Redbooks
ibm.com/redbooks

IBM Fusion HCI

Metro Disaster Recovery Use Case

Pallavi Singh

Prakash Chandraya

Seema Pandhre

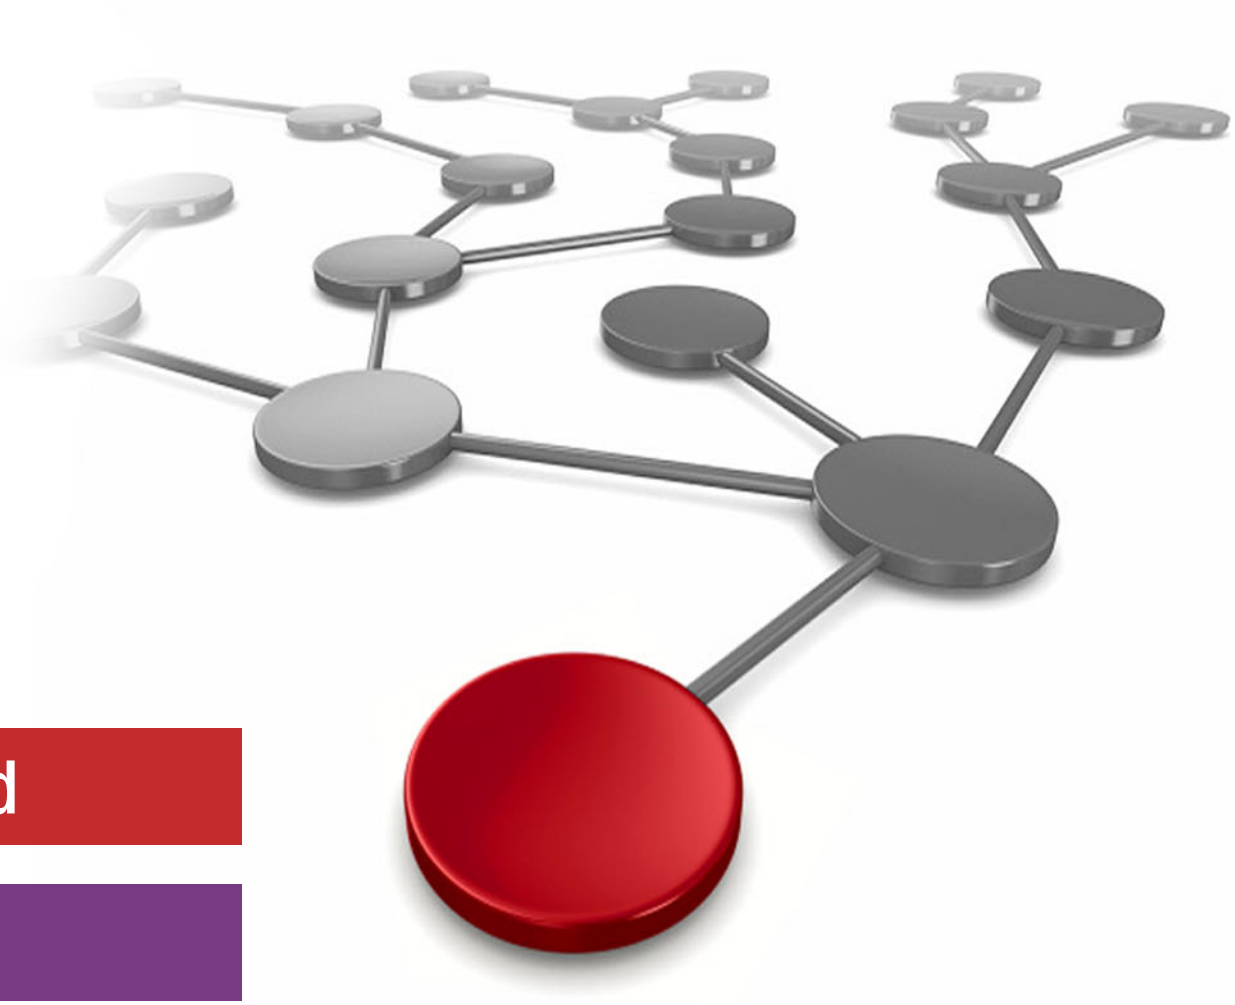

Cloud

Storage

IBM

Redpaper

IBM Redbooks

IBM Fusion HCI: Metro Disaster Recovery Use Case

September 2024

REDP-5708-01

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

Second Edition (September 2024)

This edition applies to Version 2, Releases 7 and 8 of IBM Fusion HCI.

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.

Redbooks (logo)  ® IBM® Redbooks®

The following terms are trademarks of other companies:

Red Hat, Ansible, 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

Metro disaster recovery (DR) provides two-way synchronous data replication between IBM® 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 DR solution is only available for metropolitan distance data centers with 40 milliseconds latency or less.

Note: The procedures described in this paper are for IBM Fusion HCI 2.7.x and 2.8.x.

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

Authors

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

Prakash Chandraya is an Advisory Software Engineer and works as a quality lead with IBM Systems Labs. He has 18 years of experience in product quality assurance. Prakash is a prior 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 17 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.

Pallavi Singh is a Senior Technical Staff Member and works as an architect with IBM System Labs. She has 20 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.

Special thanks to the following IBM Software Engineers:

Kaustubh Katruwar, IBM Systems, 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 (DR). He is a storage security enthusiast with inclination towards secure engineering practices and security posture.

Pallavi Joshi, IBM Systems, is a software engineer with the IBM Systems Labs. Her experience spans across back-end development with knowledge of Cloud orchestration and Disaster Recovery. She is skilled in developing RESTful APIs and web services for efficient data management with focus on microservices architecture and service mesh implementation for cloud-native applications. She is an enthusiast working towards latest trends in hybrid and multi-cloud adoption.

Tejshri Bhagwat, IBM Systems, is a software engineer with the IBM Systems Labs. Her experience spans across back-end development with knowledge of storage architectures, storage backup and replication, and storage resiliency.

Purnanand Kumar, IBM Systems, is a software engineer with the IBM Systems Labs. His experience spans across back-end development with knowledge of storage architectures and networking protocols. He is skilled in developing RESTful APIs and web services for efficient data management with focus on microservices architecture and service mesh implementation for cloud-native applications. He is a storage enthusiast working towards latest trends in hybrid and multi-cloud adoption.

Annaraya Sangappa Narasagond, IBM Systems, is a software engineer with the IBM Systems Development Labs. His experience spans across back-end development with knowledge on storage architectures and disaster recovery solutions. He is skilled in developing Kubernetes APIs and operators to orchestrate system functionalities in reference to microservices architecture for cloud native applications

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

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 Fusion HCI: Metro Disaster Recovery Use Case
as created or updated on September 25, 2024.

September 2024, Second Edition

This revision includes the following new and changed information.

New information

- The procedures described in this paper are for IBM Fusion HCI 2.7.x and 2.8.x.
- Global Data Platform Installation as a service
- Global Data Platform Configuration for Metro-DR pairing through services page

Changed information

- UI modifications for Disaster Recovery Topology page
- Multiple ways to enable and disable disaster recovery for applications
- UI modifications for Disaster Recovery Replicated applications page
- UI modifications for Software download with Entitled Systems Support (ESS)
- UI modifications for failover of applications made easy
- UI modifications while installing Global Data Platform as a service
- UI modification while configuring Global Data Platform for Metro-DR pairing

1

Metro disaster recovery deployment models

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

1.1 Introduction

Metro DR (Disaster Recovery) provides two-way synchronous data replication between IBM 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 DR solution is only available for metropolitan distance data centers with 40-millisecond latency or less.

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

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

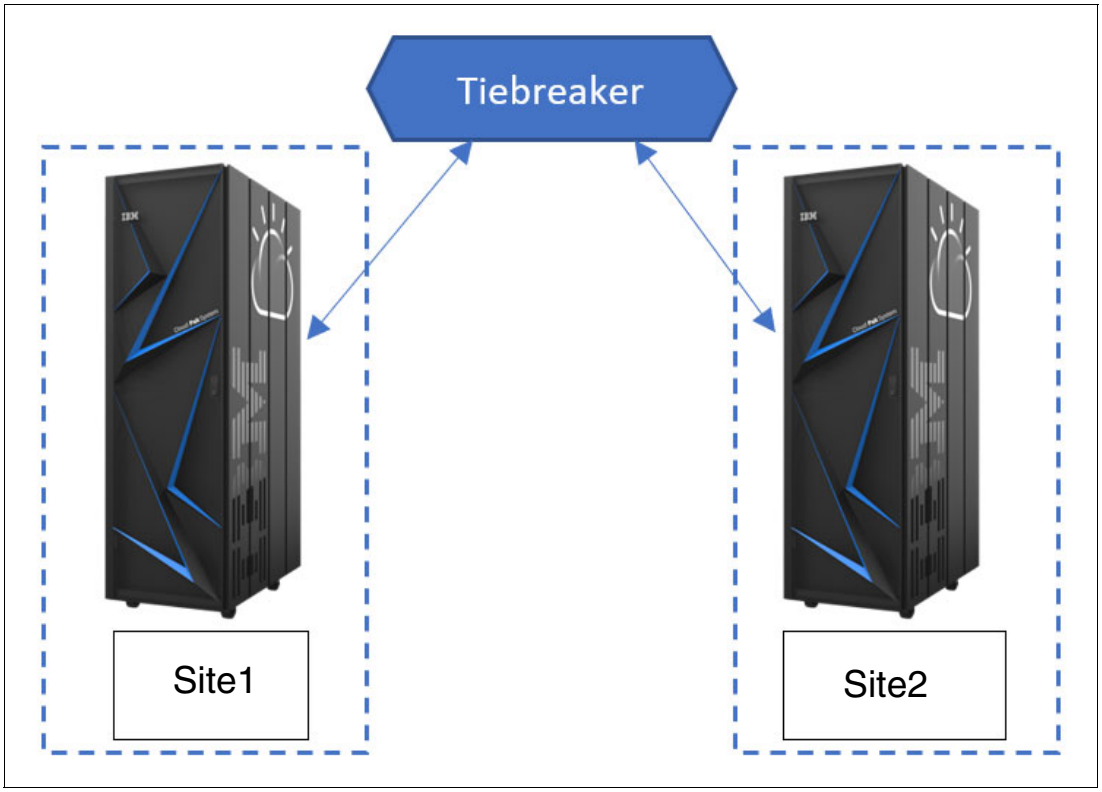

Figure 1-1 Metro DR configuration

IBM Fusion HCI is based on IBM Storage Scale and uses its “stretch cluster” feature to provide a unique active-active resiliency across data centers that are up to 150 km distant. This is achieved by spanning an IBM Storage Scale file system across two IBM 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 Fusion to synchronously replicate data to another nearby IBM Fusion, each with their own OpenShift clusters.

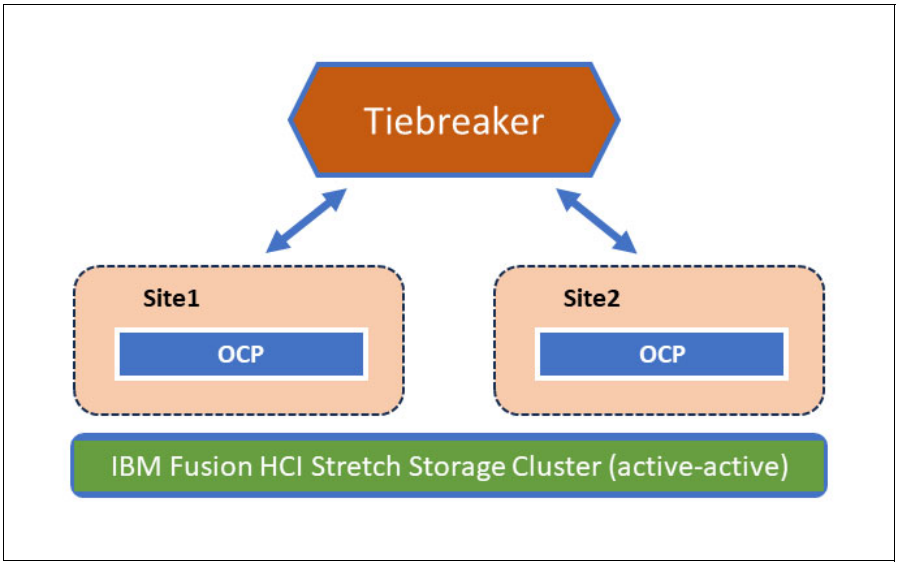

Figure 1-2 Metro DR component architecture

1.2 Deployment Models

There are two deployment models for Metro DR configuration.

