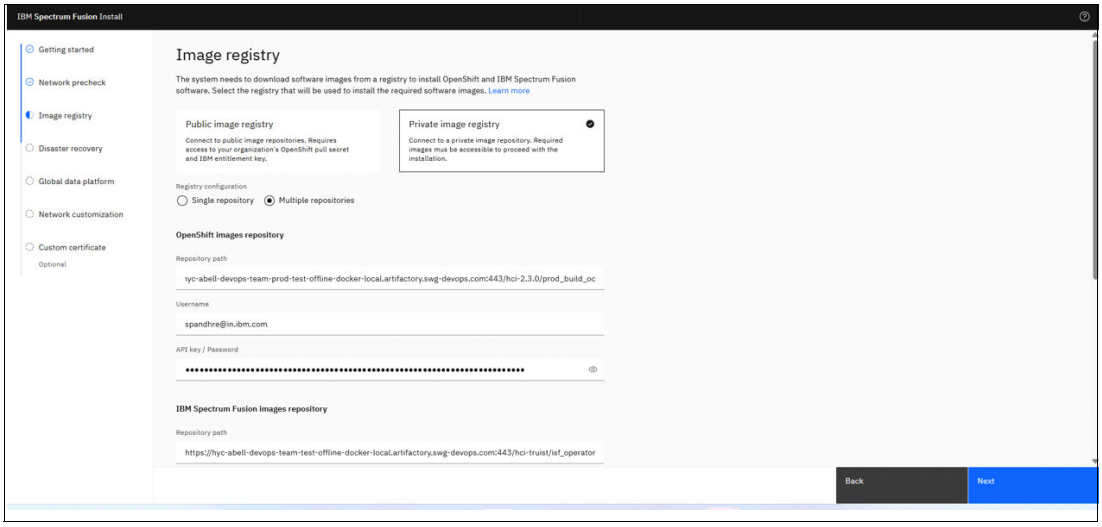

Figure 2-29 Image registry

6. To configure Disaster recovery between two sites, you need to get the connection details of Site1. The following instructions show how to get the credentials:
   a. Login to OpenShift of Site1
      i. `oc login -u kubeadmin -p <passwd of Site1> <api_url of Site1>`
   b. Switch to ibm-spectrum-fusion-ns project
      `i. oc project ibm-spectrum-fusion-ns`
   c. Get the secret name for secret `fusion-admin-controller-manager`
      i. `oc get sa fusion-admin-controller-manager -oyaml`

      From the command output, get the secret name for the `fusion-admin-controller-manager-token` secret shown here:

      Example `name: fusion-admin-controller-manager-token-9mq6b`
   d. Retrieve token from this secret mentioned in the service account
      `i. oc get secret fusion-admin-controller-manager-token-9mq6b -o yaml`
7. In the **Disaster recovery** page as shown in Figure 2-30 do the following steps:
   a. Select the 3rd tile which is the second site of the Disaster Recovery pair.
   b. Provide the Site1 details - API URL of Site1 and token retrieved in the prior step.

   Click the **Next** button to continue.

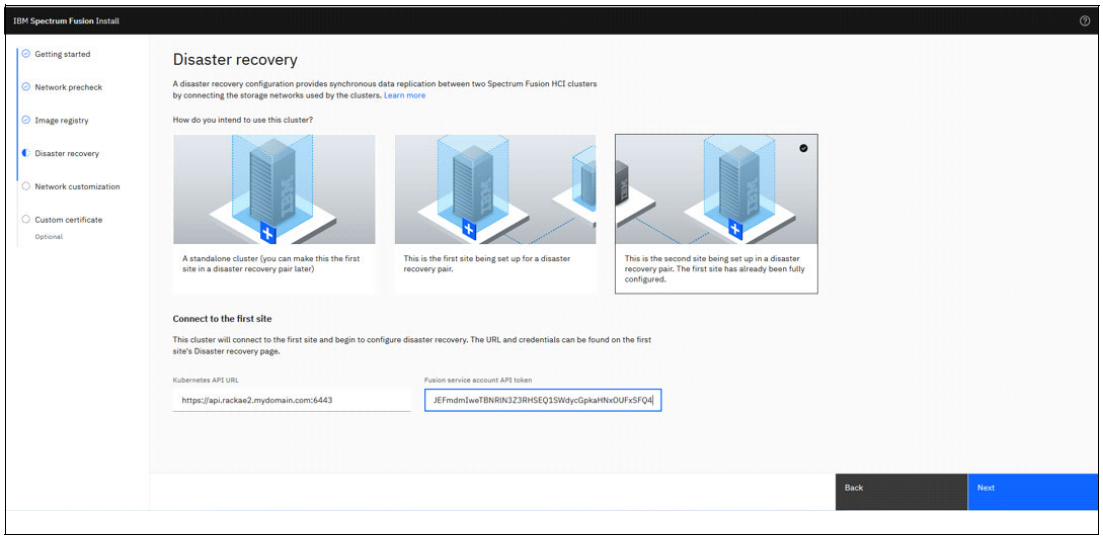

Figure 2-30 Disaster Recovery

8. On the **Global data platform** page as shown in Figure 2-31, select the appropriate building block.

   Click the **Next** button.

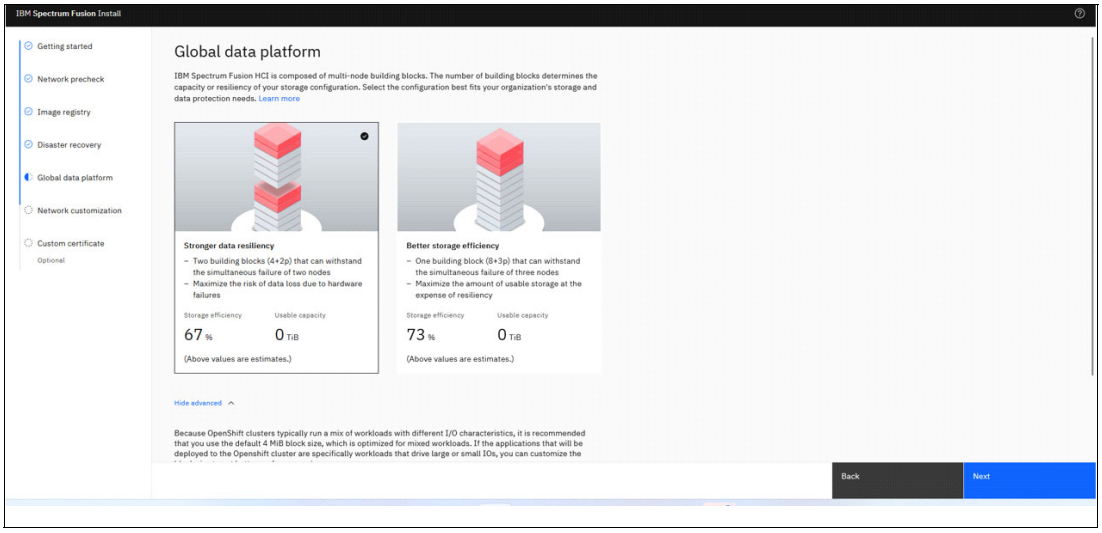

Figure 2-31 Global data platform

9. On the **Network customization** page as shown in Figure 2-32. Enter the details of the OpenShift network and storage network.

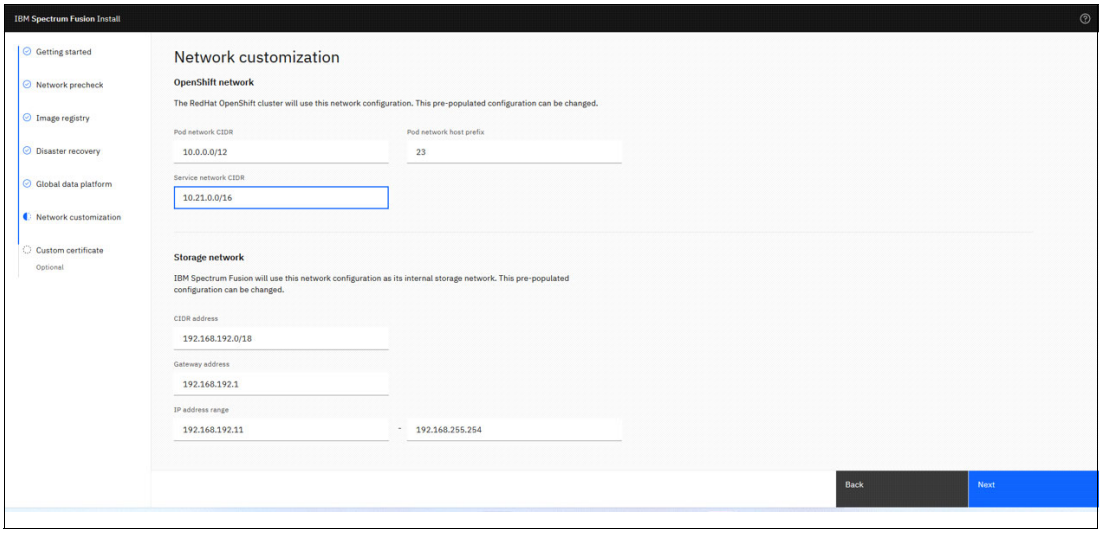

Figure 2-32 Network customization

10. On the **Custom certificate** page as shown in Figure 2-33 on page 21, provide the details of your organization certificate, if any.

    Click the **Finish** button to start the installation.

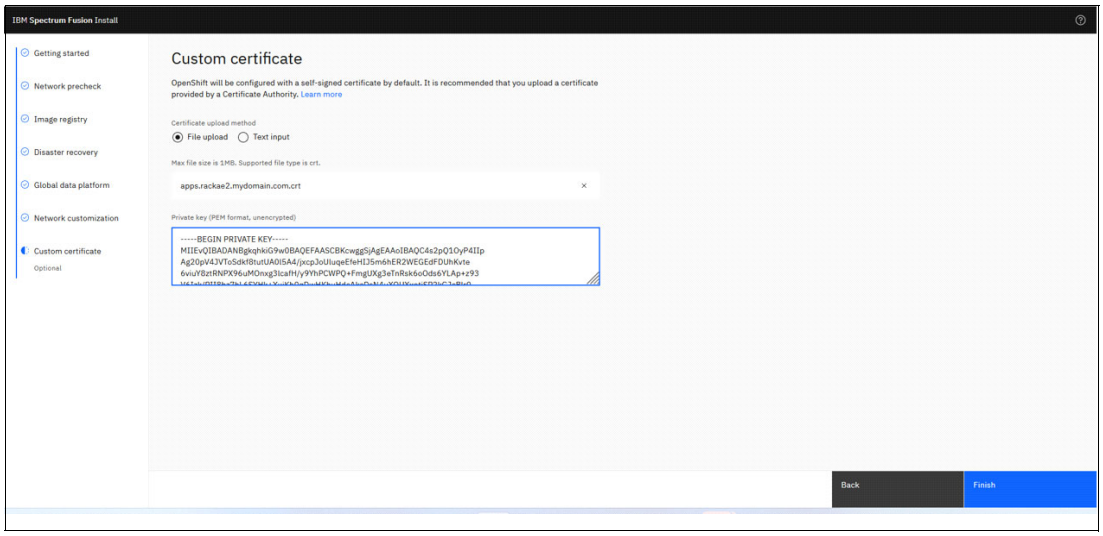

Figure 2-33 Custom certificate

11. Once the Install completes, the message is displayed as shown in Figure 2-34.

> **Note:** Ensure to download the OpenShift Password and CoreOS Key and/or copy the credentials from the page.

12. Once the OpenShift credentials are downloaded, confirm by clicking on the checkbox.
13. Now, click the **IBM Spectrum Fusion** button to proceed with the install.

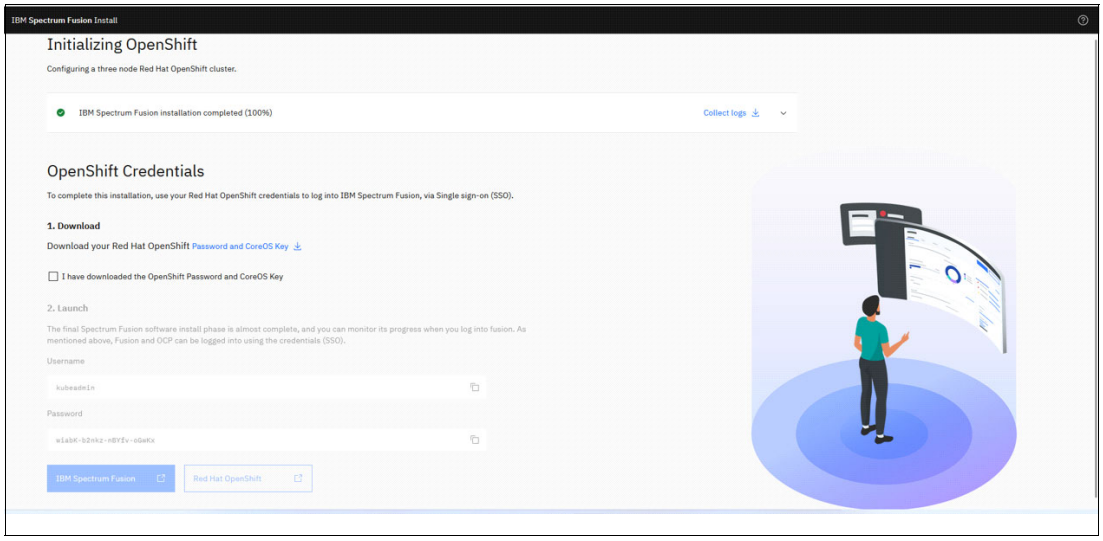

Figure 2-34 Initializing OpenShift

OpenShift configuration, Global data platform installation, and Disaster recovery connections

This section describes the steps to for the OpenShift configuration, Global data platform installation, Disaster recovery connections:

1. Wait for the OpenShift configuration, global data platform installation and disaster recovery connections to complete successfully as shown in Figure 2-35.

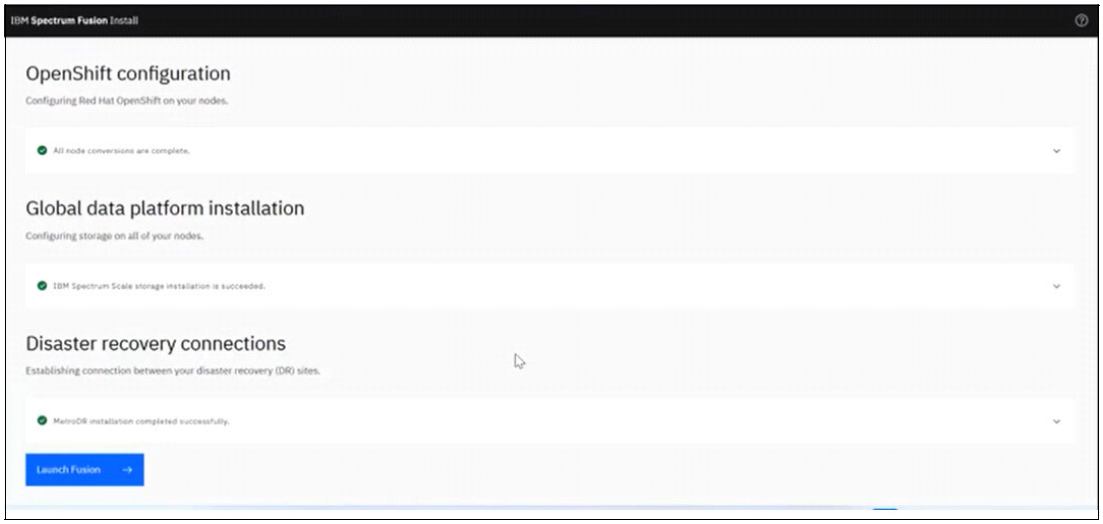

Figure 2-35 Final installation process of IBM Spectrum Fusion

2. Click the **Launch Fusion** button in Figure 2-35 to go to the **IBM Spectrum Fusion Quick start** page.
3. Now, we are ready to use IBM Spectrum Fusion as shown in Figure 2-36.