1.2.1 Deployment Type #1

As a client, you buy a single IBM Fusion appliance. Later, as the business requirements increase, you need another IBM Fusion appliance as well if you want to achieve DR. In this case as shown in Figure 1-3, you convert the first appliance into Site1 (Refer to 2.1.3, “Site2 Installation” on page 25 for the steps **to configure Global Data Platform**) and then continue with the installation of the second appliance as Site2.

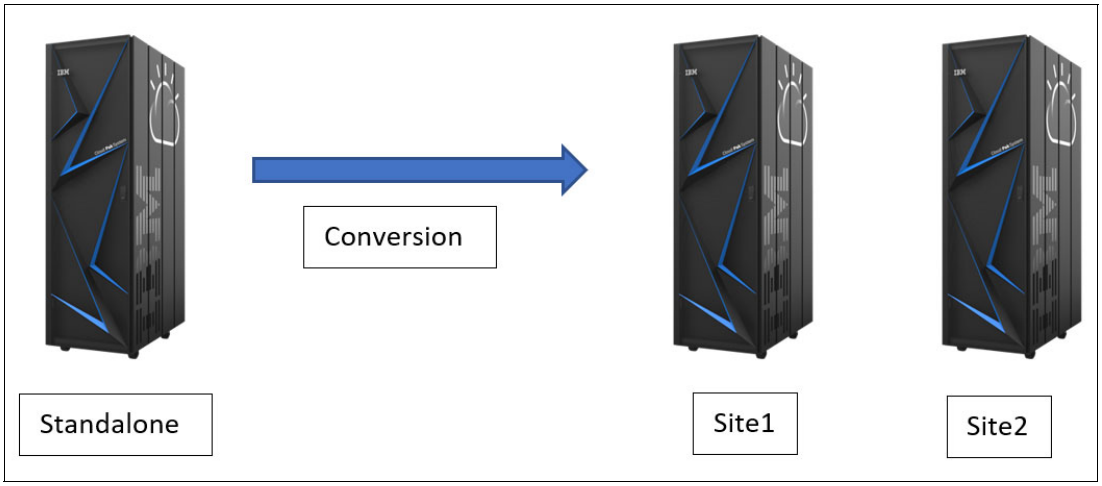

Figure 1-3 Deployment model Type #1 for setting up DR

1.2.2 Deployment Type #2

As a client, based on business requirements, you plan to have DR. In this case, you would have to purchase two IBM Fusion appliances, see Figure 1-4. The installation of the two appliances will be done in sequence. During IBM 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.

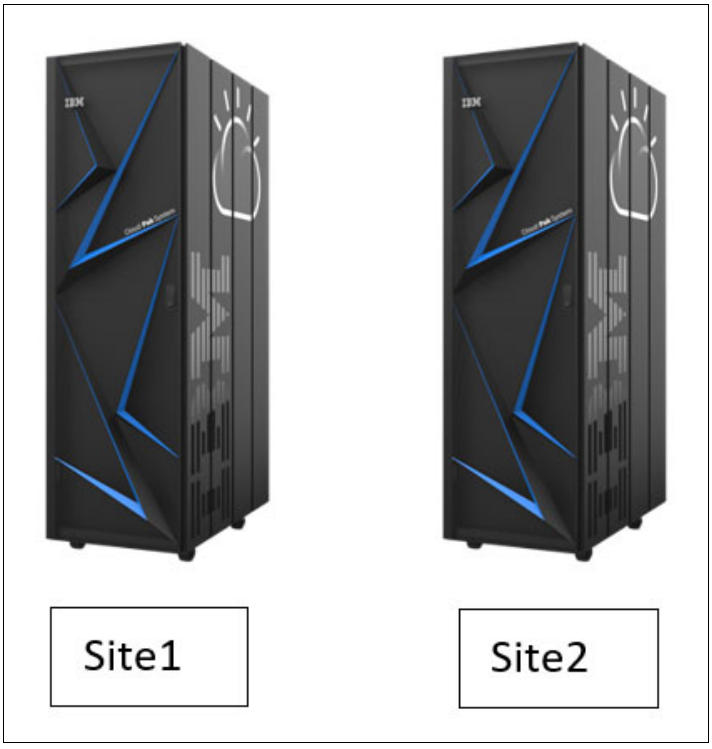

Figure 1-4 Deployment model Type #2 for setting up DR

2

IBM Fusion HCI Metro disaster recovery installation

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

2.1 Installation

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

2.1.1 Prerequisites

Before you begin installation, you need to ensure that the IBM Fusion HCI appliance is set up in the data center as per guidance from IBM.

1. For more information for IBM Fusion HCI 2.7 installation, see:
   - Planning and prerequisites R2.7.x for instructions to setup the appliance in Data Center.
   - General Metro-DR prerequisites R2.7.x for instructions to setup network and connectivity between Metro DR clusters.
2. For more information for IBM Fusion HCI 2.8 installation, see:
   - Planning and prerequisites R2.8.x for instructions to setup the appliance in Data Center.
   - General Metro-DR prerequisites R2.8.x for instructions to setup network and connectivity between Metro DR clusters.

2.1.2 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. Select the Deployment Type as **Single Rack** and click **Next** as shown in Figure 2-1.

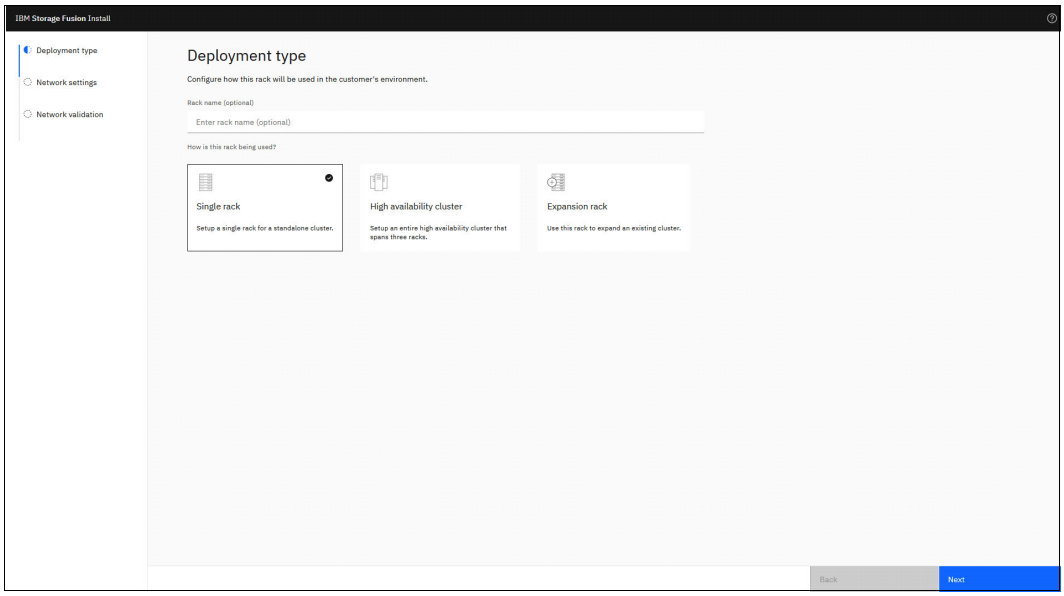

Figure 2-1 Deployment Type

2. Enter the details of the Network configuration as shown in Figure 2-2.

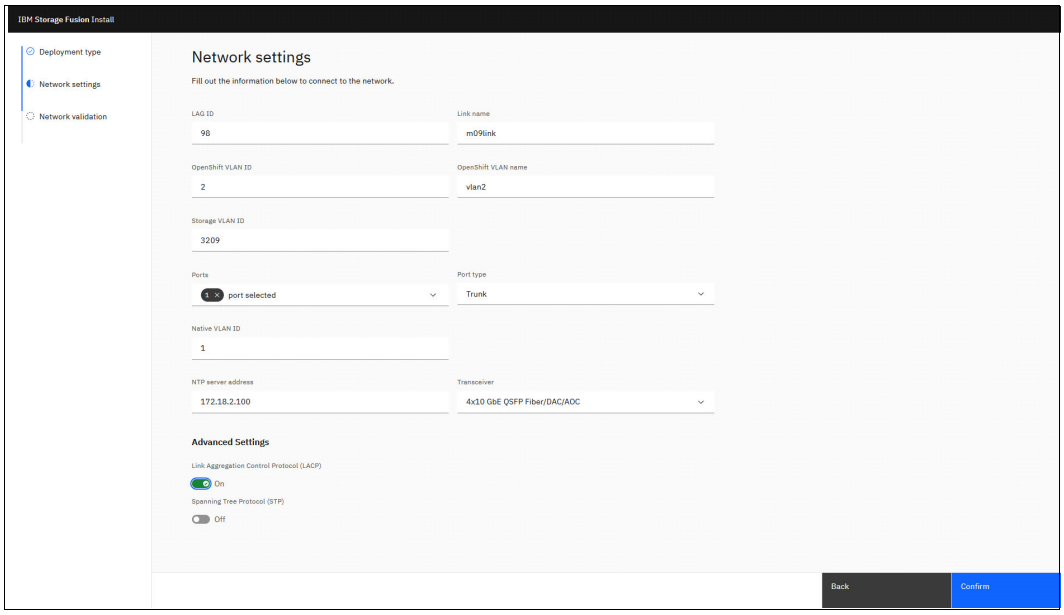

Figure 2-2 Network settings

3. Click on the **Confirm** button as shown in Figure 2-2.

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

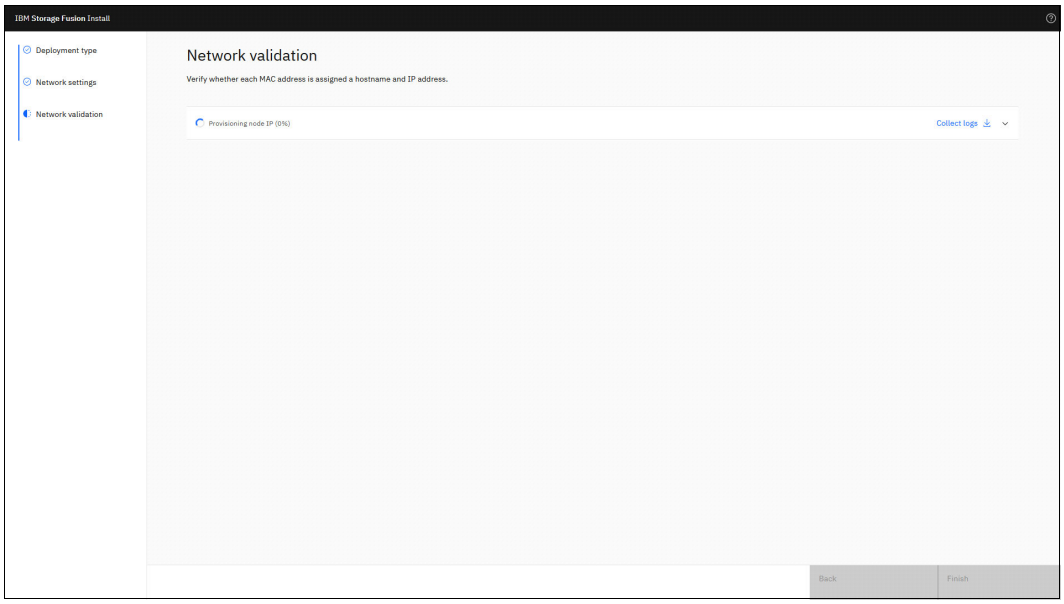

Figure 2-3 Network validation

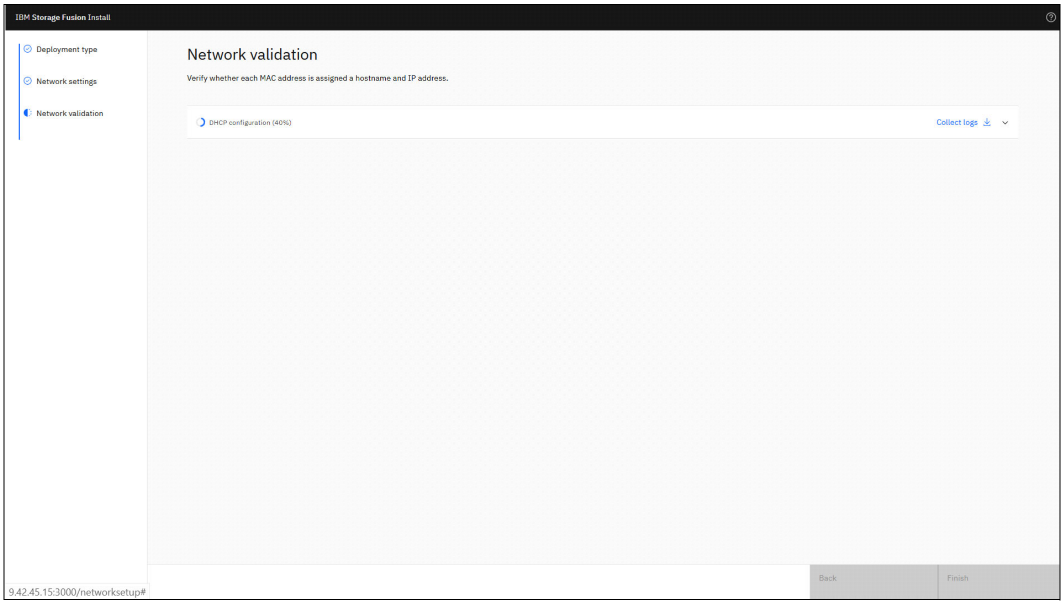

Figure 2-4   Network validation - DHCP configuration

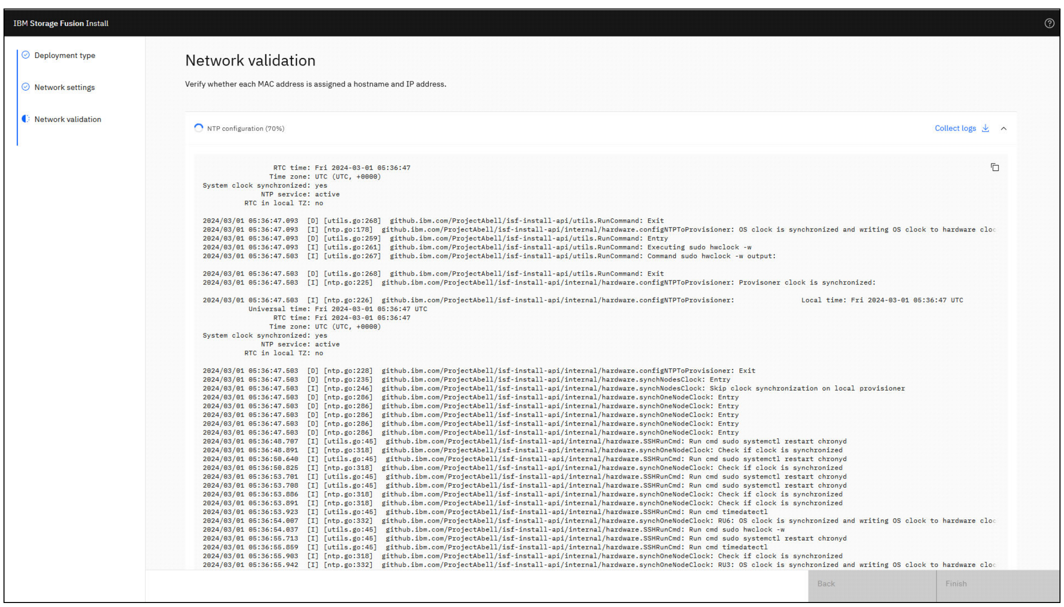

Figure 2-5   Network validation - NTP configuration

4. The network configuration is completed successfully as shown in Figure 2-6 on page 9.

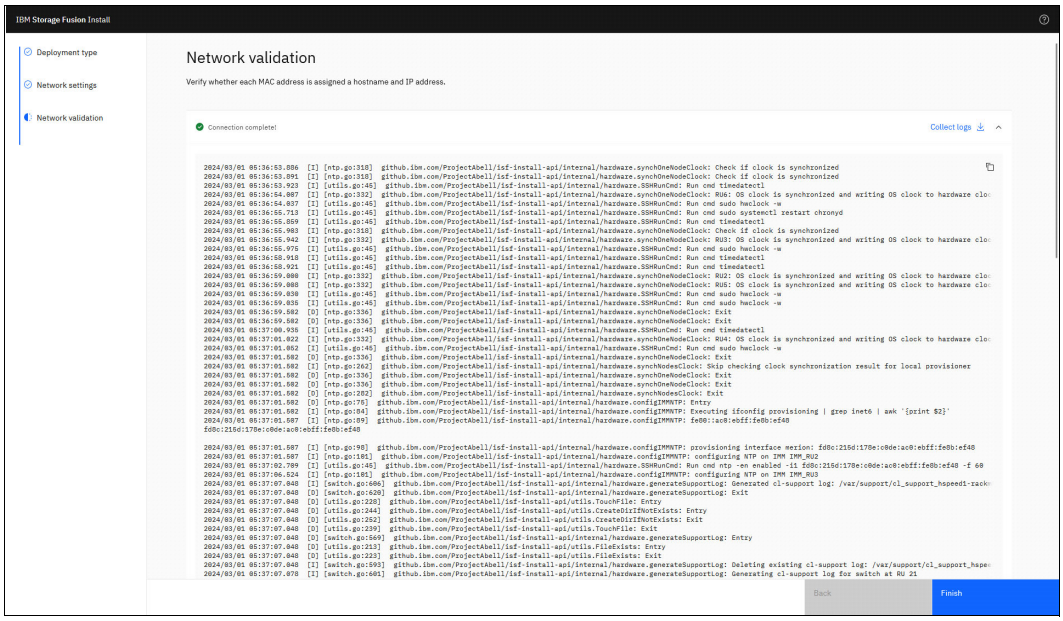

Figure 2-6 Network validation - Successful completion.

5. Click the **Finish** button. This page will continue with the installation as shown in Figure 2-7.

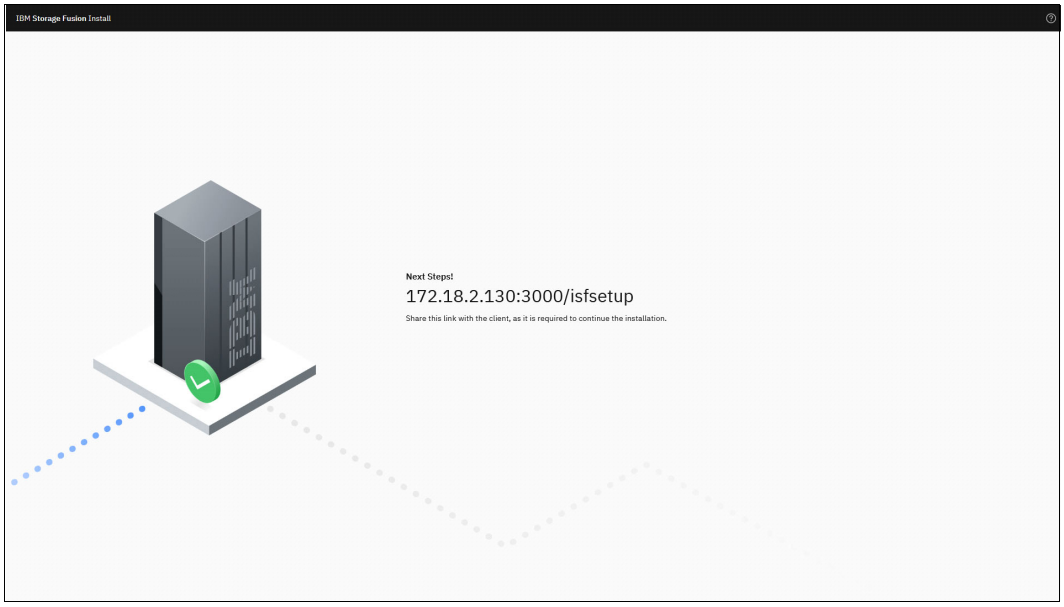

Figure 2-7 Next Steps

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

1. To proceed with the IBM 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 on page 10.

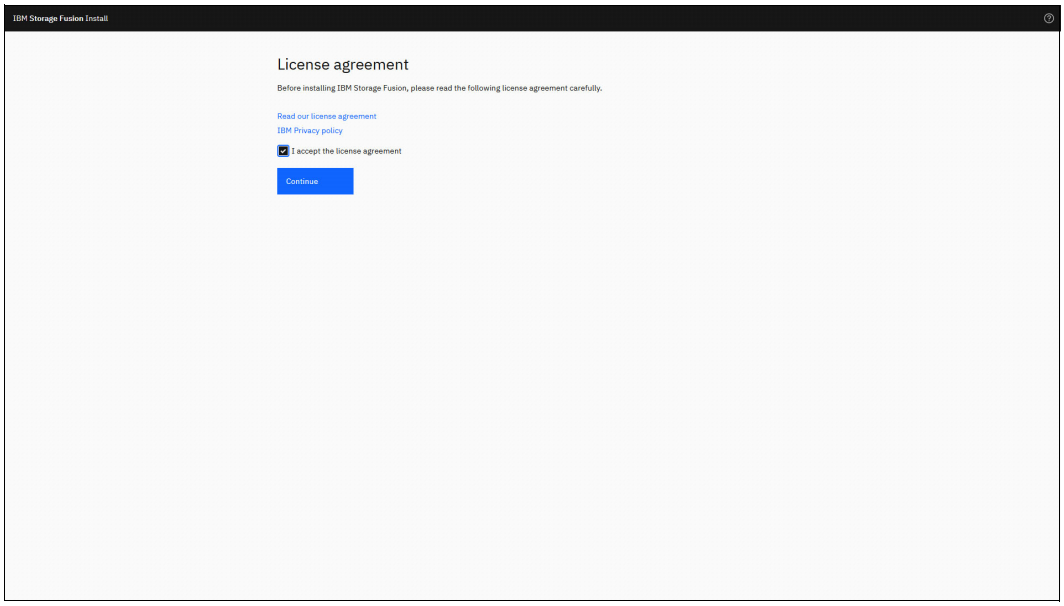

Figure 2-8 License agreement

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

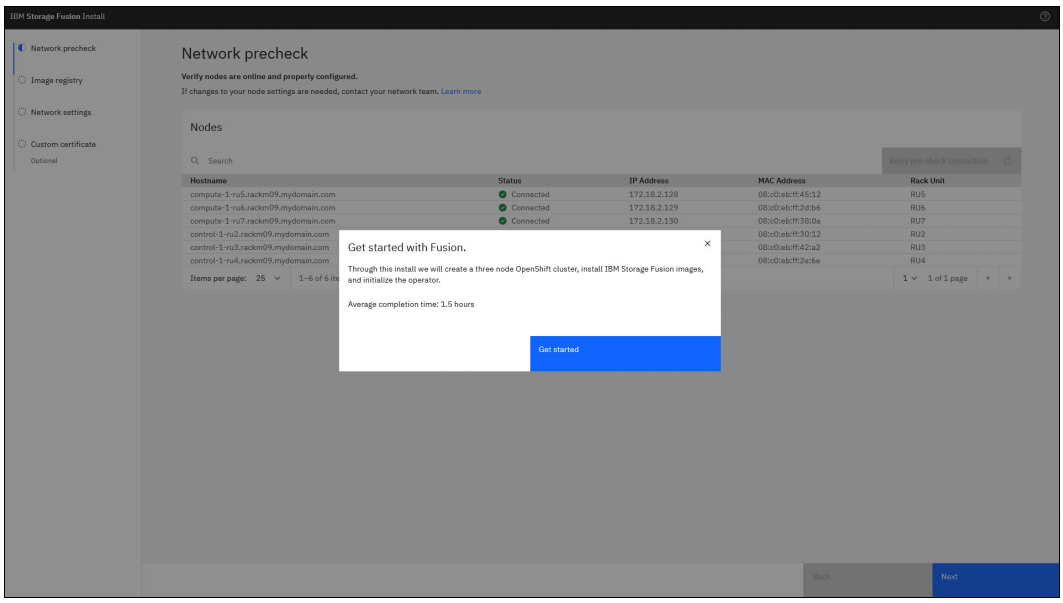

Figure 2-9 Network precheck

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 on page 11.