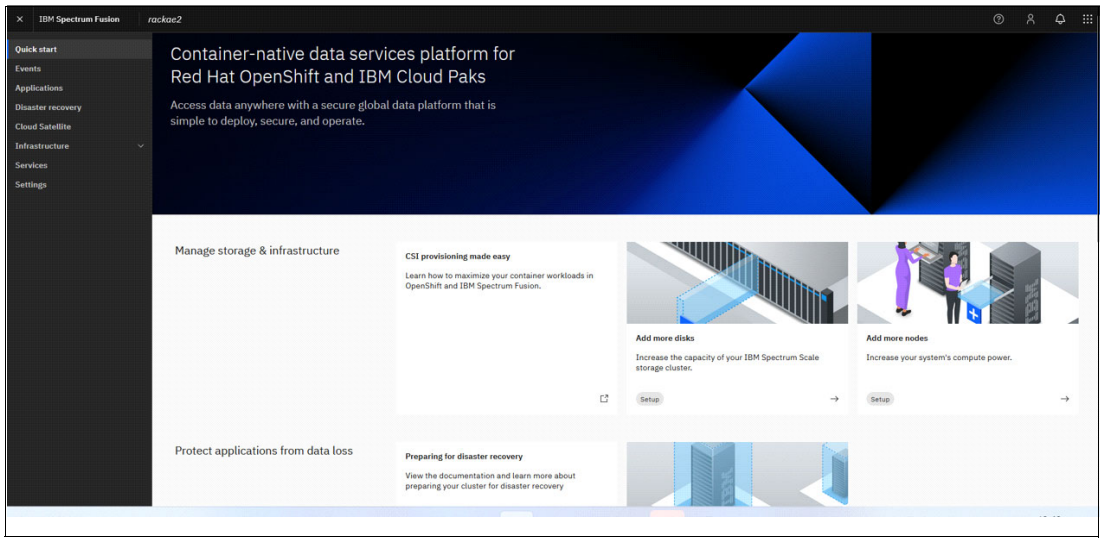

Figure 2-36 IBM Spectrum Fusion Quick start page

Disaster recovery

On the **IBM Spectrum Fusion** page, in the left pane menu, click the **Disaster recovery** option. The **Disaster recovery** page is shown as Figure 2-37.

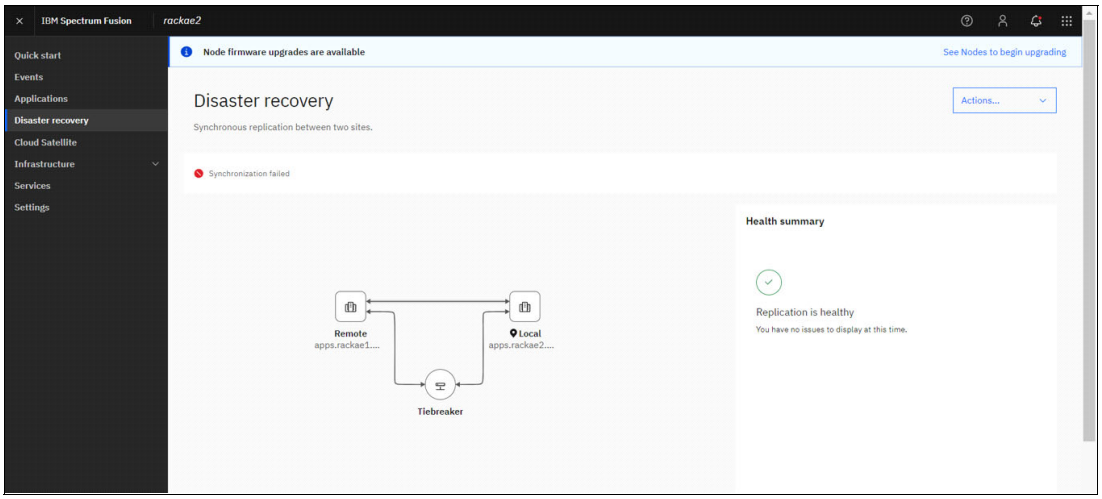

Figure 2-37 Disaster recovery

2.1.5 Tiebreaker installation

A special tiebreaker node is hosted at a third site and is used to determine which cluster is in charge of the data in the event that the network between the two clusters is severed. Configuring a Metro sync DR topology requires several network connections to be made between the two clusters and the tiebreaker.

The following are the installation requirements:

- Hardware requirements are CPU 2 cores, Memory 4G, a raw disk with less than 20 GB
- For software requirements, see https://www.ibm.com/docs/en/spectrum-scale/5.1.5?topic=gpfs-software-requirements
- For tiebreaker allowed ports, see https://www.ibm.com/docs/en/spectrum-fusion/2.4?topic=planning-firewall-requirements-spectrum-fusion-hci

Download the IBM Spectrum Scale Data Management 5.1.X.0 from IBM Entitled System Support. The following steps show the download and installation:

1. Login into the Entitled Systems Support (ESS) portal as shown in Figure 2-38.

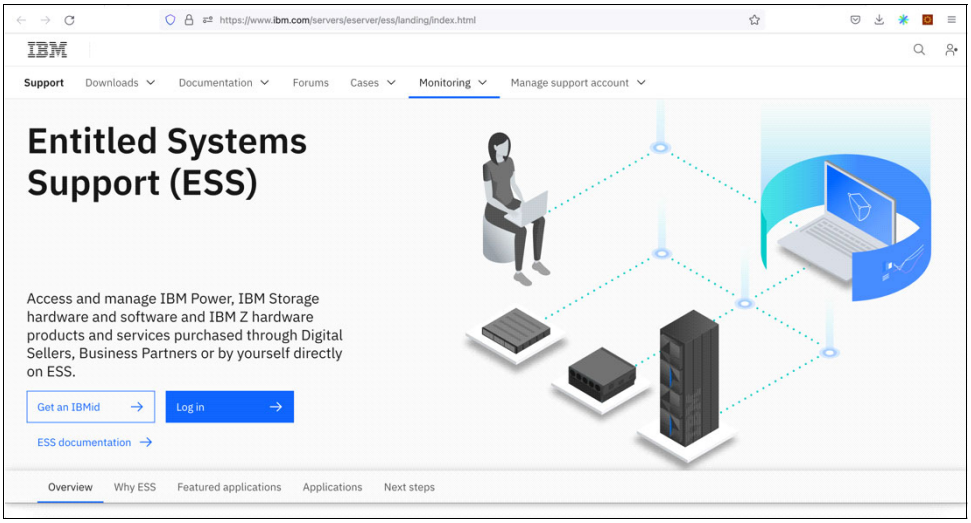

Figure 2-38 Entitled Systems Support (ESS) portal

2. Click on **My Entitled Software** link in the page as shown in Figure 2-39.

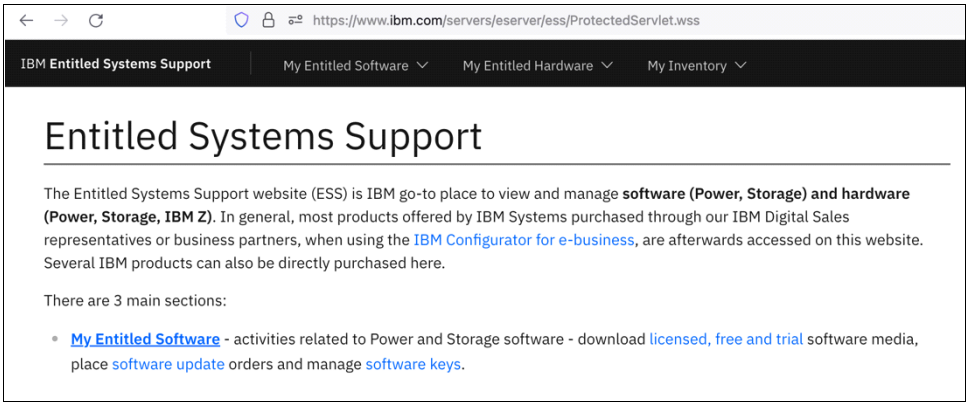

Figure 2-39 Entitled Systems Support website

3. Click on **Software Downloads** link as shown in Figure 2-40 on page 25.

This section of the ESS website focuses on activities related to **Power and Storage software products** - downloading, ordering and updating software media and keys.

You **need to be registered for one or more customer numbers** to use most of the options, as access to customer numbers is what governs everything you are entitled to do. Options related to customer number access can be found under **My Profile**.

The following activities are available within **My Entitled Software**:

- **Software Updates** - IBM i, AIX and Linux - place update orders for products you have an active software maintenance for
- **Software Downloads** - IBM i, AIX, Linux, Storage - download all your entitled software based on authorized customer numbers
- **IBM i Evaluation and NLV Download** - IBM i - download any available free, trial and secondary national language version software media
- **IBM i Software Keys** - IBM i - download keys for licensed IBM i products, manage your IBM i temporary licenses (5733-ITL) and download interim keys for supported products