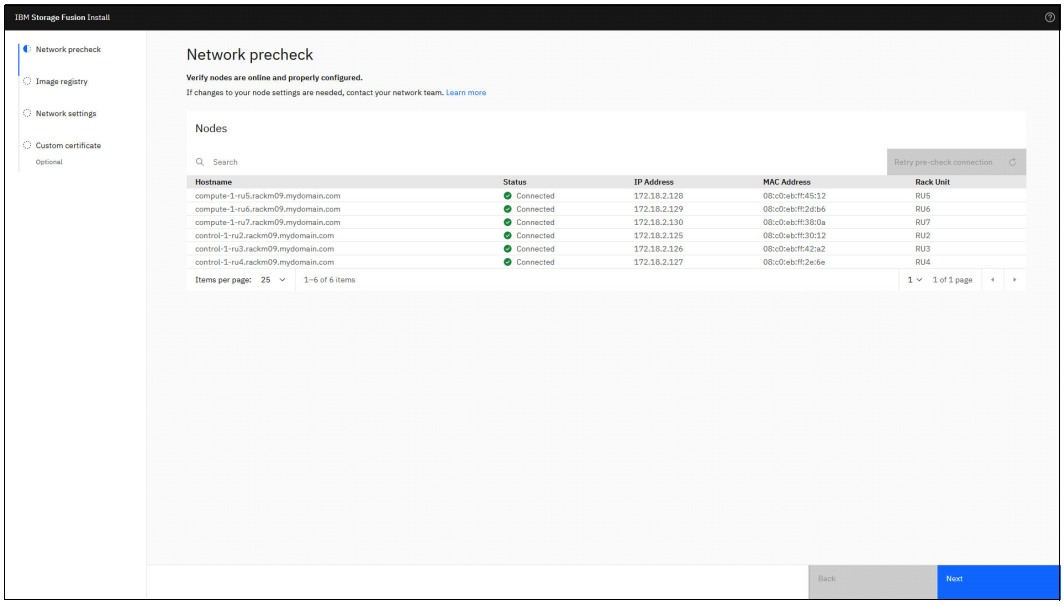

Figure 2-10 Network precheck details

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

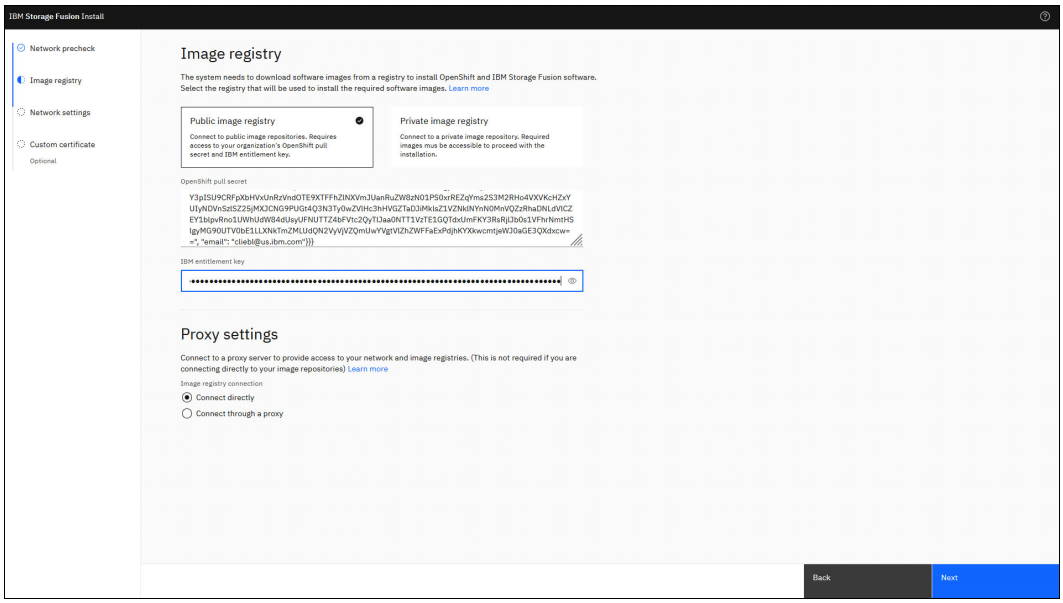

Figure 2-11 Image registry

6. In the **Network Settings** page as shown in Figure 2-12 on page 12, choose the **Yes** option to use Global Data Platform to enable Metro Disaster Recovery and choose **First Cluster** for the question **Is this the first or second cluster in the Metro DR relationship?**

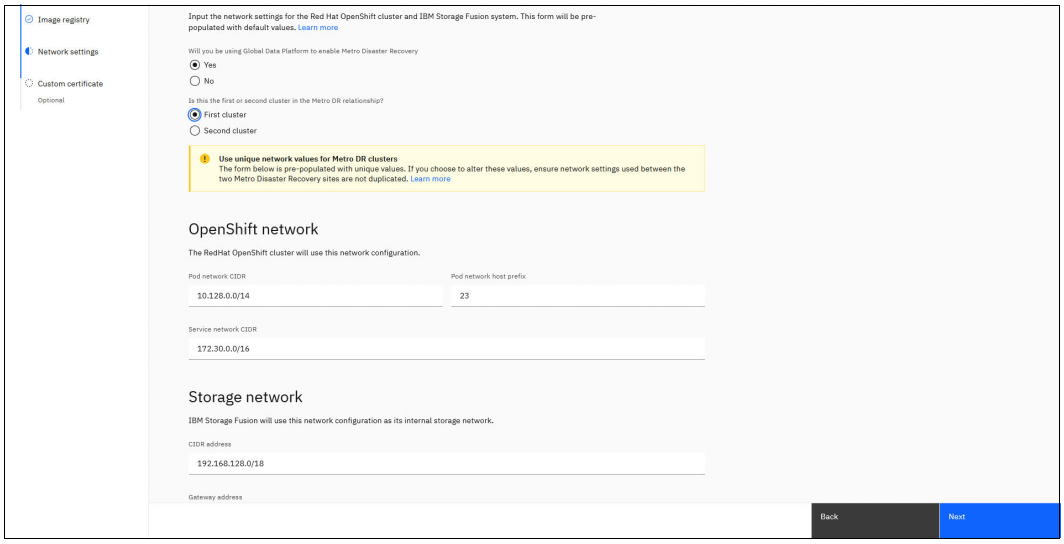

Figure 2-12 Network settings

Also, as shown in Figure 2-13 **Network settings**, enter the details of the OpenShift network and storage network.

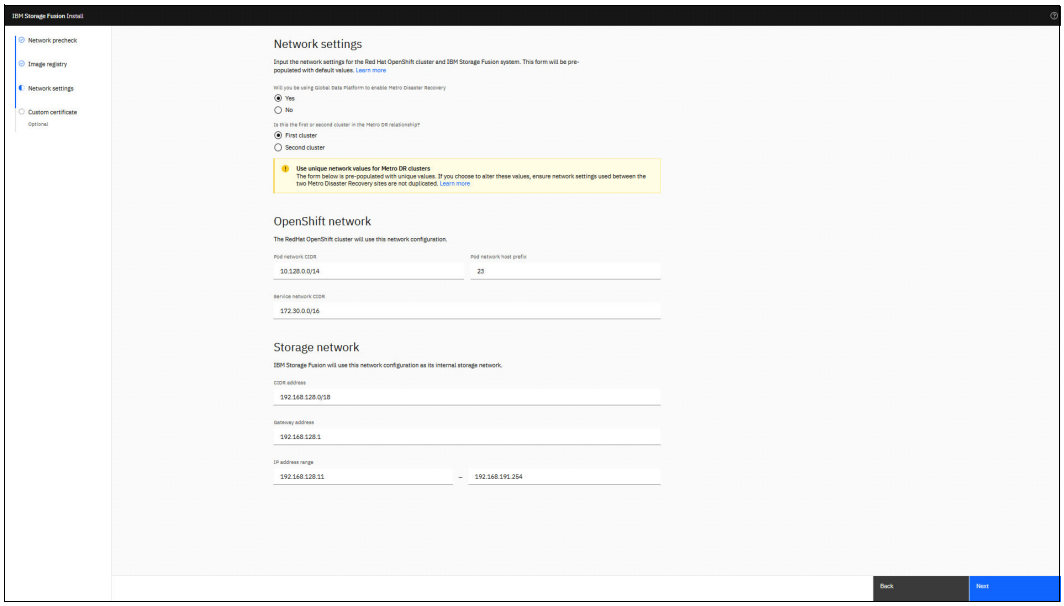

Figure 2-13 Network settings

7. On the **Custom certificate** page as shown in Figure 2-14 on page 13, provide the details of your organization certificate, if any or else set the **Configure Now** to **No** and click **Finish** to start the installation as shown in Figure 2-15 on page 13.