Figure 2-40 Software Downloads

4. Search for the product **5771** as shown in Figure 2-41 and Figure 2-42.

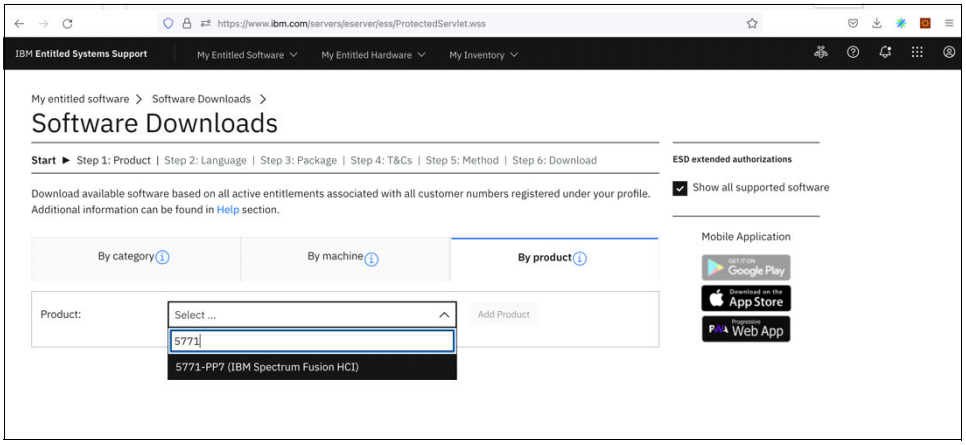

Figure 2-41 Search Product

5. Select the product **5771-PP7 (IBM Spectrum Fusion HCI)** as shown in Figure 2-42.

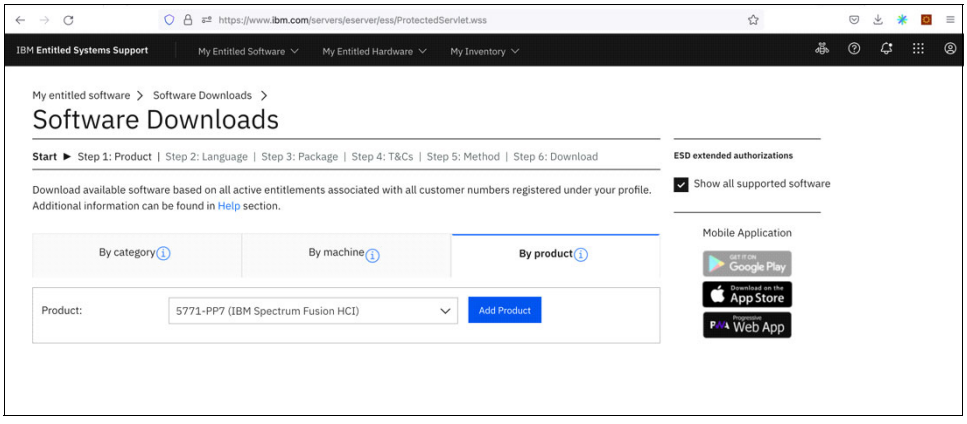

Figure 2-42 Select the product

6. The selected product is displayed as shown in Figure 2-43. Click on the **Continue** button.

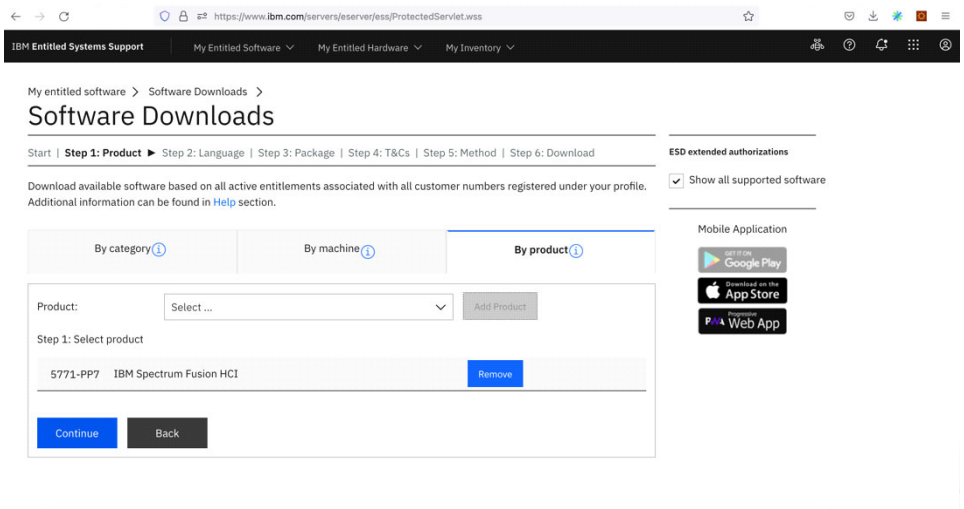

Figure 2-43 Selected product display

7. Click the checkbox as shown in Figure 2-44. Click the **Continue** button.

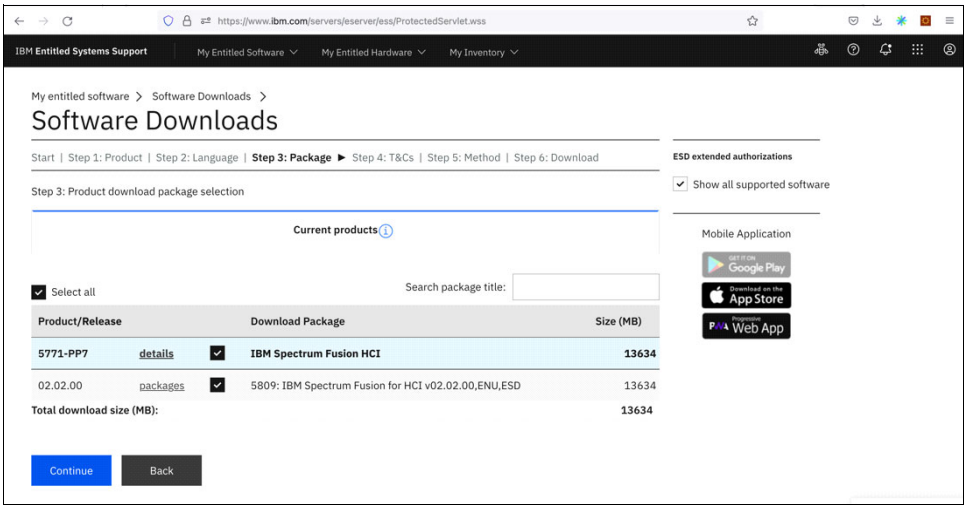

Figure 2-44 Confirm the selected products

8. Read the License Terms and click on the **I agree** button as shown in Figure 2-45 on page 27.

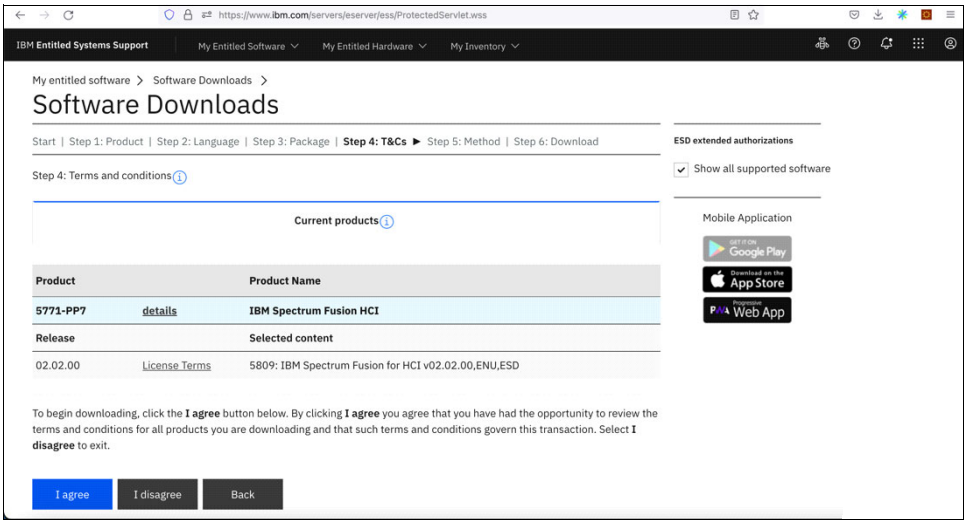

Figure 2-45 License Terms confirmation

9. Select the download method as shown in Figure 2-46 and click on the **Continue** button.

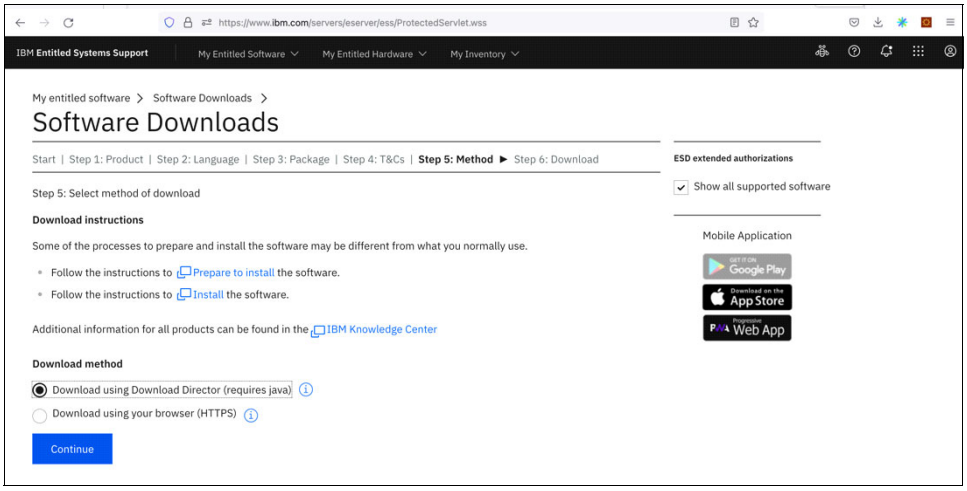

Figure 2-46 Download method

10. Review the download details as shown in Figure 2-47 and click on **Download now** button.

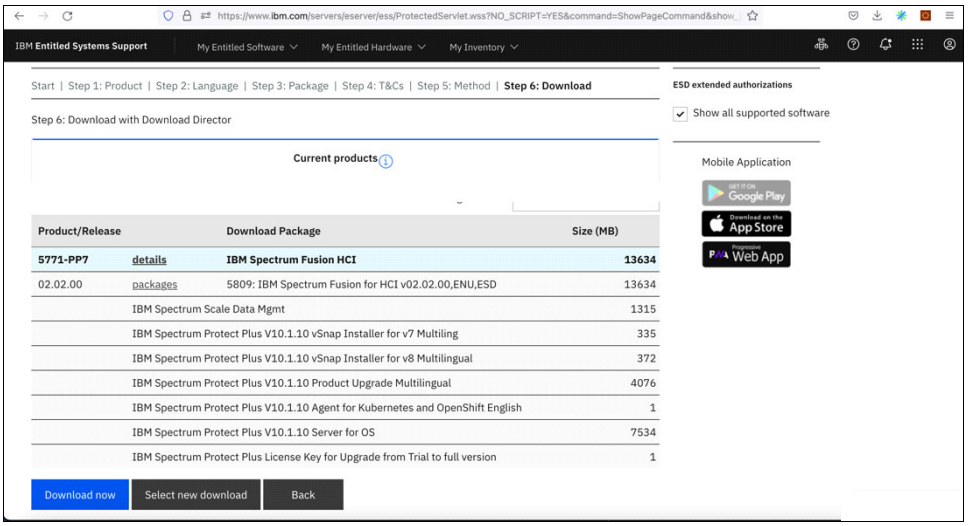

Figure 2-47 Start download

11. Once the download is complete, navigate to the download folder. List the contents of the directory. You should be able to see the files as listed in Example 2-1.

Example 2-1 Downloaded files

```
[root@metroDRAG2-tiebreaker-2 tmp]# ls
Spectrum_Scale_Data_Management-5.1.6.0-x86_64-Linux-install.sh
systemd-private-552311a87a9244faafdc49ef361fcbee-chronyd.service-bUy7QJ
```

12. Change the permissions of the file as shown in Example 2-2.

Example 2-2 Permission change

```
[root@metroDRAG2-tiebreaker-2 tmp]# chmod +x Spectrum_Scale_Data_Management-
5.1.6.0-x86_64-Linux-install.sh
```

13. Start the installation as shown in Figure 2-48 on page 29 and Figure 2-49 on page 30. Input the details as requested.

```
[root@metroDRAG2-tiebreaker-2 tmp]# ./Spectrum_Scale_Data_Management-5.1.6.0-x86_64-Linux-install.sh
Extracting License Acceptance Process Tool to /usr/lpp/mmfs/5.1.6.0 ...
tail -n +660 ./Spectrum_Scale_Data_Management-5.1.6.0-x86_64-Linux-install.sh | tar -C /usr/lpp/mmfs/5.1.6.0 -xvz --exclude=installer
--exclude=*_rpms --exclude=*_debs --exclude=*rpm  --exclude=*tgz --exclude=*deb --exclude=*tools* 1> /dev/null

Installing JRE ...

If directory /usr/lpp/mmfs/5.1.6.0 has been created or was previously created during another extraction,
.rpm, .deb, and repository related files in it (if there were) will be removed to avoid conflicts with the ones being extracted.

tail -n +660 ./Spectrum_Scale_Data_Management-5.1.6.0-x86_64-Linux-install.sh | tar -C /usr/lpp/mmfs/5.1.6.0 --wildcards -xvz
ibm-java*tgz 1> /dev/null
tar -C /usr/lpp/mmfs/5.1.6.0/ -xzf /usr/lpp/mmfs/5.1.6.0/ibm-java*tgz
Defaulting to --text-only mode.

Invoking License Acceptance Process Tool ...
/usr/lpp/mmfs/5.1.6.0/ibm-java-x86_64-80/jre/bin/java -cp /usr/lpp/mmfs/5.1.6.0/LAP_HOME/LAPApp.jar com.ibm.lex.lapapp.LAP
-l /usr/lpp/mmfs/5.1.6.0/LA_HOME -m /usr/lpp/mmfs/5.1.6.0 -s /usr/lpp/mmfs/5.1.6.0  -text_only

LICENSE INFORMATION

The Programs listed below are licensed under the following
License Information terms and conditions in addition to the
Program license terms previously agreed to by Client and
IBM. If Client does not have previously agreed to license
terms in effect for the Program, the International Program
License Agreement (i125-3301-15) applies.

Program Name (Program Number):
IBM Spectrum Scale Data Management Edition 5.1.6.0 (5737-
F34)
IBM Spectrum Scale Data Management Edition 5.1.6.0 (5641-
DM1)

Press Enter to continue viewing the license agreement, or
enter "1" to accept the agreement, "2" to decline it, "3"
to print it, "4" to read non-IBM terms, or "99" to go back
to the previous screen.
1

License Agreement Terms accepted.

Extracting Product RPMs to /usr/lpp/mmfs/5.1.6.0 ...
tail -n +660 ./Spectrum_Scale_Data_Management-5.1.6.0-x86_64-Linux-install.sh | tar -C /usr/lpp/mmfs/5.1.6.0 ...

ganesha_rpms/sles15 gpfs_rpms/rhel7 gpfs_rpms/rhel8 gpfs_rpms/rhel9 gpfs_rpms/sles15 object_rpms/rhel8 smb_rpms/rhel7
cloudkit gpfs_debs gpfs_rpms manifest 1> /dev/null
   - Public_Keys
   - ansible-toolkit
   - cloudkit/dependencies
   - ganesha_debs/ubuntu/ubuntu20
   - ganesha_debs/ubuntu/ubuntu22
   - gpfs_debs/ubuntu/ubuntu20
   - gpfs_debs/ubuntu/ubuntu22
   - hdfs_rpms/rhel/hdfs_3.1.1.x
   - hdfs_rpms/rhel/hdfs_3.2.2.x
   - hdfs_rpms/rhel/hdfs_3.3.x
   - smb_debs/ubuntu/ubuntu20
   - smb_debs/ubuntu/ubuntu22
   - zimon_debs/ubuntu/ubuntu20
   - zimon_debs/ubuntu/ubuntu22
   - ganesha_rpms/rhel7

   - zimon_rpms/sles15
   - cloudkit
   - gpfs_debs
   - gpfs_rpms
   - manifest

Removing License Acceptance Process Tool from /usr/lpp/mmfs/5.1.6.0 ...
rm -rf  /usr/lpp/mmfs/5.1.6.0/LAP_HOME /usr/lpp/mmfs/5.1.6.0/LA_HOME

Removing JRE from /usr/lpp/mmfs/5.1.6.0 ...
rm -rf /usr/lpp/mmfs/5.1.6.0/ibm-java*tgz
```

Figure 2-48 Install

Figure 2-49 Install (continuation)

14. Navigate to the ansible toolkit directory as shown in Example 2-3.

Example 2-3 Ansible toolkit directory

```
[root@metroDRAG2-tiebreaker-2 tmp]# cd /usr/lpp/mmfs/5.1.6.0/ansible-toolkit
```

15. View the contents of the directory as shown in Example 2-4. Verify if you can see a file named spectrumscale.

Example 2-4 Directory contents of ansible-toolkit directory

```
[root@metroDRAG2-tiebreaker-2 tmp]# cd /usr/lpp/mmfs/5.1.6.0/ansible-toolkit
[root@metroDRAG2-tiebreaker-2 ansible-toolkit]# ls
README       ansible  cli            documentation  externallibs       license
spectrumscale
```

16. Install the toolkit as shown in Example 2-5 using the following command:

```
./spectrumscale setup -s <IP of the scale cluster master node>
```

Example 2-5 Ansible-toolkit install

```
[root@metroDRAG2-tiebreaker-2 ansible-toolkit]# ./spectrumscale setup -s
10.11.123.47
[ INFO  ] Installing prerequisites for install node
[ INFO  ] Installing Ansible version 2.9.15.
[ INFO  ] Install Toolkit setup type is set to Spectrum Scale (default). If an ESS
is in the cluster, run this command to set ESS mode: ./spectrumscale setup -s
server_ip -st ess
[ INFO  ] Your ansible controller node has been configured to use the IP
10.11.123.47 to communicate with other nodes.
[ INFO  ] Port 10080 will be used for package distribution.
[ INFO  ] SUCCESS
[ INFO  ] Tip : Designate protocol, nsd and admin nodes in your environment to use
during install:./spectrumscale -v node add <node> -p  -a -n
[root@metroDRAG2-tiebreaker-2 ansible-toolkit]#
```

17. Get the secrets of the following Spectrum Scale keys as shown in Example 2-6.

Example 2-6 Secrets of IBM Spectrum Scale

```
oc get secret ibm-spectrum-scale-core-ssh-key-secret -n ibm-spectrum-scale -ojsonpath="{.data.ssh-authorizedkeys}"
oc get secret ibm-spectrum-scale-core-ssh-key-secret -n ibm-spectrum-scale -ojsonpath="{.data.ssh-privatekey}"
oc get secret ibm-spectrum-scale-core-ssh-key-secret -n ibm-spectrum-scale -ojsonpath="{.data.ssh-publickey}"
```

18. Validate the secrets from Step 17 are present in the file tiebreaker_nodedefinition.json as shown in Example 2-7.

Example 2-7 Secret validation

```
[root@metroDRAG2-tiebreaker-2 ansible-toolkit]# cd /usr/lpp/mmfs/5.1.6.0/ansible- toolkit/ansible/vars
[root@metroDRAG2-tiebreaker-2 vars]# cat tiebreaker_nodedefinition.json
```

19. Copy the yaml files as shown in Example 2-8.

Example 2-8 Copy the yaml files

```
[root@metroDRAG2-tiebreaker-2 vars]# cp
/usr/lpp/mmfs/5.1.6.0/ansible-toolkit/ansible/sample/playbook_tiebreakernode_install.yml
/usr/lpp/mmfs/5.1.6.0/ansible-toolkit/ansible/sample/set_json_variables_tb.yml
/usr/lpp/mmfs/5.1.6.0/ansible-toolkit/ansible/
```

20. Navigate to the ansible directory to run the playbook as shown in Example 2-9.

Example 2-9 Run the ansible playbook

```
[root@metroDRAG2-tiebreaker-2 vars]# cd  /usr/lpp/mmfs/5.1.6.0/ansible-toolkit/ansible/
[root@metroDRAG2-tiebreaker-2 ansible]# ansible-playbook playbook_tiebreakernode_install.yml
```

21. Ensure the playbook runs successfully as shown in Figure 2-50.

```
[root@metroDRAG2-tiebreaker-2 ansible]# ansible-playbook playbook_tiebreakernode_install.yml
[WARNING]: provided hosts list is empty, only localhost is available. Note that the implicit localhost does not match 'all'

PLAY [localhost] ***************************************************************************************************************

TASK [Gathering Facts] *********************************************************************************************************
ok: [localhost]

TASK [Read all intermediate output from Resource Details] **********************************************************************
ok: [localhost]

TASK [Check valid json file] ***************************************************************************************************
ok: [localhost] => {
    "changed": false,
    "msg": "All assertions passed"
}
…

TASK [prepare | Authorize all SSH keys] ****************************************************************************************
changed: [10.11.123.47] => (item=10.11.123.47)

TASK [shell] *******************************************************************************************************************
skipping: [10.11.123.47]

TASK [shell] *******************************************************************************************************************
skipping: [10.11.123.47]

TASK [prepare | Change the Port 12345 in ssh_config] ***************************************************************************
changed: [10.11.123.47]

TASK [prepare | Change the Port 12345 in sshd_config] **************************************************************************
changed: [10.11.123.47]

TASK [prepare | Restart the sshd service] **************************************************************************************
changed: [10.11.123.47]

TASK [Create a file in tiebreaker to avoid mounting any filesystem] ************************************************************
changed: [10.11.123.47] => (item=10.11.123.47)

PLAY RECAP *********************************************************************************************************************
```

Figure 2-50 Successful completion of ansible playbook

22. Get the secret by using the command as given in Example 2-10.

Example 2-10 Get the secret

```
[root@metroDRAG2-tiebreaker-2 ansible]# echo "/dev/vdb" | base64
```

23. Patch the secret on any of the sites of the IBM Spectrum Fusion `metrodr` pair as shown in Example 2-11.

Example 2-11 Secret patched on the IBM Spectrum Fusion rack

```
% oc patch secret isf-metrodr-config-secret -n ibm-spectrum-fusion-ns -p
'{"data":{"TieBreakerDevice":"secret from Step 22"}}'
secret/isf-metrodr-config-secret patched
```

Tiebreaker configuration from IBM Spectrum Fusion GUI

Now that we finished the install steps we are ready to finish the Metro sync DR configuration from the IBM Spectrum Fusion GUI. Follow these steps:

1. Login to the IBM Spectrum Fusion Gui.
2. Go to Disaster Recovery page.
   - Add Tiebreaker IPs and Credentials on the **Disaster recovery** page as shown in Figure 2-51 on page 33.
   - Click on the **Connect** button.

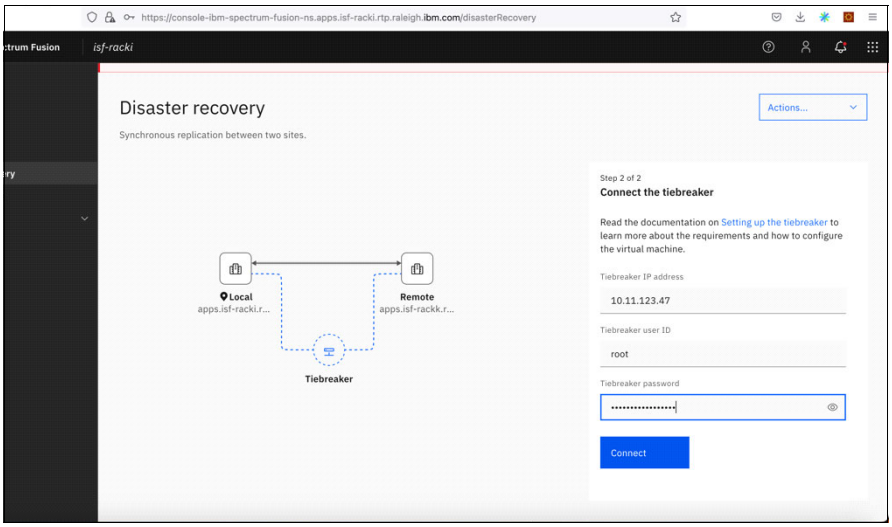

Figure 2-51 Add tiebreaker IPs and Credentials

Verify IBM Spectrum Scale status with **`mmgetstate -a`** from OCP shown in Example 2-12.

Example 2-12 Verify IBM Spectrum Scale with mmgetstate -a command

```
sh-4.4# mmgetstate -a

 Node number  Node name                                        GPFS state
------------------------------------------------------------------------
           1  control-0.daemon.ibm-spectrum-scale.stg.rackag2  active
           2  control-1.daemon.ibm-spectrum-scale.stg.rackag2  active
           3  control-2.daemon.ibm-spectrum-scale.stg.rackag2  active
           4  compute-1.daemon.ibm-spectrum-scale.stg.rackag2  active
           5  compute-2.daemon.ibm-spectrum-scale.stg.rackag2  active
           6  compute-0.daemon.ibm-spectrum-scale.stg.rackag2  active
           7  compute-0.daemon.ibm-spectrum-scale.stg.rackag3  active
           8  compute-1.daemon.ibm-spectrum-scale.stg.rackag3  active
           9  compute-2.daemon.ibm-spectrum-scale.stg.rackag3  active
          10  control-0.daemon.ibm-spectrum-scale.stg.rackag3  active
          11  control-1.daemon.ibm-spectrum-scale.stg.rackag3  active
          12  control-2.daemon.ibm-spectrum-scale.stg.rackag3  active
          13  gpfs-tiebreaker                                  active
sh-4.4#
```

3

Application failover and failback

This chapter describes the steps to set up and demonstrate the failover and failback of the WordPress application between a local site and remote site with IBM Spectrum Fusion HCI.

3.1 Application Failover/Failback between local site and remote site

In case of disaster or maintenance work on any of the sites, you need to setup the applications for recovery. IBM Spectrum Fusion provides a simple method for setting up the application(s) for disaster recovery (DR).

3.1.1 Pre-requisites for application failover/failback

Before the application(s) is setup for disaster recovery, complete the following steps:

1. Ensure the health status of the Disaster Recovery is “healthy” as shown in Figure 3-1.

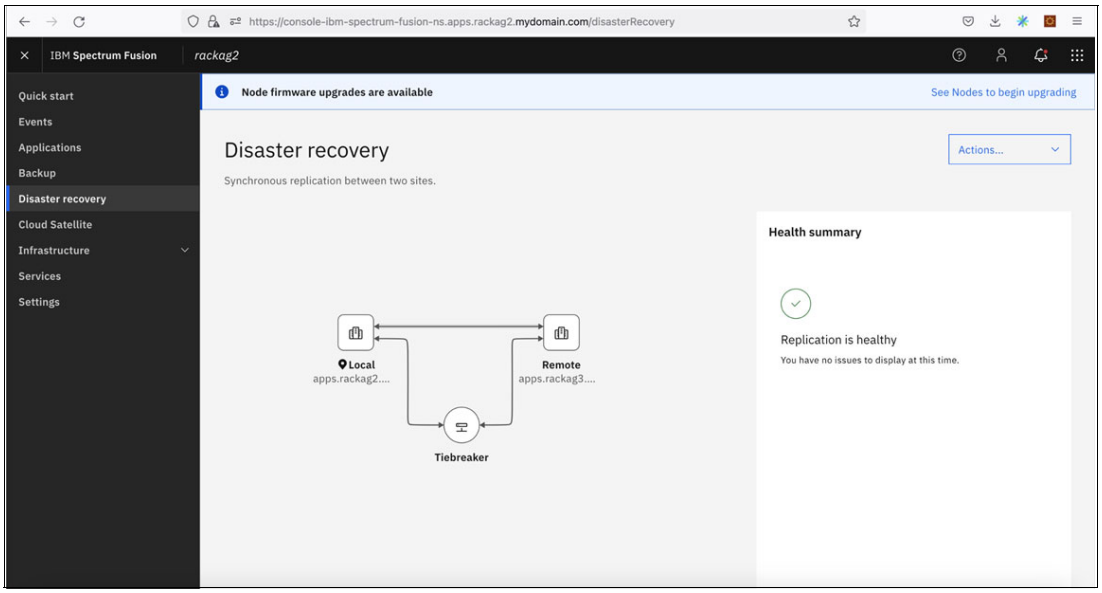

Figure 3-1 Disaster Recovery user interface

2. The application is deployed and displays on the **Applications** page in the **Local** tab as shown in Figure 3-2 on page 37.

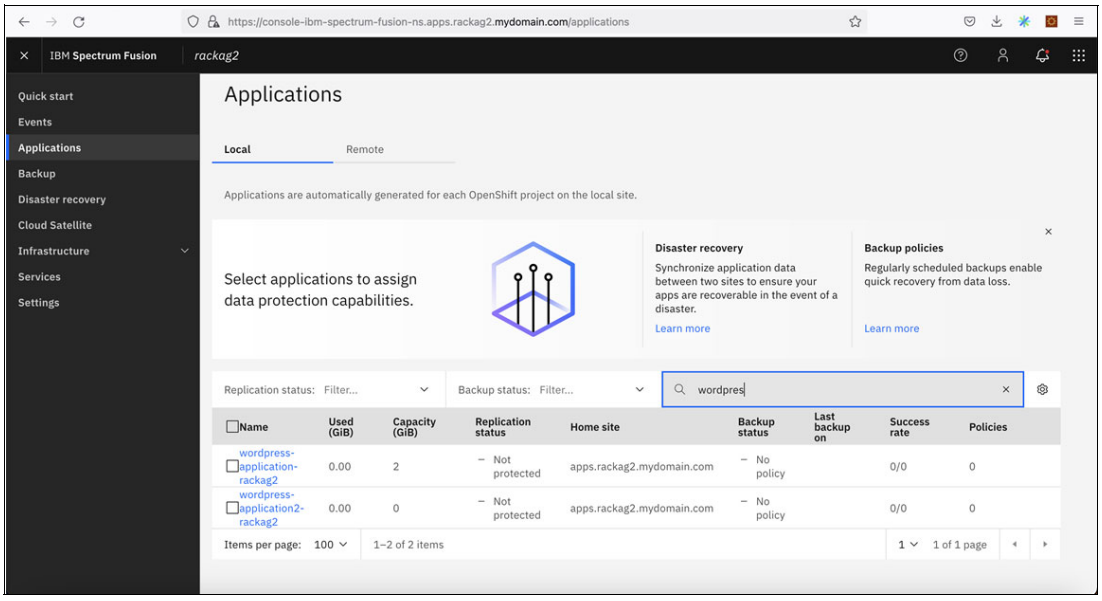

Figure 3-2 User deployed applications showing up on Applications page

3.1.2 Setting up applications for disaster recovery.

You can use the IBM Spectrum Fusion user interface to setup one or multiple applications for DR. There are multiple ways to setup one or multiple applications. This section describes the methods and steps setup the application(s) for disaster recovery.

Setting up a single application for DR

This section describes several methods of setting up a single application for DR.

Method1: Applications user interface.

Here are the steps for setting up the DR Method1 using the application user interface:

1. Go to the **Applications** page shown in Figure 3-3.
2. On the **Local** tab, you can view the applications.
3. For the application you wish to enroll for disaster recovery, go to the end of that row and click on the three dots to open the menu.
4. Click on **Add disaster recovery**.

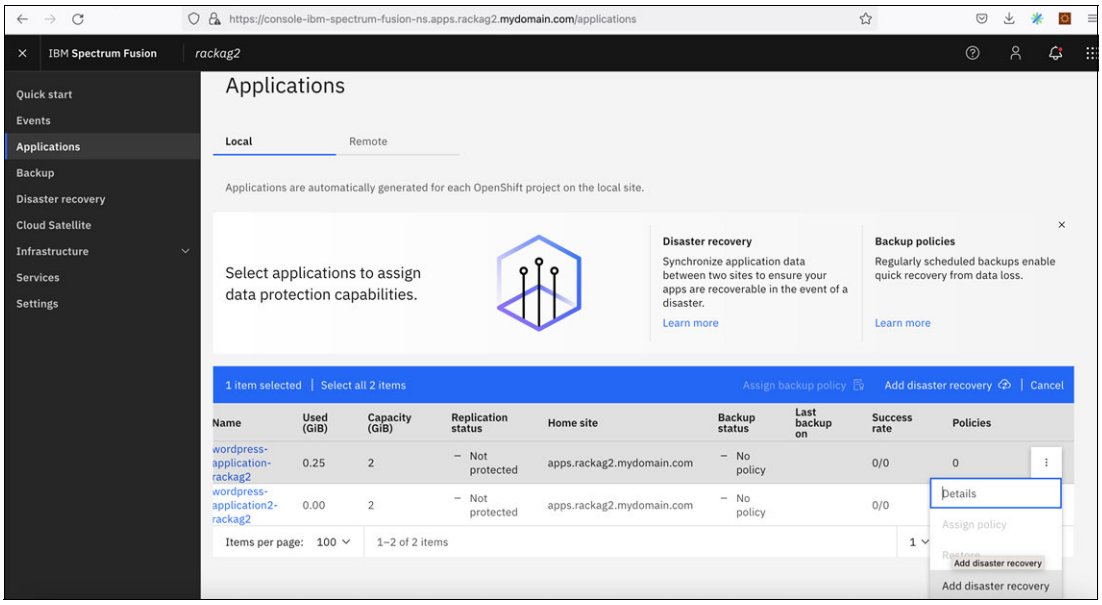

Figure 3-3 Enroll application for disaster recovery from Applications page

Method 2: Applications details page

Here are the steps for setting up the DR Method2 using the **Application details** page.

- Click on the application from the **Applications** page.
- The **Application details** page opens up as shown in Figure 3-4 on page 39.
- Click on **Actions** to open the menu.
- Click on **Add disaster recovery** from the drop-down menu.

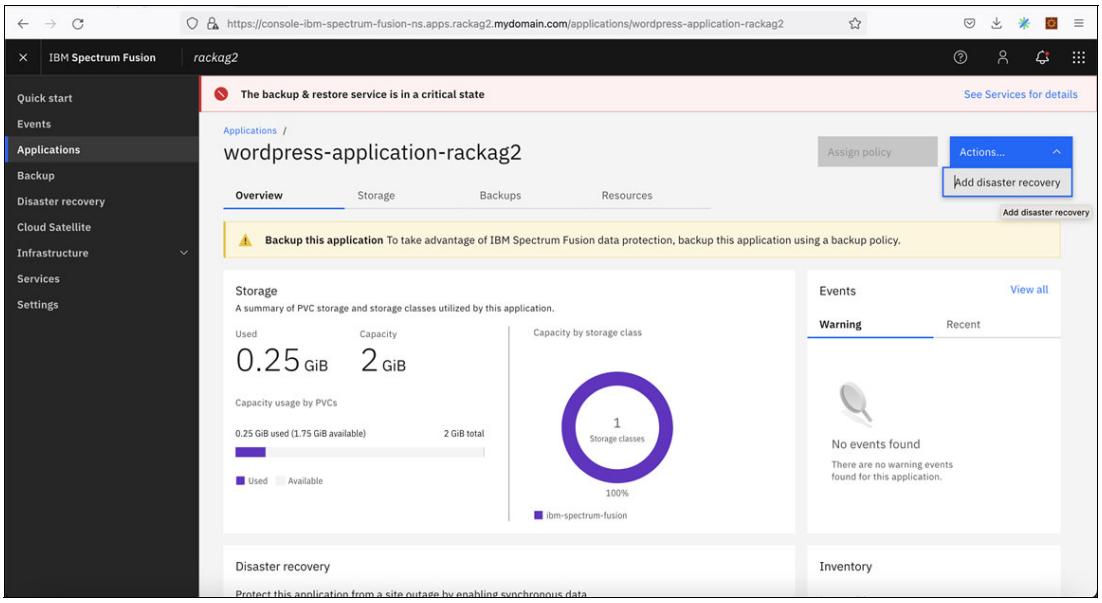

Figure 3-4 Application details page

Method 3: Application overview page

You can also setup an application for disaster recovery from the **Overview** tab of the **Applications details** page.