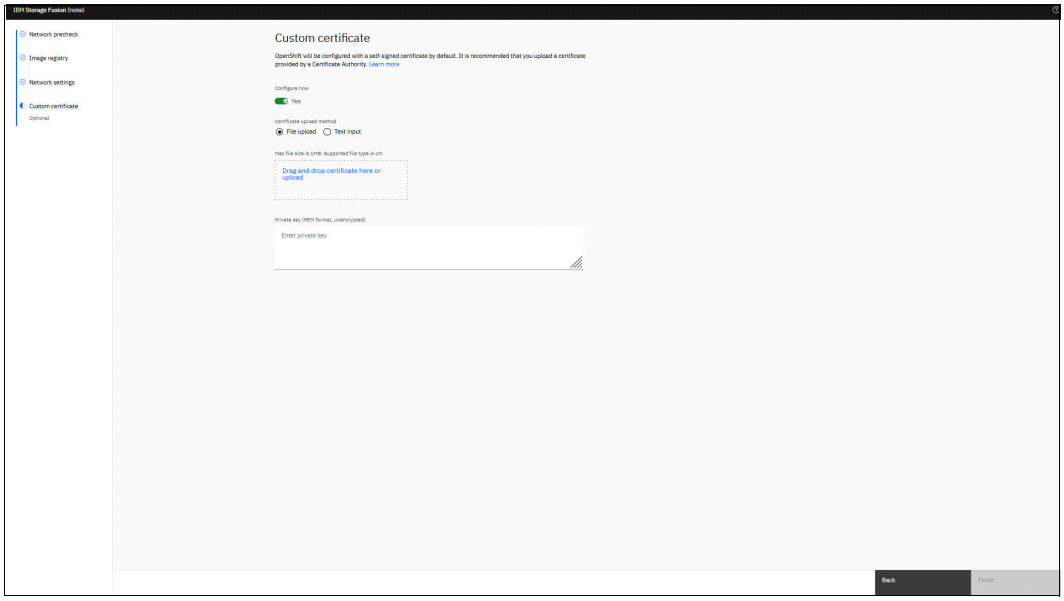

Figure 2-14 Custom certificate using your organization's certificate

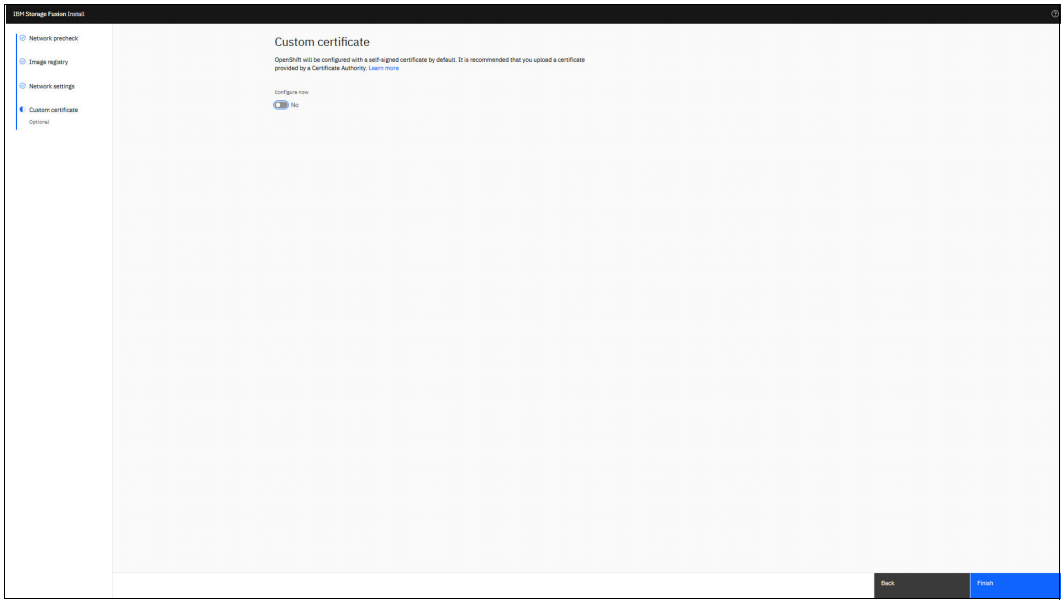

*Figure 2-15 Custom certificate with selection of **No** to **Configure now***

8. Once the Install completes, the button to download the password and CoreOS key is in the enabled state.
9. Click the button **Download Password and CoreOS key** shown in Figure 2-16 on page 14.
10. Now, click on the **IBM Storage Fusion** button to proceed with the install.

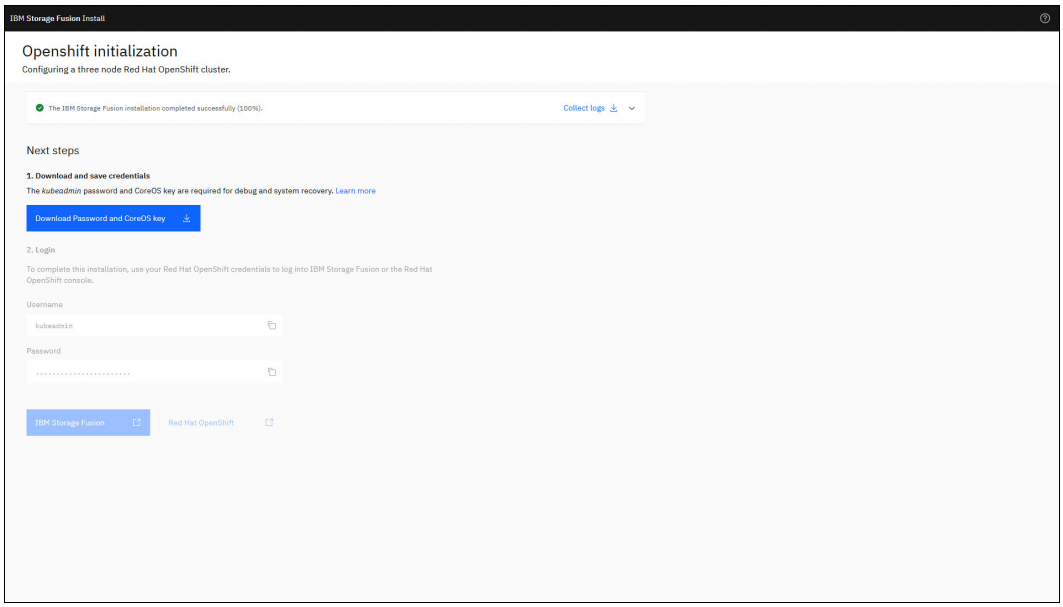

Figure 2-16 OpenShift Initialization

Once the prior steps are completed, a three-node OpenShift Cluster is deployed and IBM Storage Fusion software is installed on it.

Click on **IBM Storage Fusion** button to launch the IBM Storage Fusion user interface as shown in Figure 2-17.

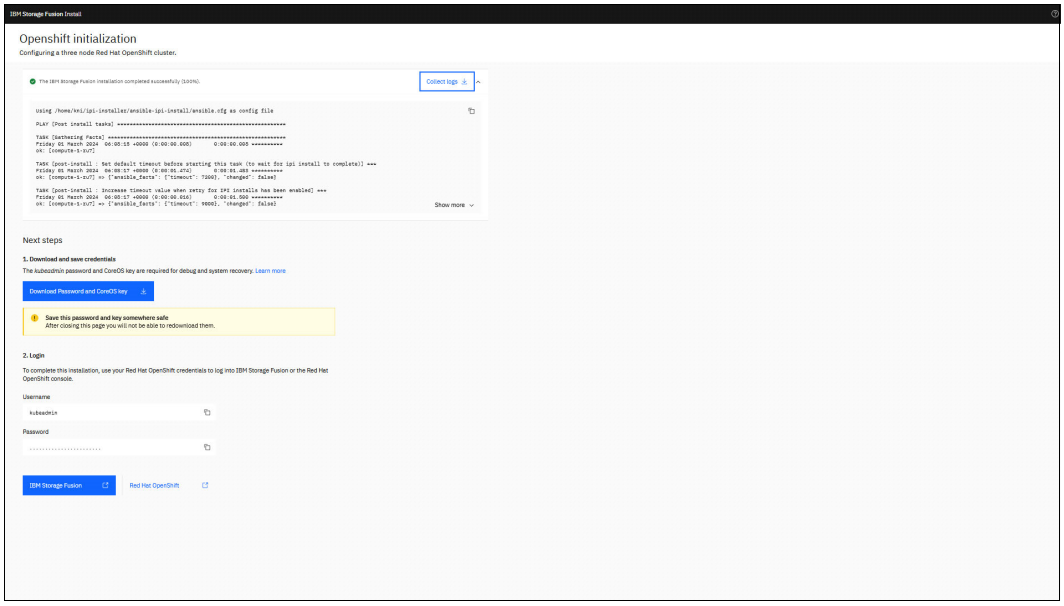

Figure 2-17 OpenShift Initialization - launch IBM Storage Fusion user interface

Addition of remaining nodes to the OpenShift Cluster

This section describes the steps for adding the remaining nodes to the OpenShift cluster.

From the left menu, navigate to the **Infrastructure** → **Nodes** page shown in Figure 2-18 on page 15.

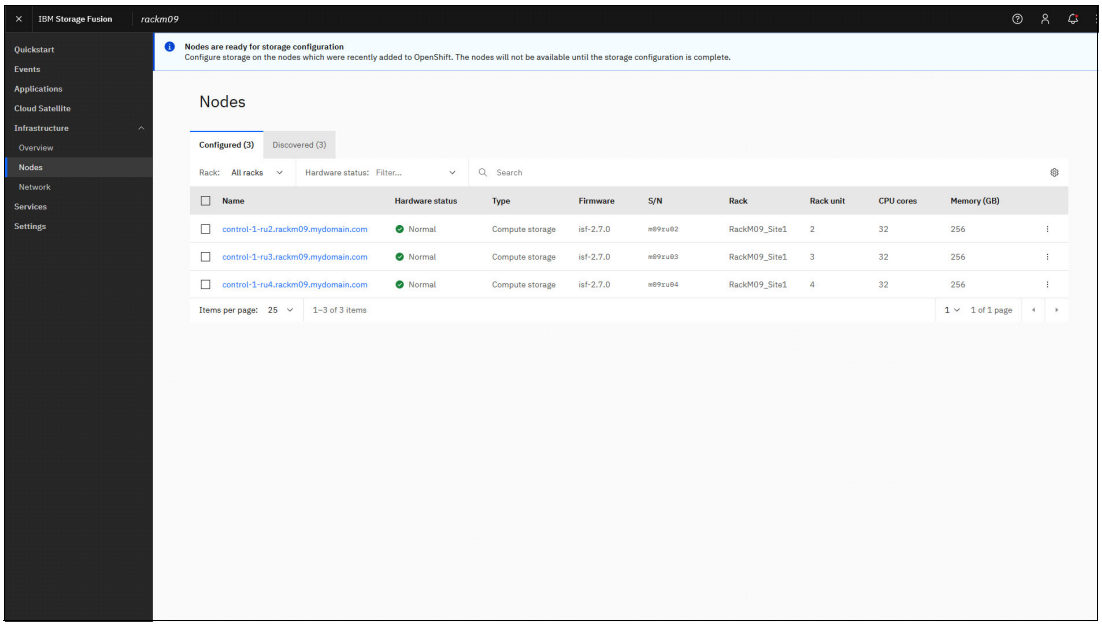

Figure 2-18 Nodes Page - Configured Tab

Click on the **Discovered** tab shown in Figure 2-19.

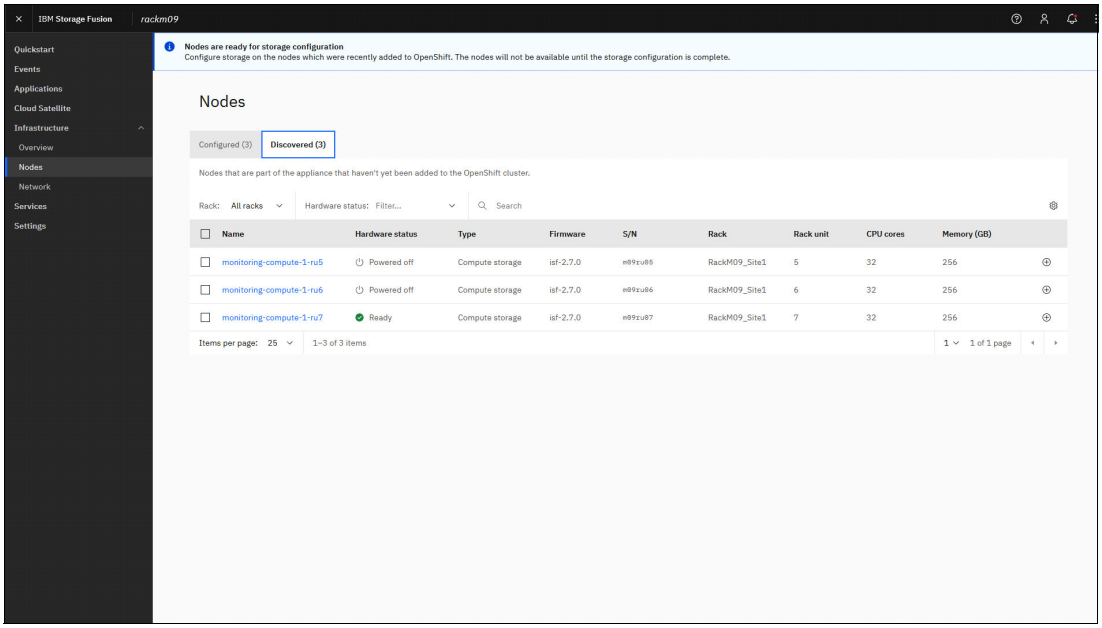

Figure 2-19 Nodes page - Discovered Tab

Select all the nodes from the **Discovered** tab shown on Figure 2-20 on page 16.