Here are the steps for setting up the DR Method3 using the **Overview** tab:

- Open the **Application details** page.
- Click on **Overview** tab.
- In the Disaster recovery section, click on **Add disaster recovery** button as shown in Figure 3-5.

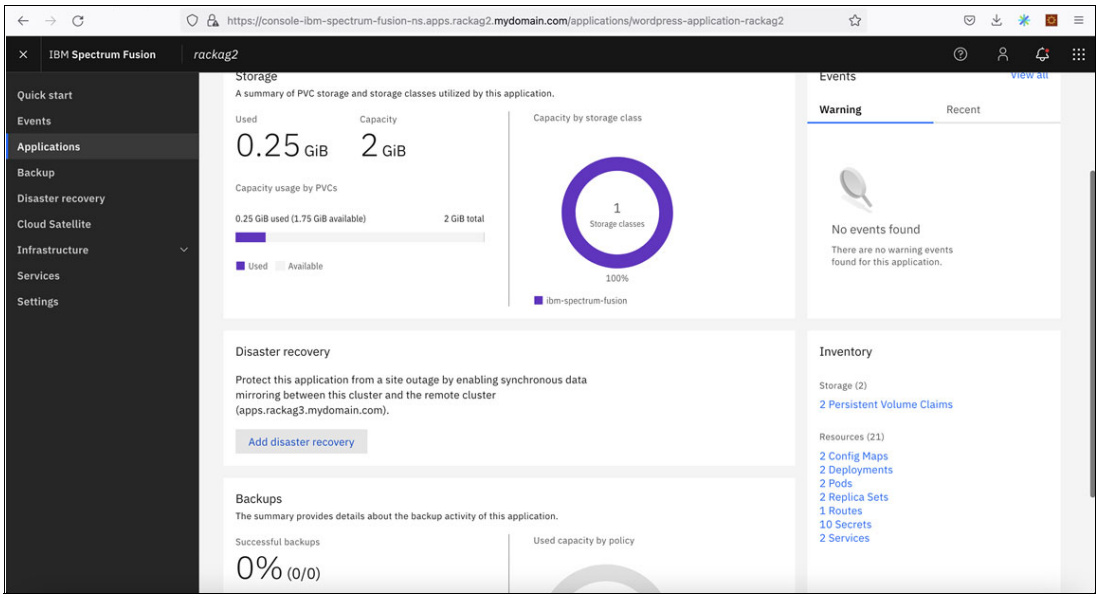

Figure 3-5 Application Overview tab Add disaster recovery button

Setting up multiple applications for DR

You can enroll multiple applications for DR simultaneously described in the following steps:

1. Open the **Applications** page.
2. On the **Local** tab, click the **checkbox** next to each Application name as shown in Figure 3-6.
3. Click on the **Add disaster recovery** button to enroll the selected applications for disaster recovery.

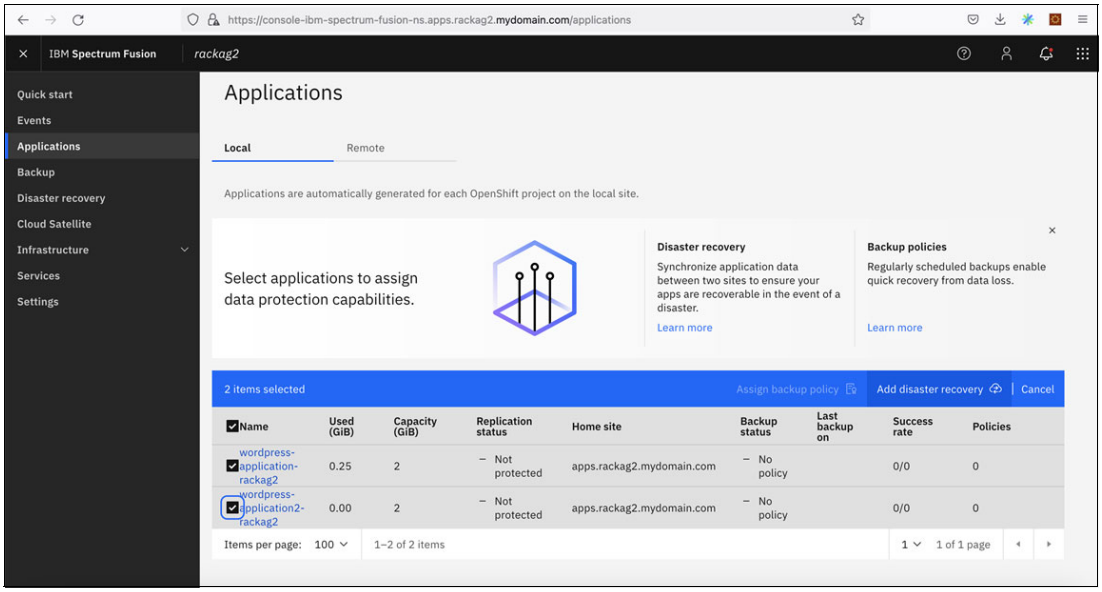

Figure 3-6 Enroll multiple applications for disaster recovery

Add application(s) for Disaster recovery

Once the **Add disaster recovery** button is clicked using any of the methods described, the DR enrollment process is initiated. In the background, the storage is replicated and synchronized across the two sites. The persistent volumes associated with the application is shared and becomes visible across both the sites.

Follow these steps to complete the DR enrollment process:

1. You will see a dialog box for confirmation as shown in Figure 3-7 on page 41.
2. Click on the **Add** button to confirm the completion of the action.

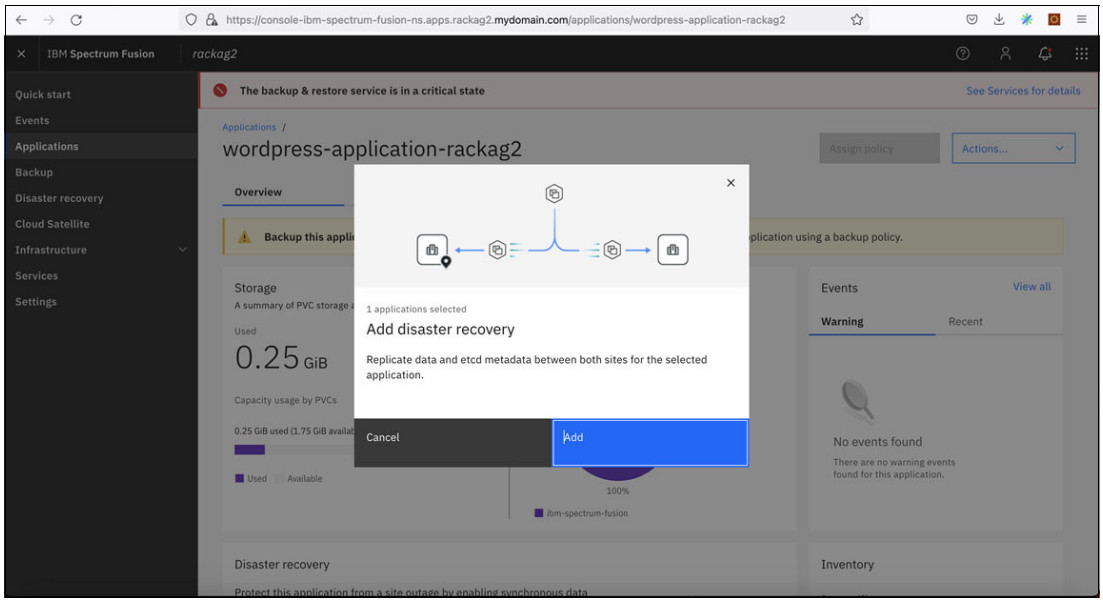

Figure 3-7 Confirmation dialog box for Disaster Recovery

3. The disaster recovery enrolment for the application(s) is started and as the message **Adding disaster recovery** is displayed for the number of applications selected as shown in Figure 3-8.

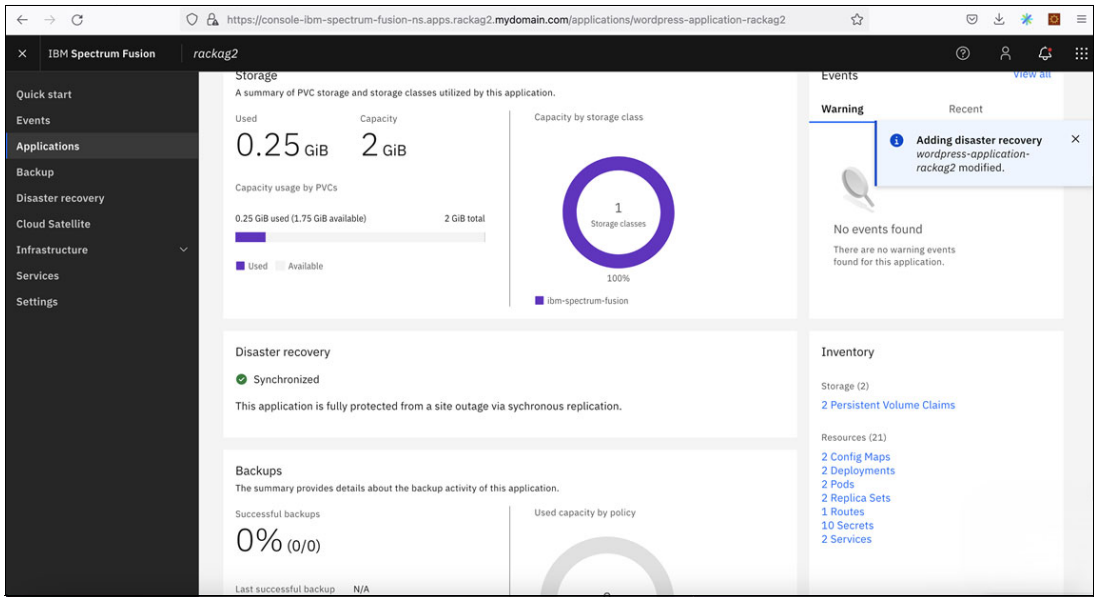

Figure 3-8 Disaster recovery enrollment process in progress

4. Once the enrollment is complete, the application can be viewed as **Synchronized**. The replication status is shown as **Synchronized** on both sites.
   a. On the local site, it is seen on the following pages.
      i. On the **Local** tab of the **Applications** page as shown in Figure 3-9.

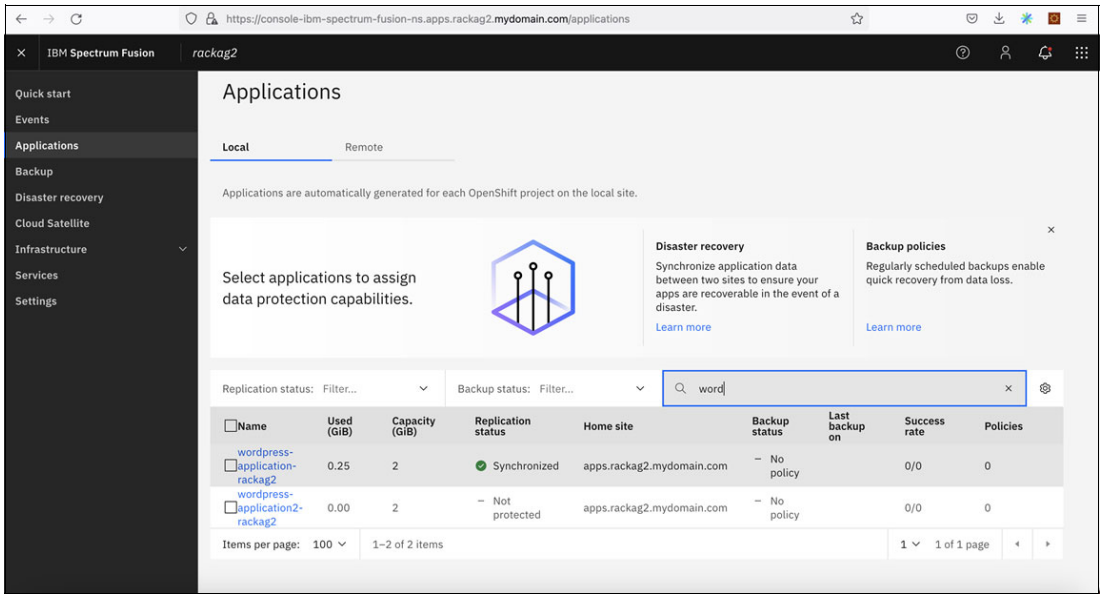

Figure 3-9 Local site - Applications page

      ii. On the **Application details** page in the **Overview** tab as shown in Figure 3-10.

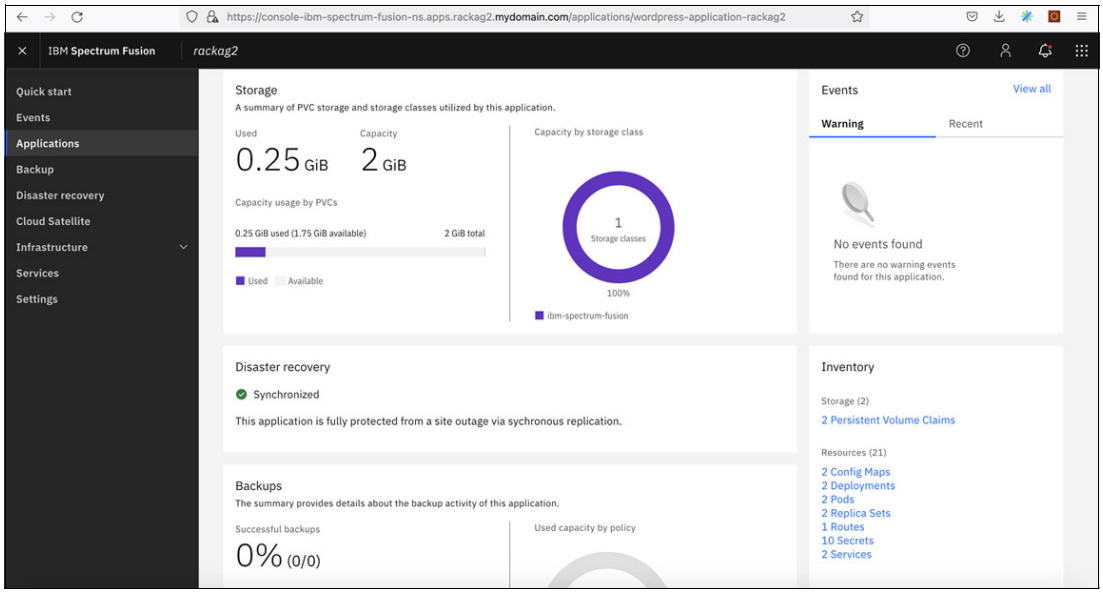

Figure 3-10 Application details page - Overview tab

   b. On the remote site, go to the **Applications** page.
      i. Go to the **Remote** tab of remote site as shown in Figure 3-11 on page 43.

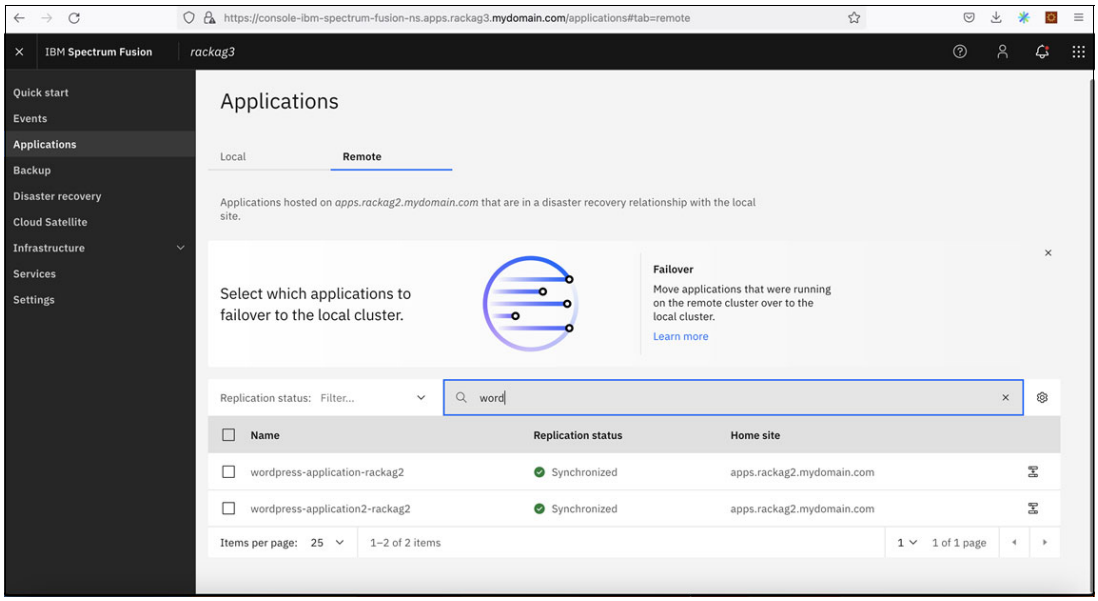

Figure 3-11 Remote site - Applications page

3.1.3 Failover

The applications are enrolled for disaster recovery and the persistent volumes are shared across the sites. In case, Site1 is unavailable due to unknown reasons or for planned maintenance work, the applications can still be accessed from Site2.

View the applications for failover by following these steps:

1. Login to IBM Spectrum Fusion of the remote site.
2. Go to **Applications** page.
3. Click on **Remote** tab.

4. View the applications as shown in Figure 3-12.
   a. The **Replication status** should show as **Synchronized**.
   b. The Home site should reflect the local site server URL.

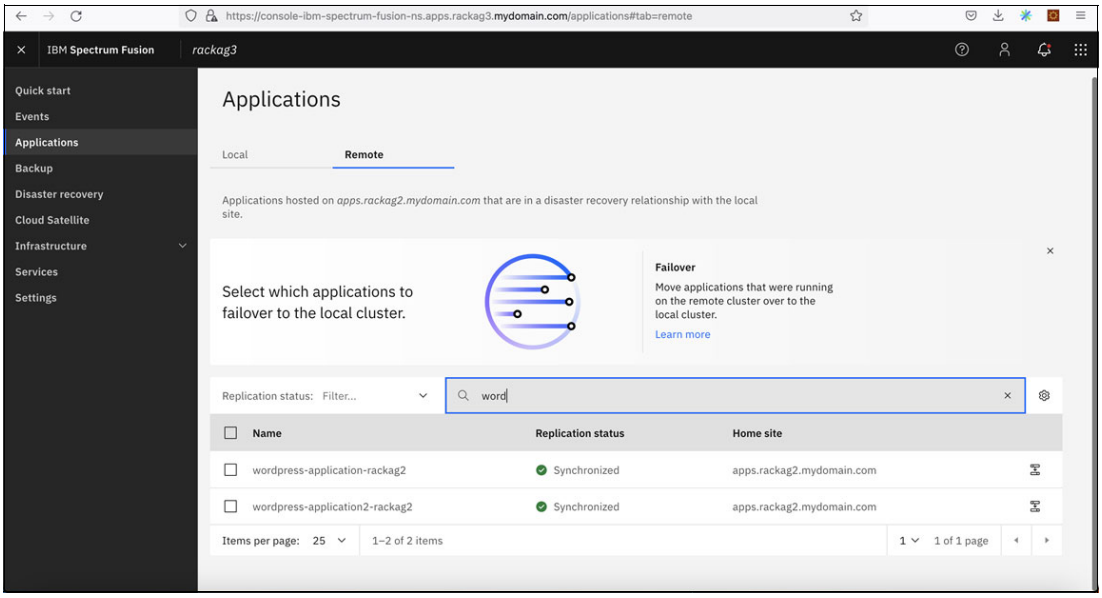

Figure 3-12 View applications for failover on the remote site

Initiate the failover process

This section describes the steps of the failover process:

1. Select an application for failover on the remote site as shown in Figure 3-13.

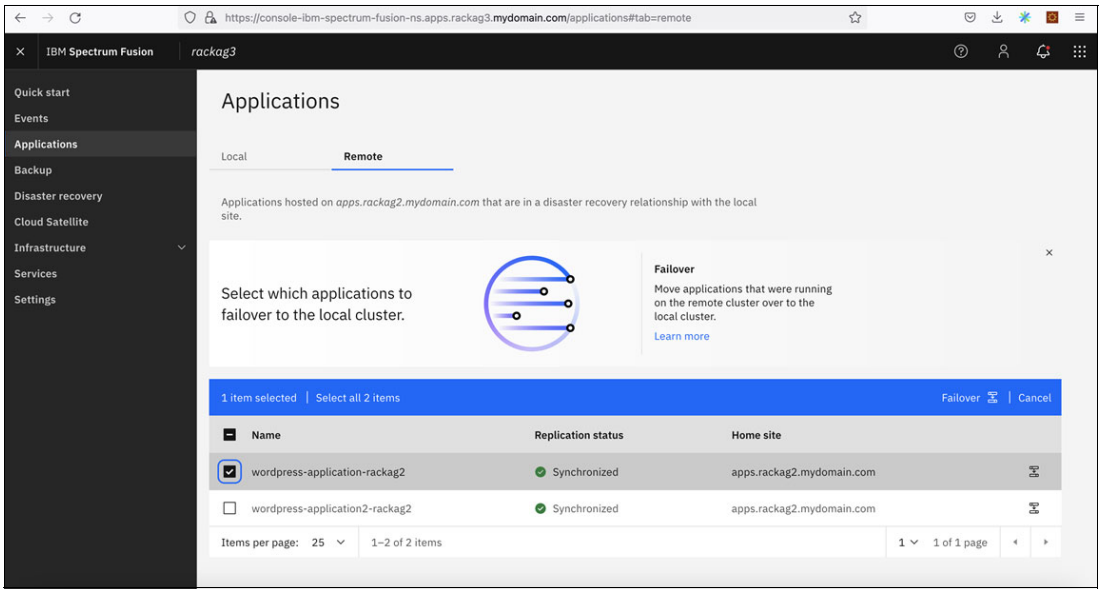

Figure 3-13 Failover process

2. Click on **the Failover** button to initiate failover. The dialog box will appear as shown in Figure 3-14 on page 45.

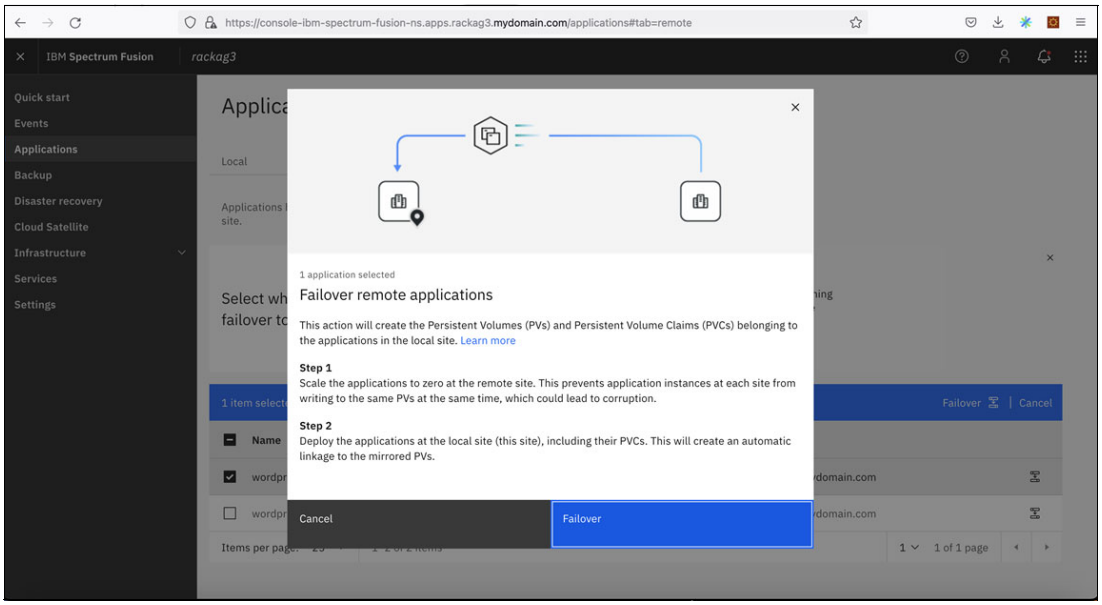

Figure 3-14 Confirmation message for failover process initiation.

3. The replication status for the application is changed from **Synchronized** to **Failover in progress** on the **Remote** tab of the remote site as shown in Figure 3-15.

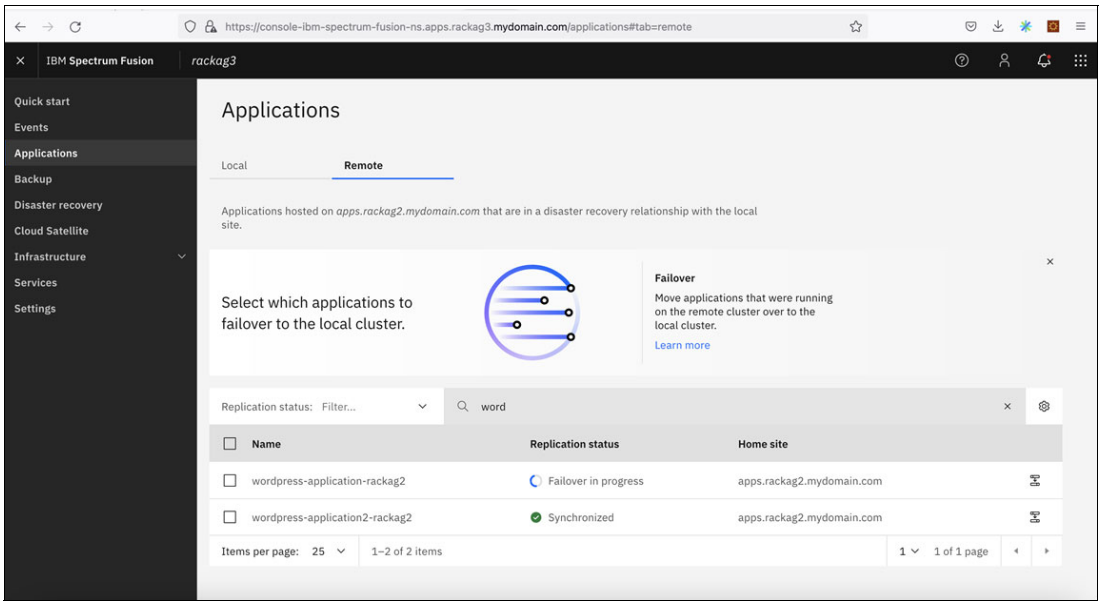

Figure 3-15 Failover progress

4. Once the failover is complete for the application, the Failover complete message is displayed as shown in Figure 3-16.