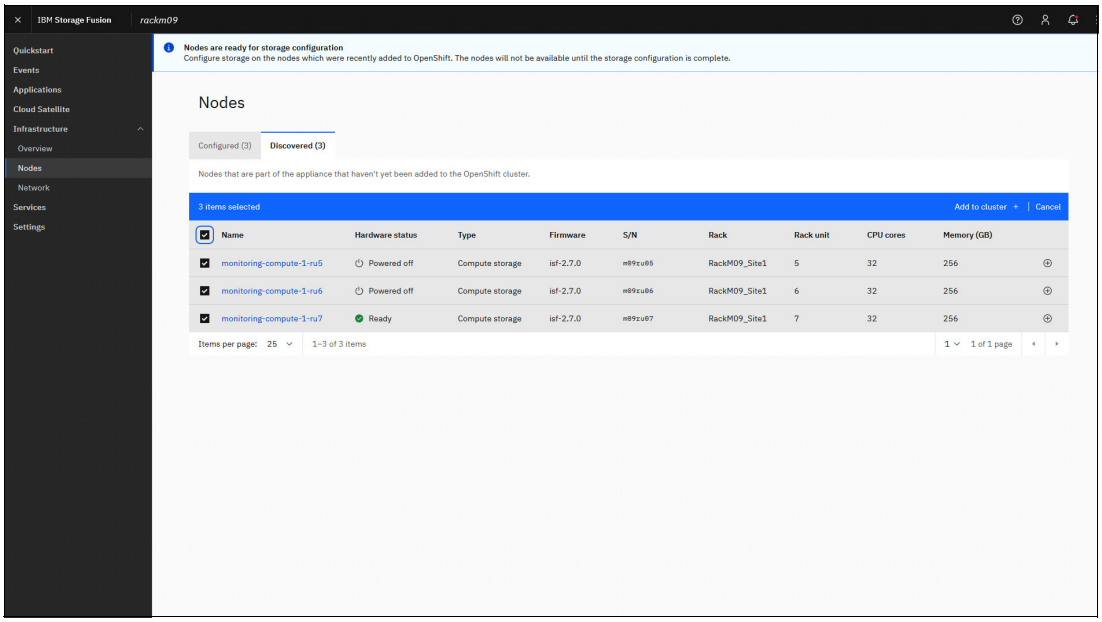

Figure 2-20 Nodes page - Discovered Tab

Click **Add to Cluster** button. This opens up a dialog box listing the nodes.

Click on **Add** button to begin the node addition process as shown in Figure 2-21.

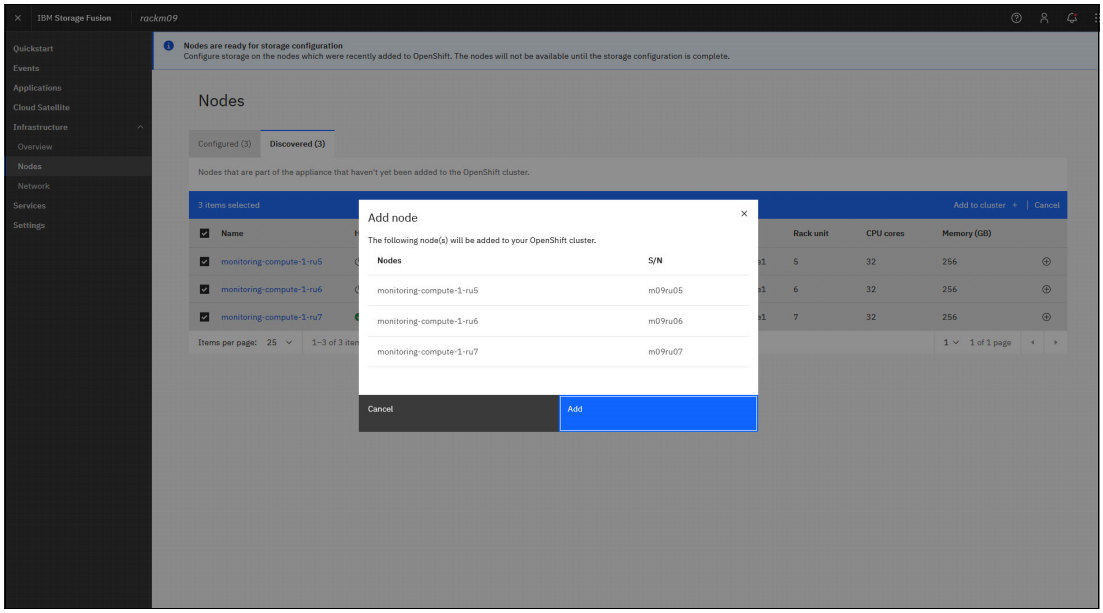

Figure 2-21 Nodes page - Add Node Dialog Box

As shown in Figure 2-22 on page 17, the nodes are moved from **Discovered** page to **Configured** page. The progress of nodes addition is shown on this page.

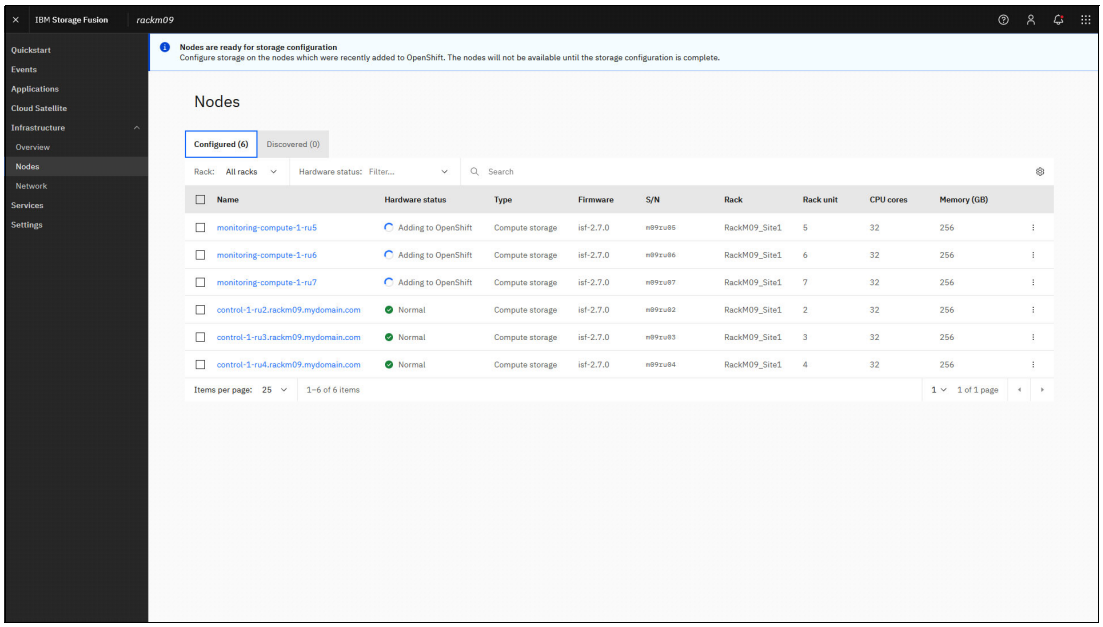

Figure 2-22 Nodes page - Configured Tab

Once the node addition process completes, the nodes status changes to **Normal** as shown in Figure 2-23.

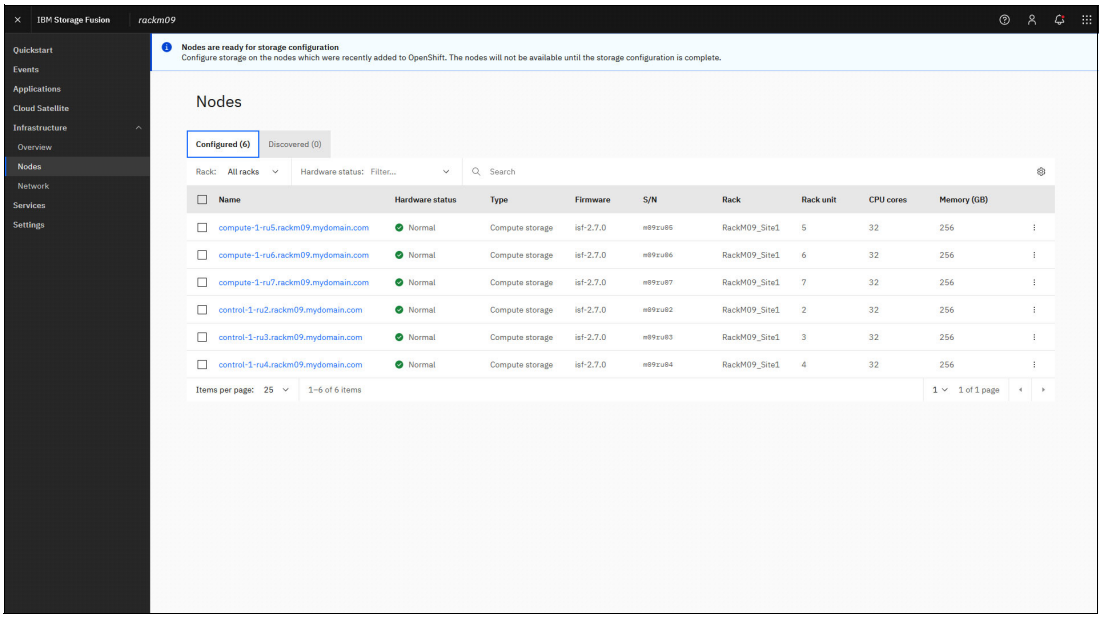

Figure 2-23 Nodes page - Configured Tab with additional nodes

Global Data Platform Installation

From the left menu, navigate to the **Services** page as shown in Figure 2-24 on page 18.

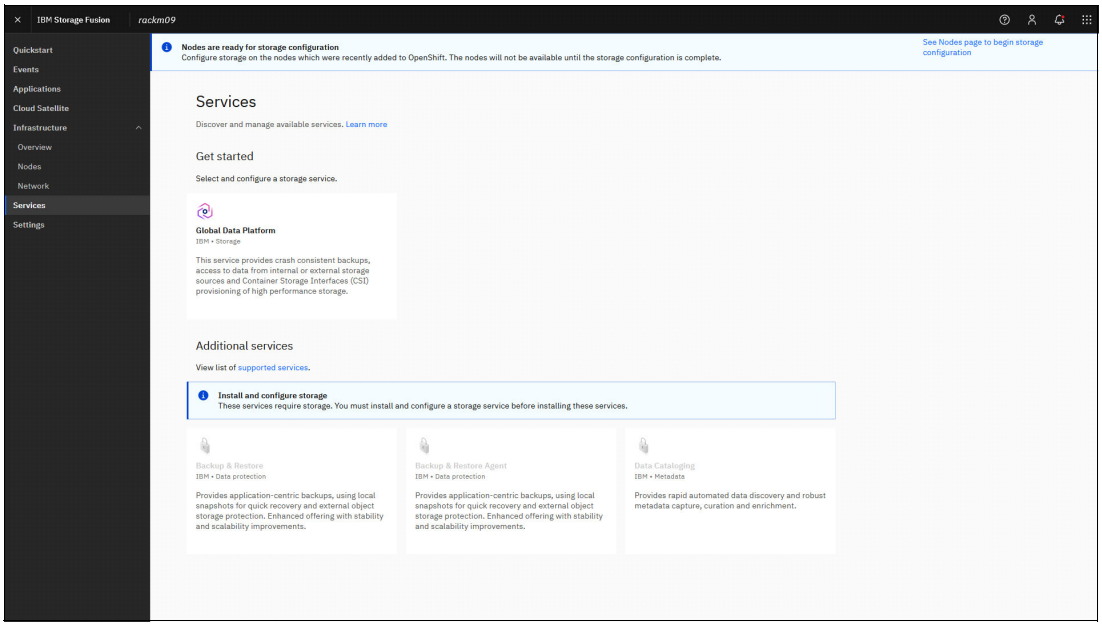

Figure 2-24 Services page

Click on **Global Data Platform** tile, a dialog box opens up. Click on **Install** button as shown in Figure 2-25.