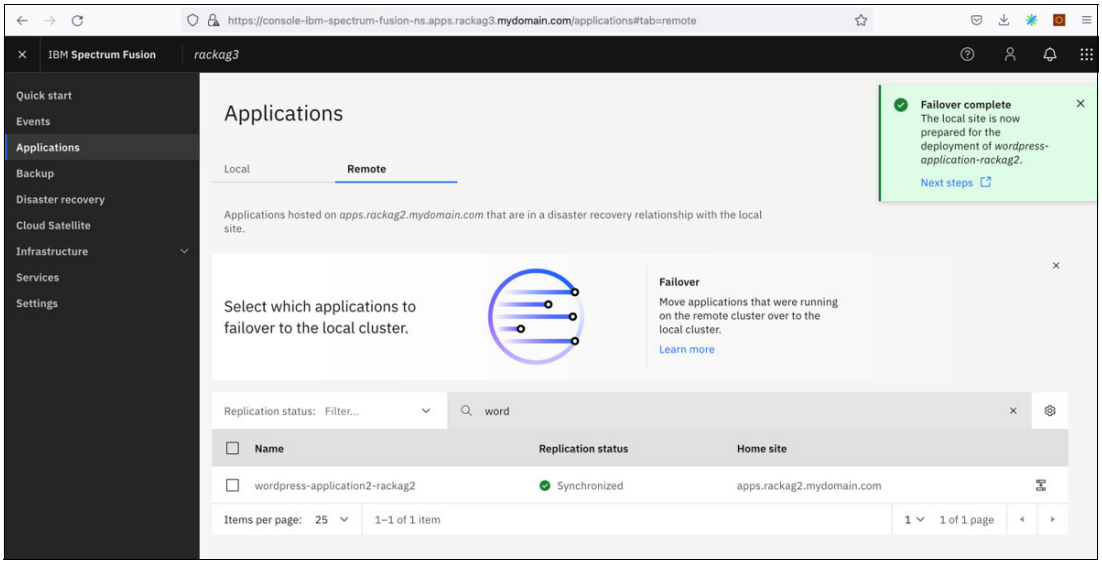

Figure 3-16 Failover completion message

5. The failed over application(s) is removed from the **Remote** tab of the remote site as shown in Figure 3-17.

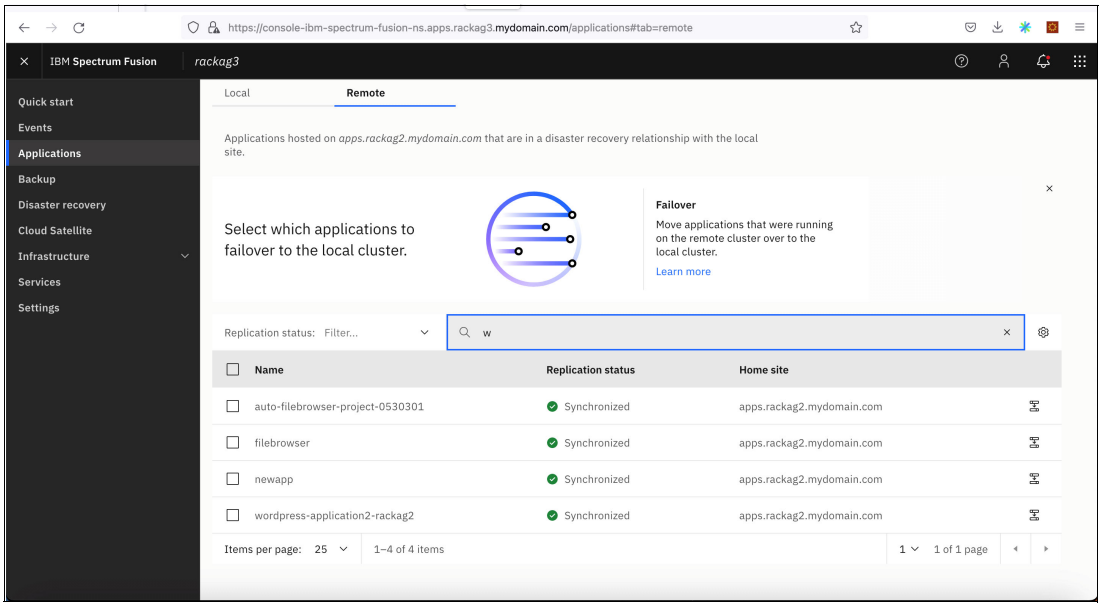

Figure 3-17 Remote tab

6. Failed over application(s) appears on **Local** tab of the remote site as shown in Figure 3-18 on page 47.

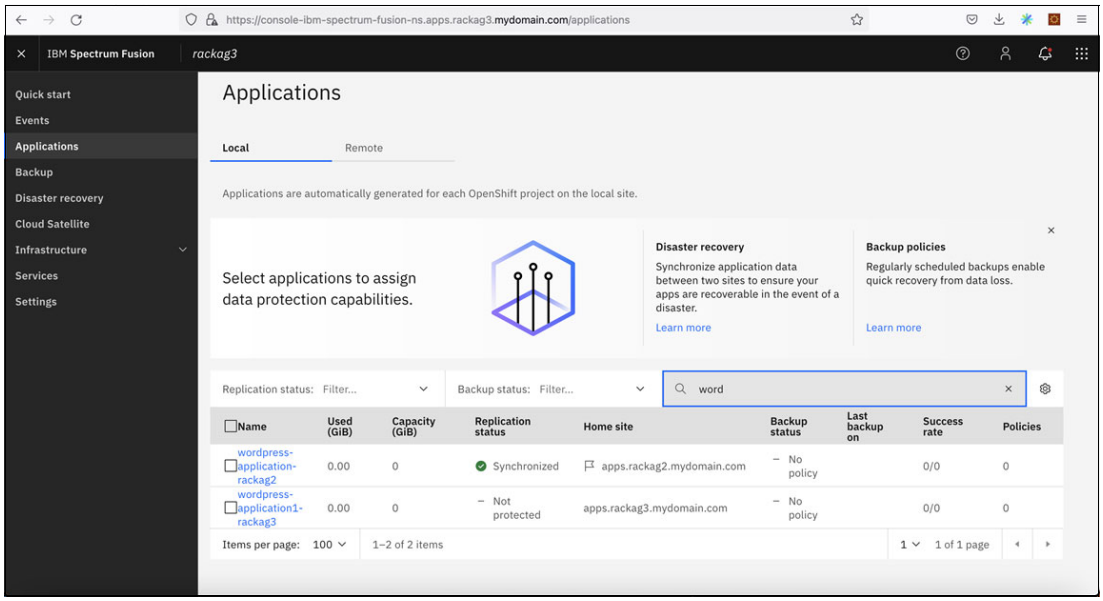

Figure 3-18 Local tab of Remote site of Application page

7. Now, redeploy the failed over application(s) on remote site as shown in Figure 3-19.

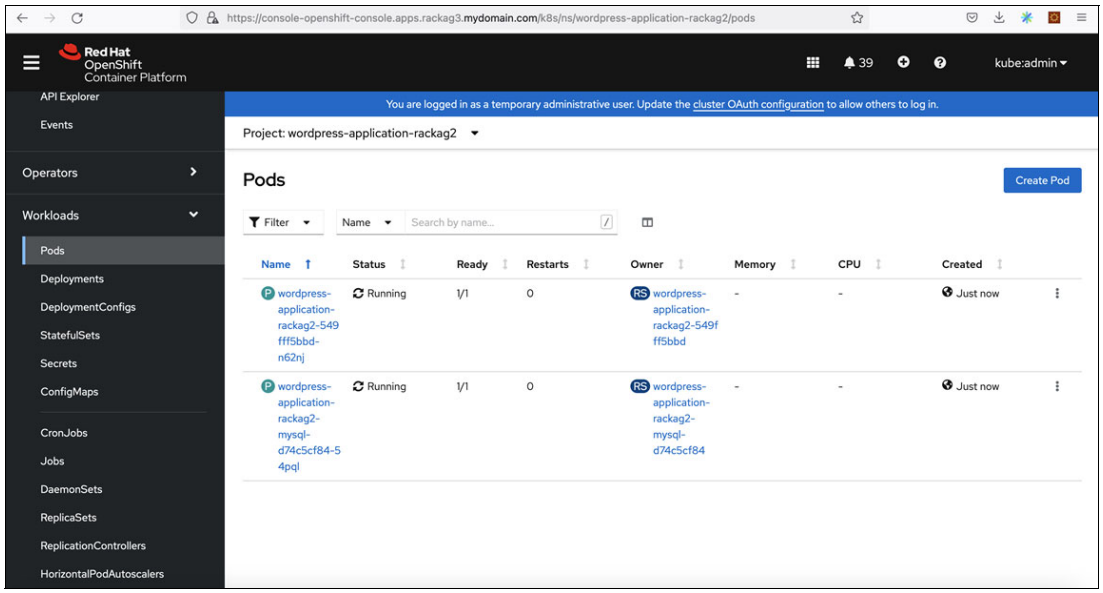

Figure 3-19 Remote site - Application deployment from OpenShift console

8. Access redeployed application(s) from remote site as shown in Figure 3-20.

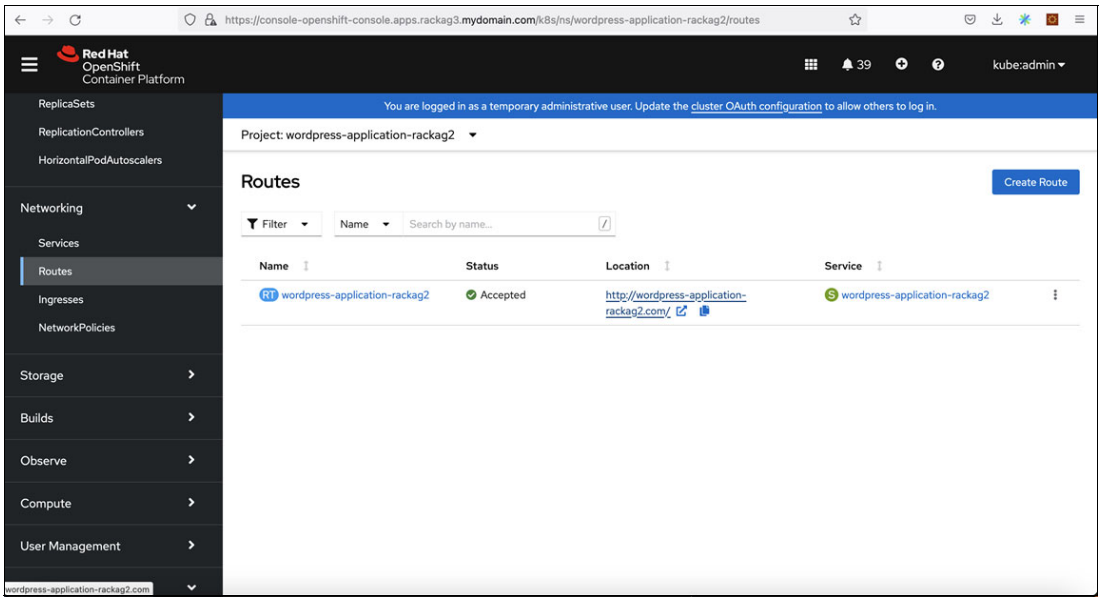

Figure 3-20 Remote site - OpenShift Console - Application route

9. View failed over application on remote site as shown in Figure 3-21.

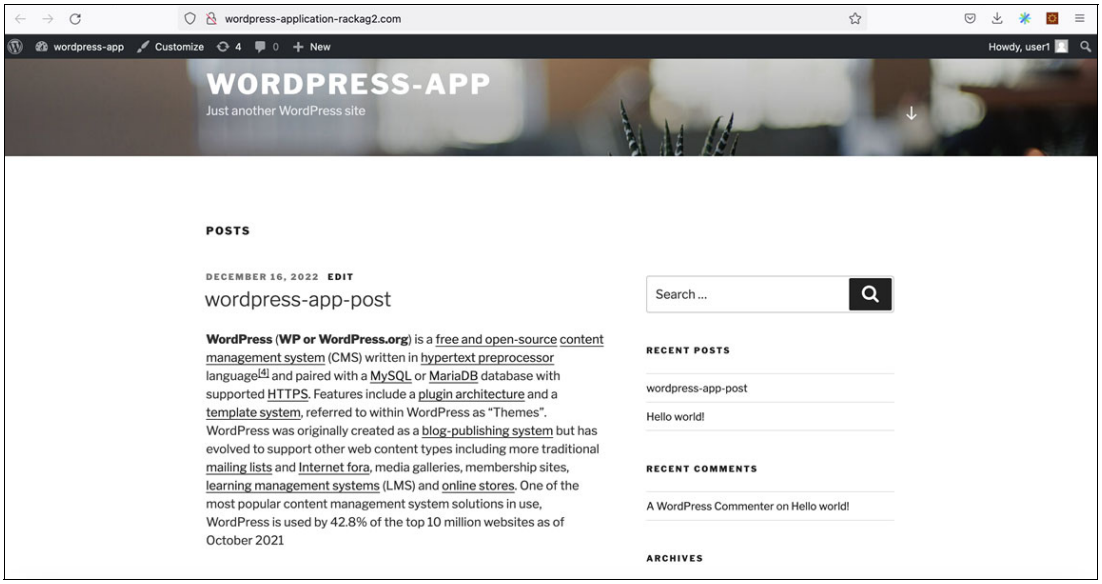

Figure 3-21 Application access from remote site

3.1.4 Failback

Once the Site1 is back online, you may want to failback the applications from remote site (Site2) to local site (Site1).

View failed over application(s)

To view the failed over application(s) from local to remote site, follow the these steps:

1. Login to IBM Spectrum Fusion of local site (Site1).
2. Go to **Applications** page.
3. Go to the **Remote** tab.
4. Check for application(s) with **Replication status** as **Synchronized**.
5. Also, check the name of **Home** site. It should reflect the Site1 as shown in Figure 3-22.

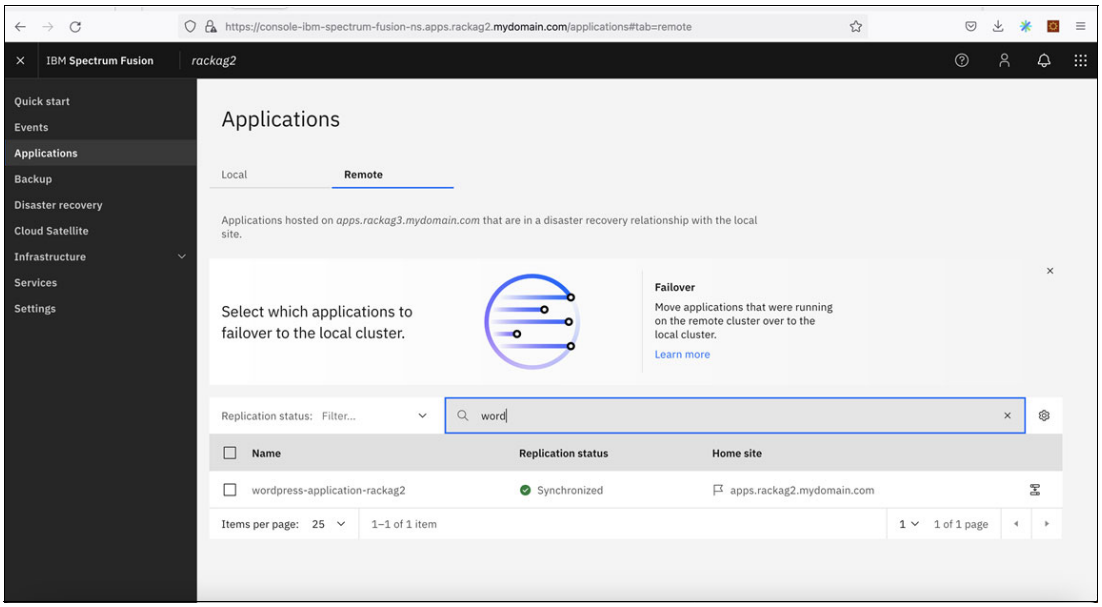

Figure 3-22 View failed over applications

Initiate failback process

Follow these steps to failback the application(s) from the remote site (Site2):

1. Login to IBM Spectrum Fusion of local site (Site1).
2. Go to the **Applications** page.
3. Go to the **Remote** tab.
4. Select application(s) for failback from the **Remote** tab of the local site.

5. Click **Failover** button as shown in Figure 3-23.

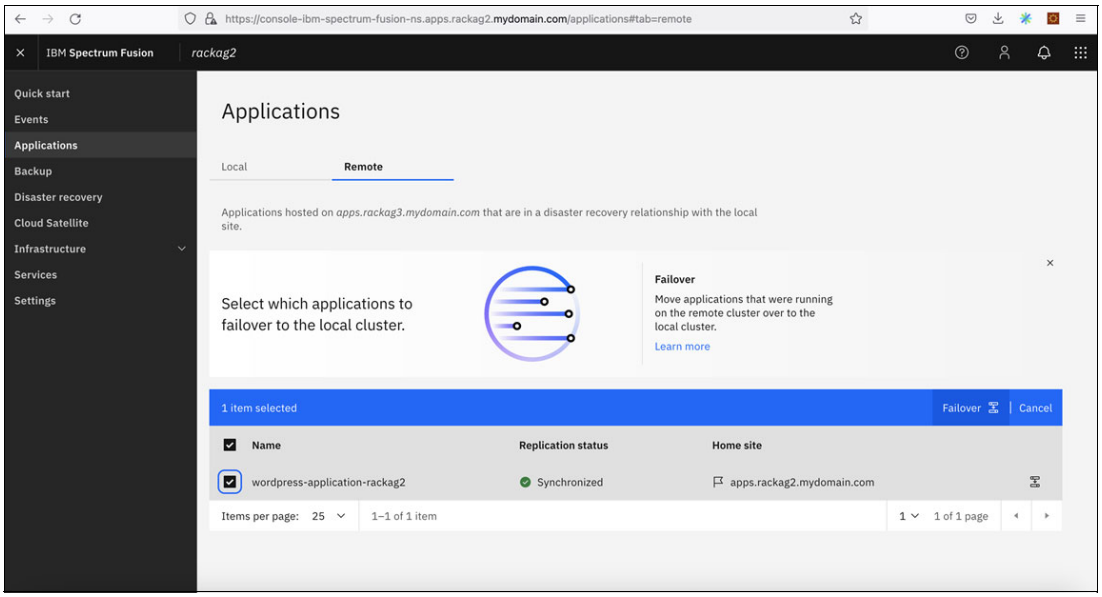

Figure 3-23 Select the application for failback and initiate failback process

6. The Failback dialog box message appears as shown in Figure 3-24.
7. Click on **Failover** button to initiate failback.

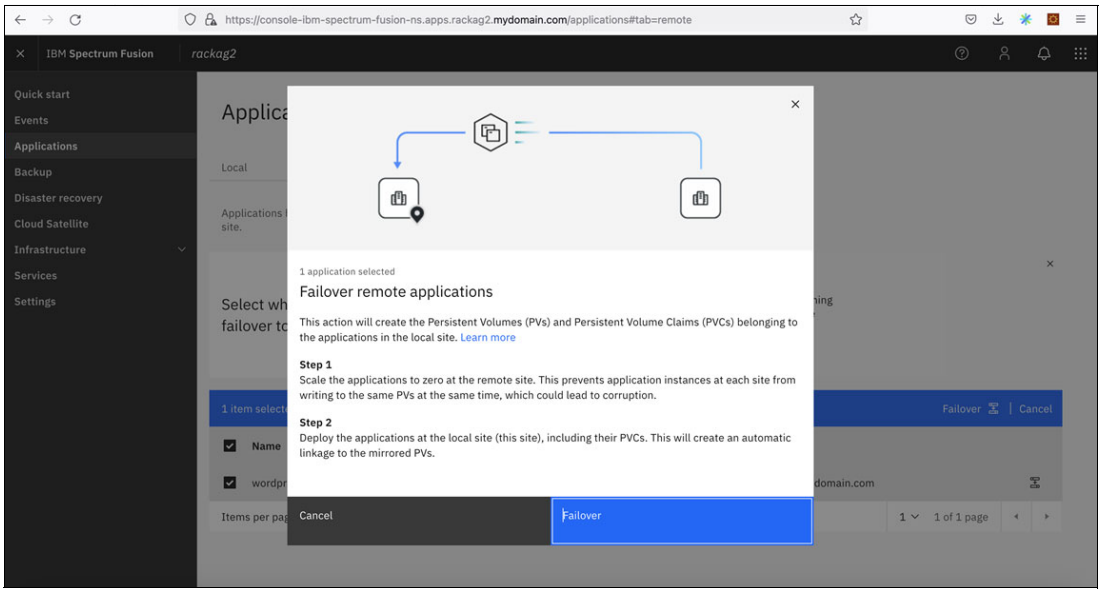

Figure 3-24 Failback dialog box

8. The message “failover initiated” is displayed for the failback initiated application(s) as shown in Figure 3-25 on page 51.

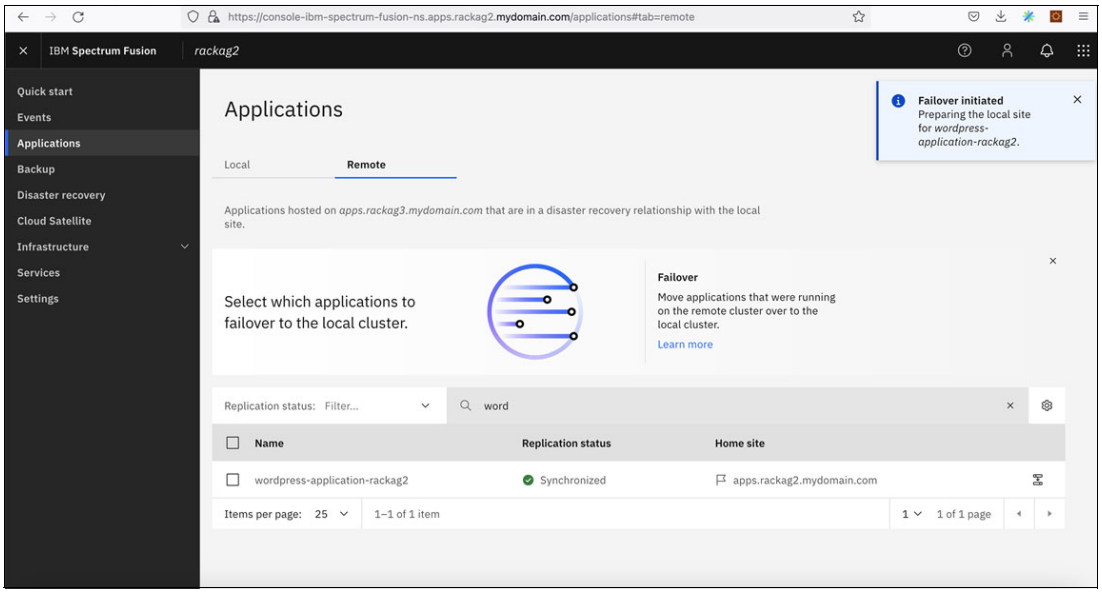

Figure 3-25 Failback message

9. The replication status is changed from **Synchronized** to **Failback in progress** as shown in Figure 3-26. Wait for the failback operation to complete.

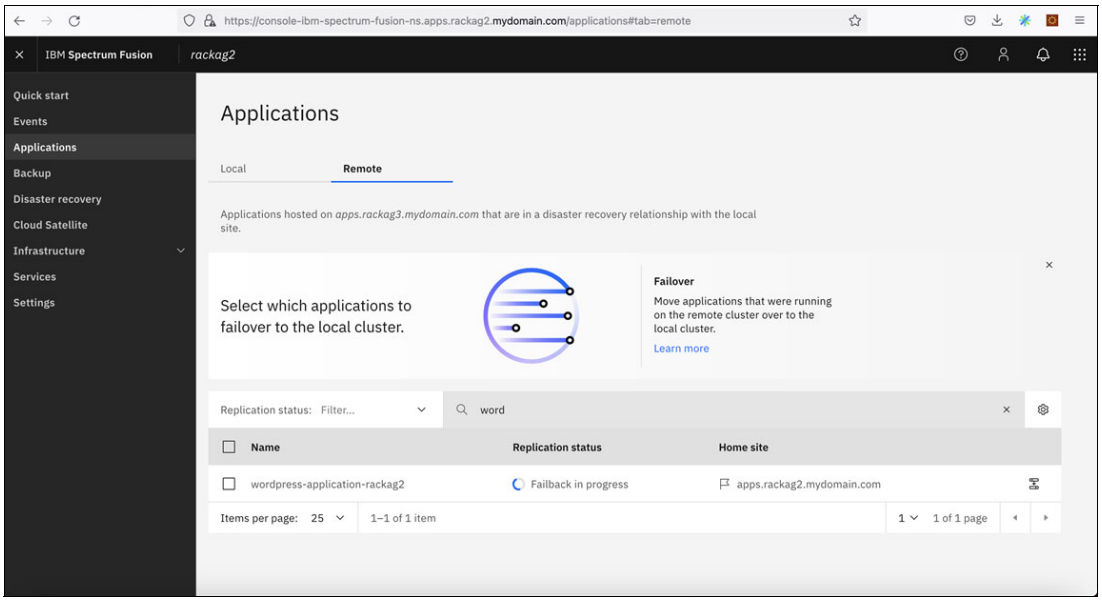

Figure 3-26 Failback operation progress is displayed

10. When the failback completes, the application is removed from the **Remote** tab of the **Applications** page as shown in Figure 3-27.

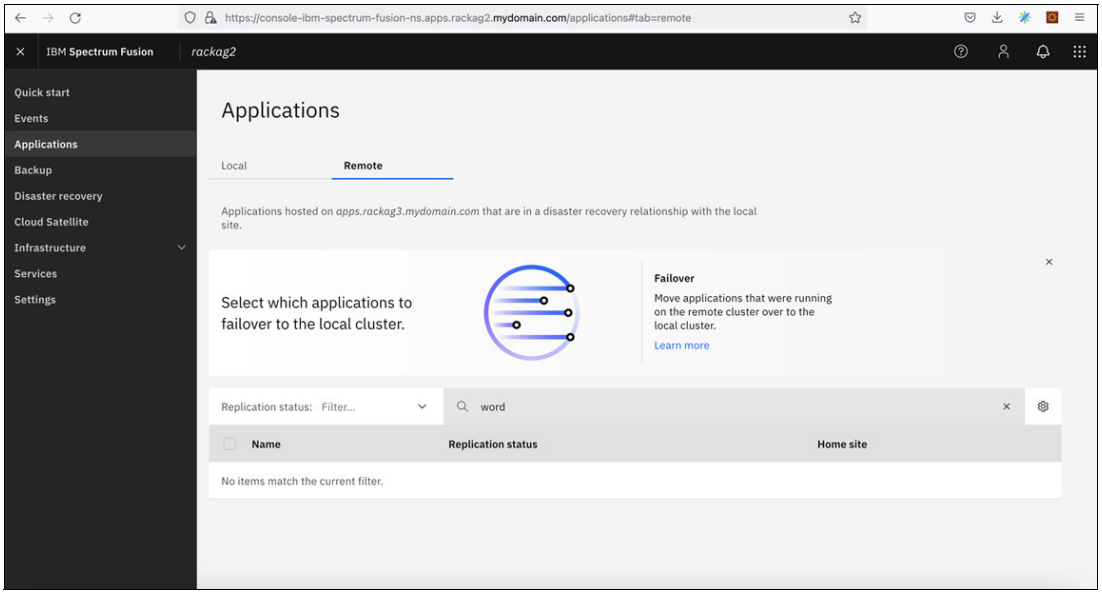

Figure 3-27 Failed back application(s) removed from Remote tab of local site

11. Go to the **Local** tab of the **Applications** page. Check for the failed back application. It should reappear on **Local** tab of local site as shown in Figure 3-28.

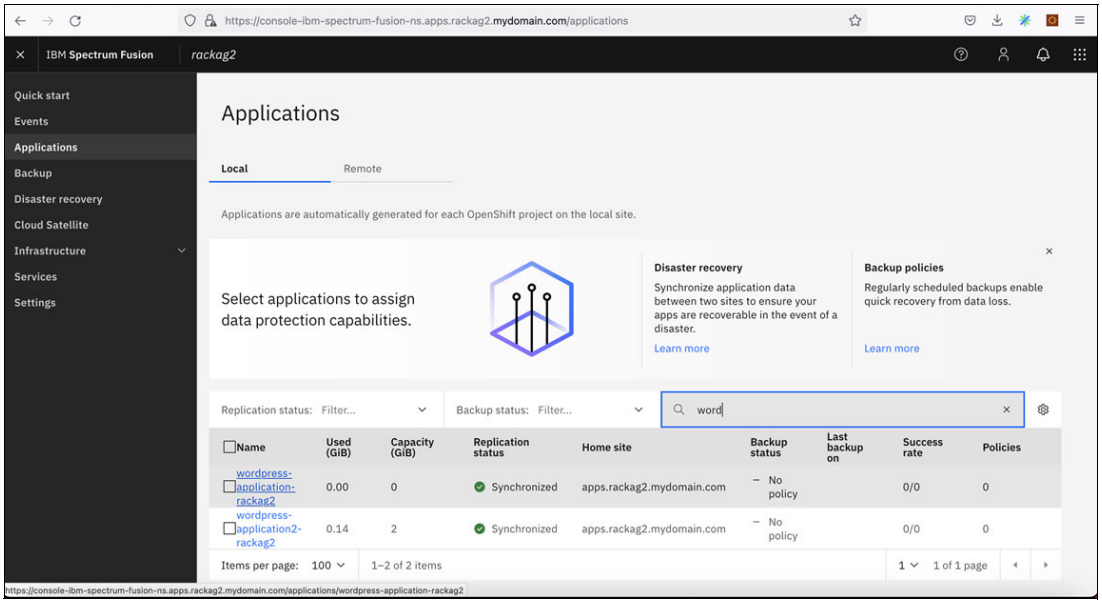

Figure 3-28 Check failed back application status on Local tab of local site

12. Access the redeployed application(s) on local site to ensure the application is up and running as shown in Figure 3-29 on page 53.

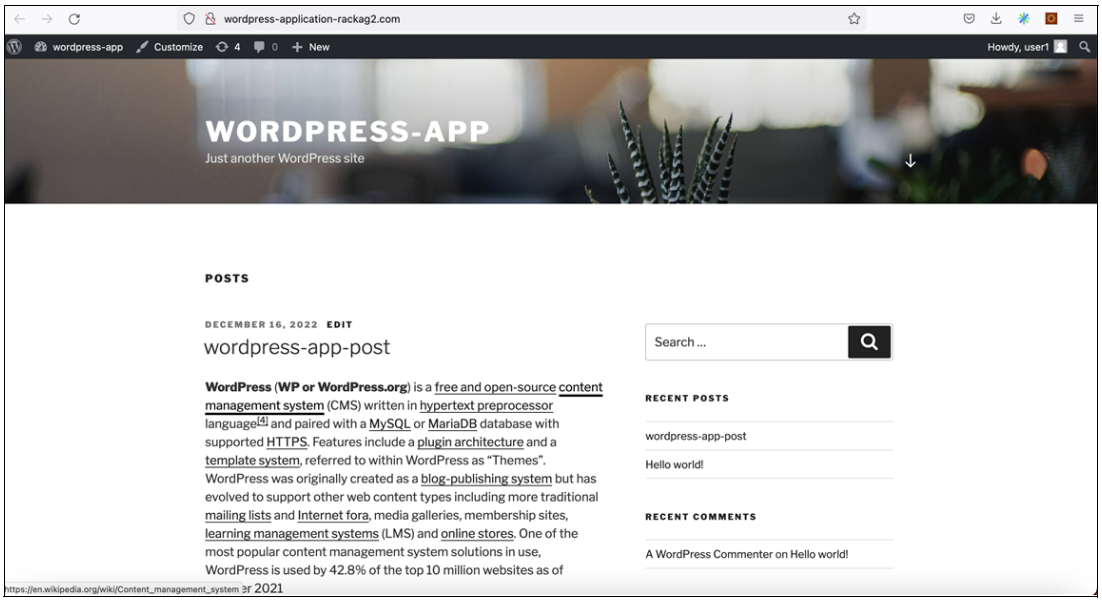

Figure 3-29 View failed back application(s) on local site

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.

- *IBM Storage Fusion Backup and Restore for Cloud Pak for Data*, REDP-5706
- *IBM Storage Fusion Product Guide*, REDP-5688

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:

- 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 Spectrum Fusion documentation

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

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

  https://www.ibm.com/products/storage-fusion
- IBM Storage Fusion HCI 2.5.2 Metro-DR

  https://www.ibm.com/docs/en/storage-fusion/2.5?topic=disaster-recovery

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

REDP-5708-00

ISBN 0738461164

Printed in U.S.A.