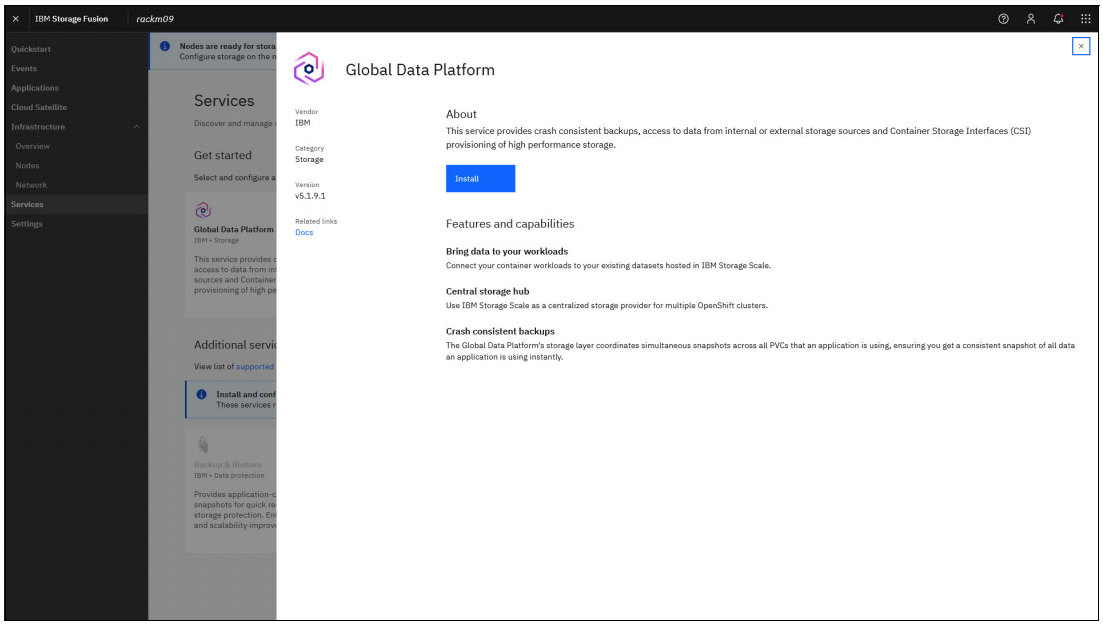

Figure 2-25 Services page - Global Data Platform Dialog Box

The **Install service** dialog box opens up, Click on **Install** button shown in Figure 2-26 on page 19.

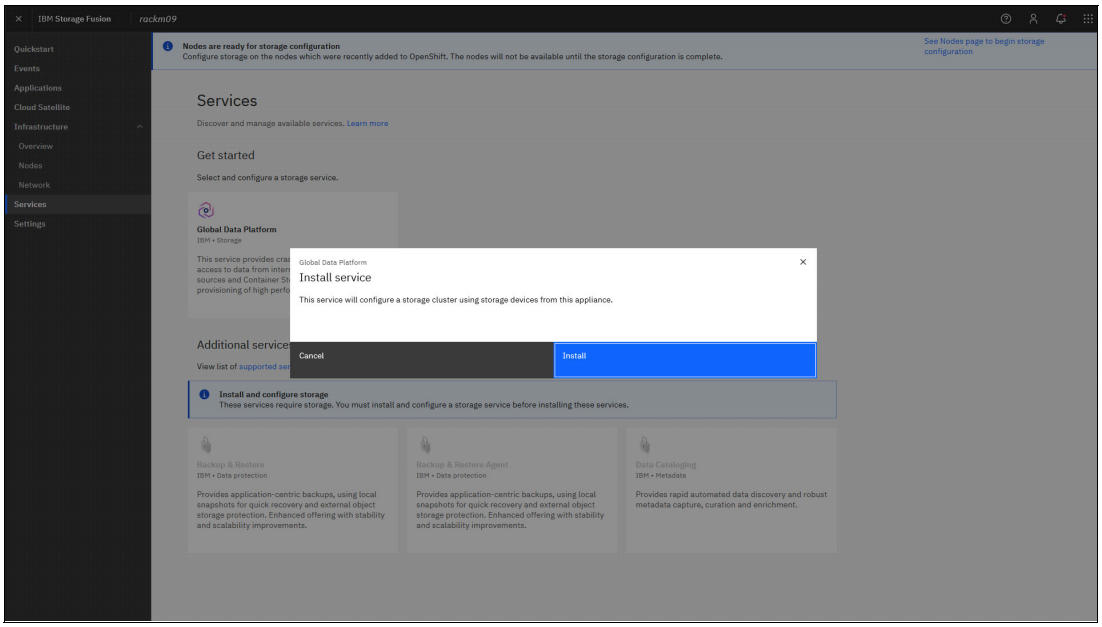

Figure 2-26 Services page - Install Services Dialog Box

Global Data Platform service installation is in progress as shown in Figure 2-27.

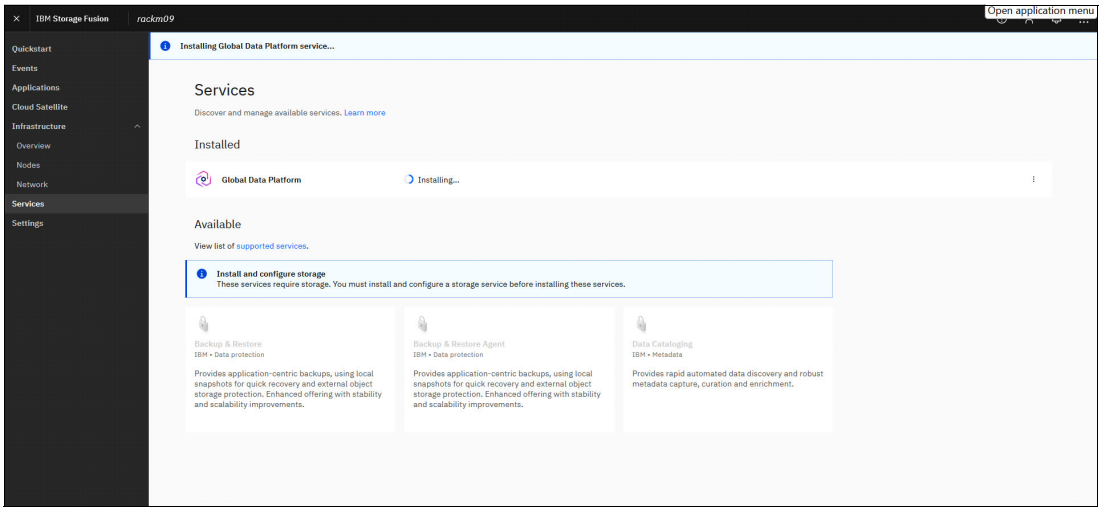

Figure 2-27 Services page - Global Data Platform Installation Progress

Once the installation is complete, the **Global Data Platform service** shows **Healthy** as shown in Figure 2-28 on page 20.

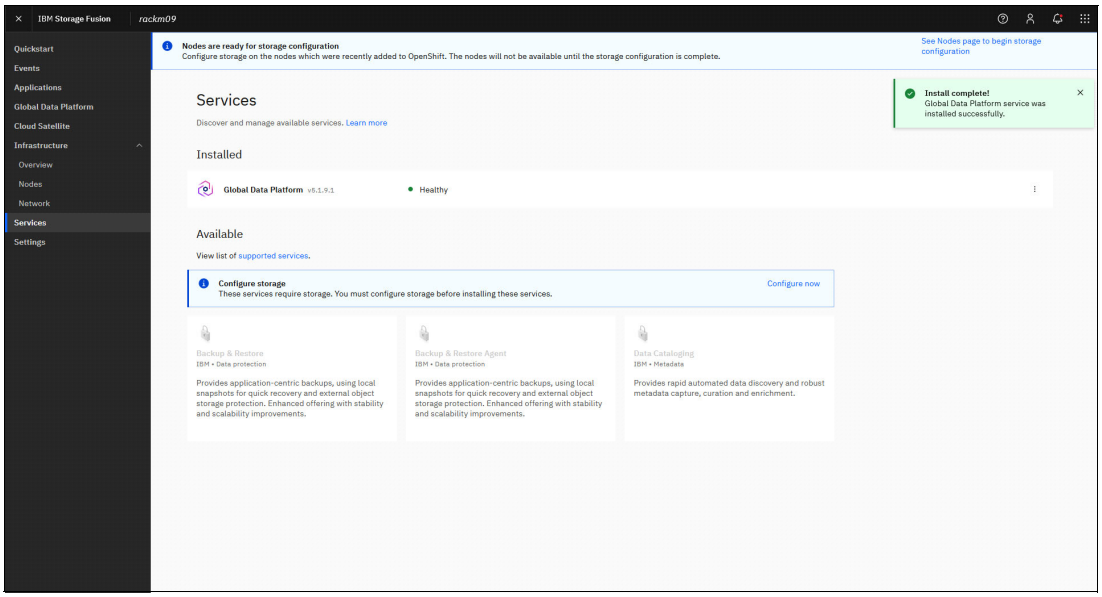

Figure 2-28 Services page – Global Data Platform Installation Completion

Global Data Platform Configuration

In the left menu, navigate to the **Services** page, click on **Configure now** as shown in Figure 2-29.

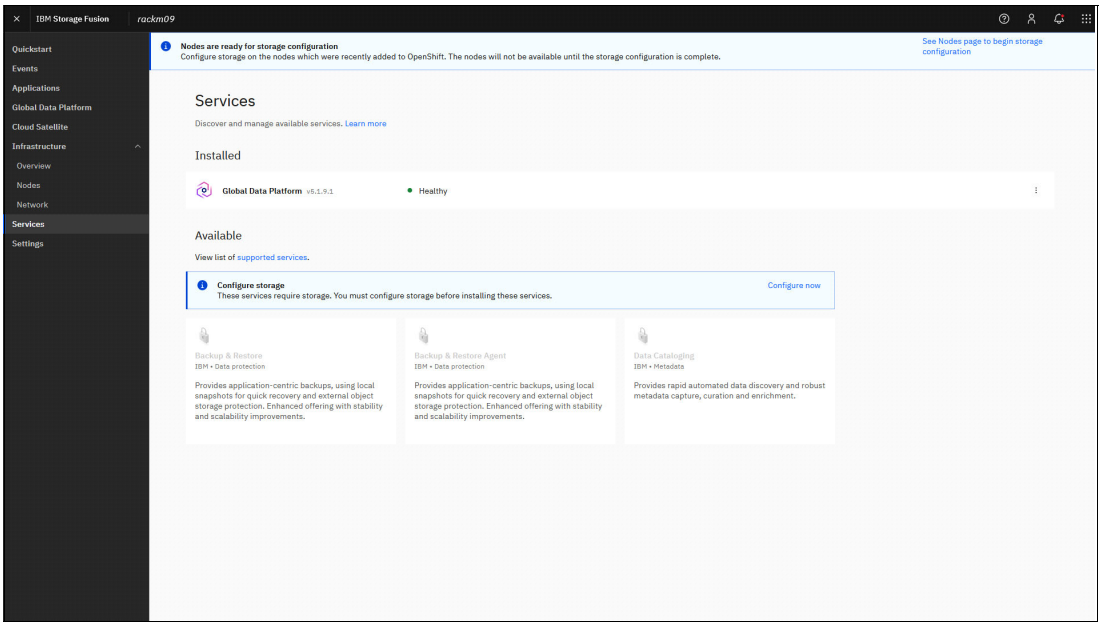

Figure 2-29 Services page

Selecting **Configure now** opens up the **Global Data Platform** page as shown in Figure 2-30 on page 21.

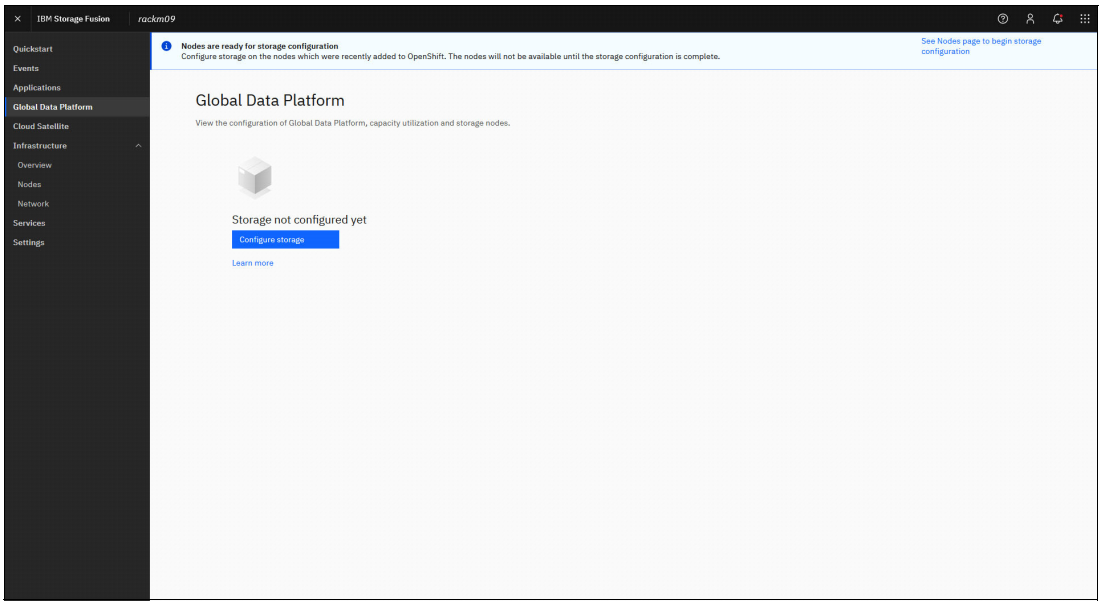

Figure 2-30 Global Data Platform Page

Click **Configure storage** and toggle on **Set up Disaster Recovery** as shown in Figure 2-31.

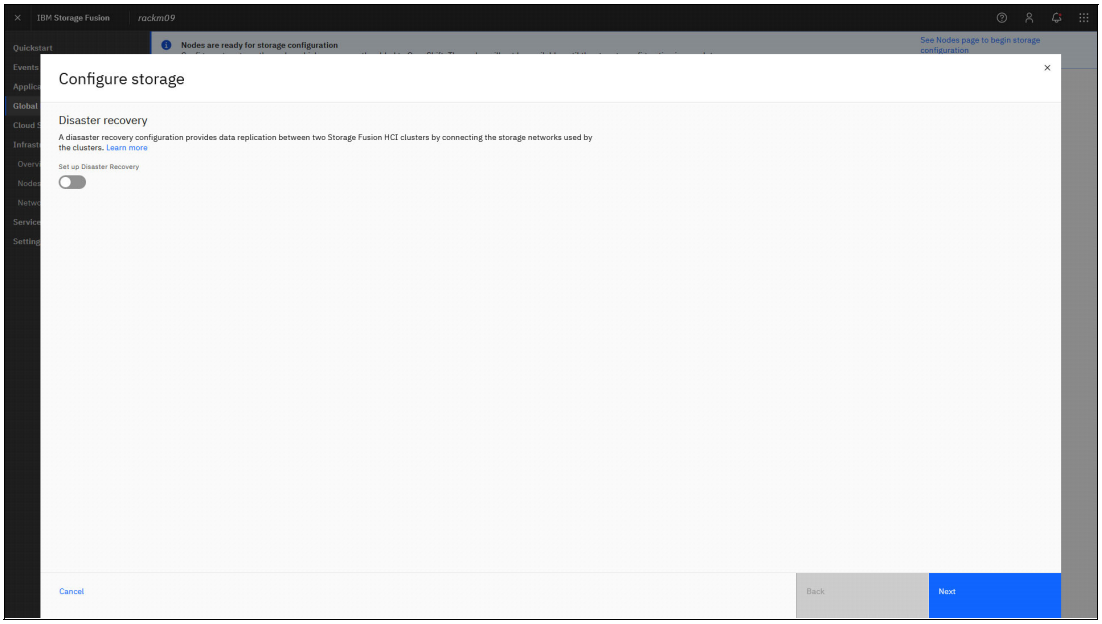

Figure 2-31 Global Data Platform Page - Configure Storage

Select **Metro** configuration and click **Next** as shown in Figure 2-31.

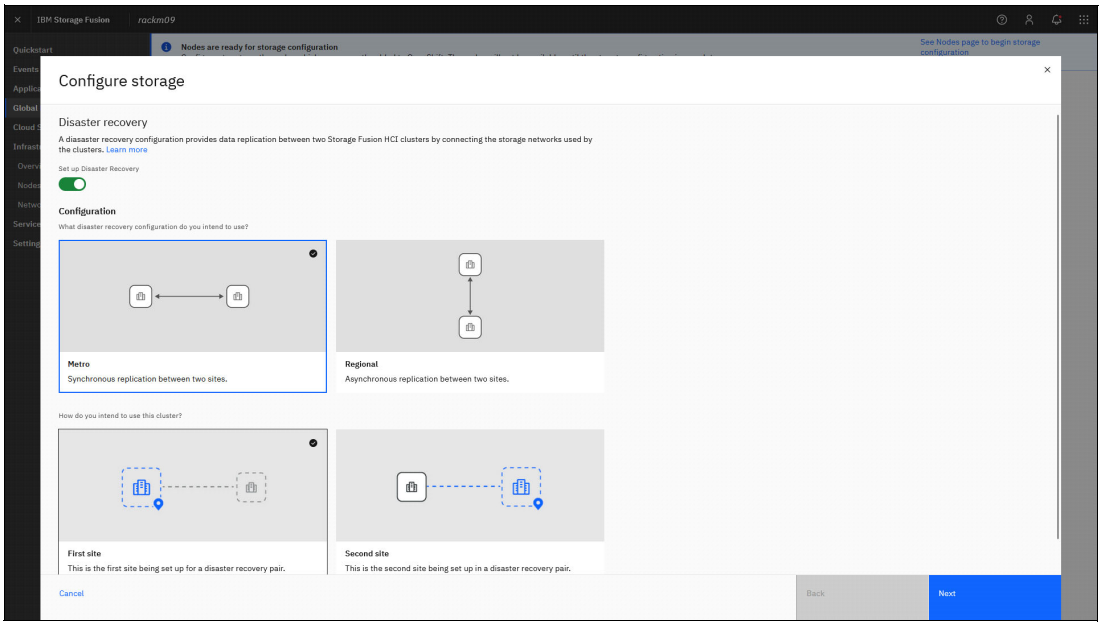

Figure 2-32 Global Data Platform Page - Configure Storage

Select the Recovery Group as per the nodes available in the Fusion HCI appliance. Click **Show advanced section** to select the Block size as shown in Figure 2-33.

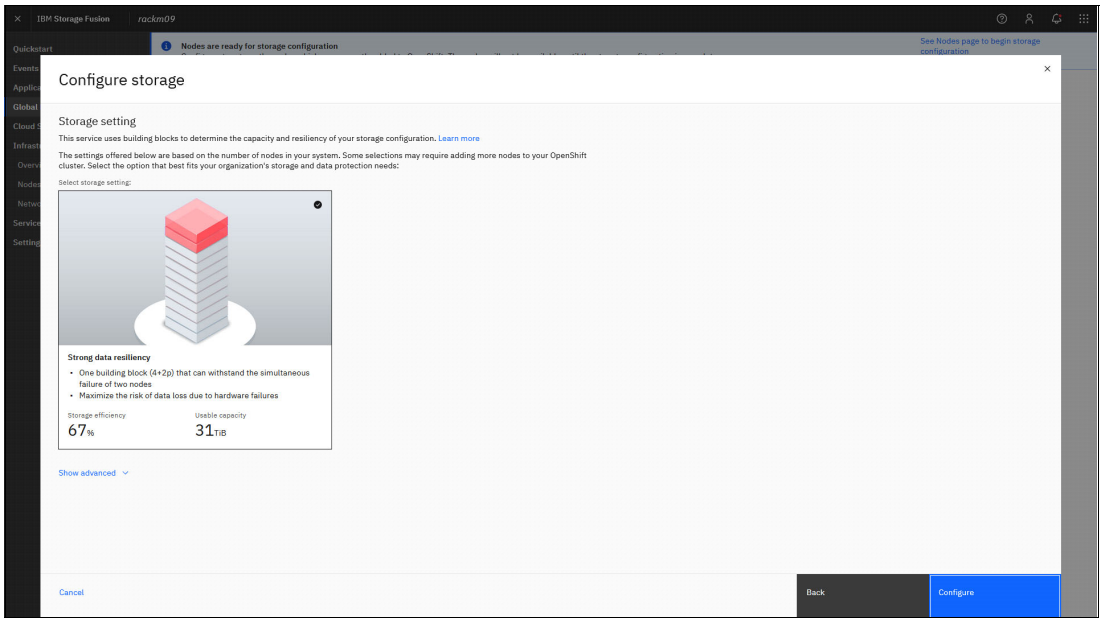

Figure 2-33 Global Data Platform Page - Configure Storage

Select the block size and click on **Configure** as shown in Figure 2-34 on page 23.

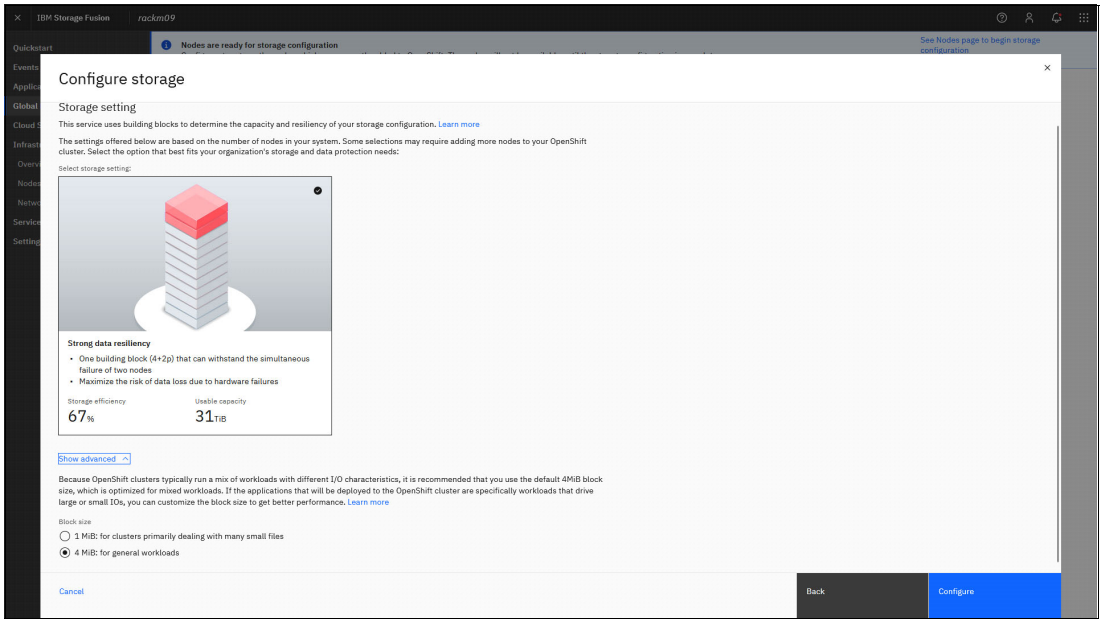

Figure 2-34 Global Data Platform Page – Configure Storage

The message **Configuring Storage** appears as shown in Figure 2-35.

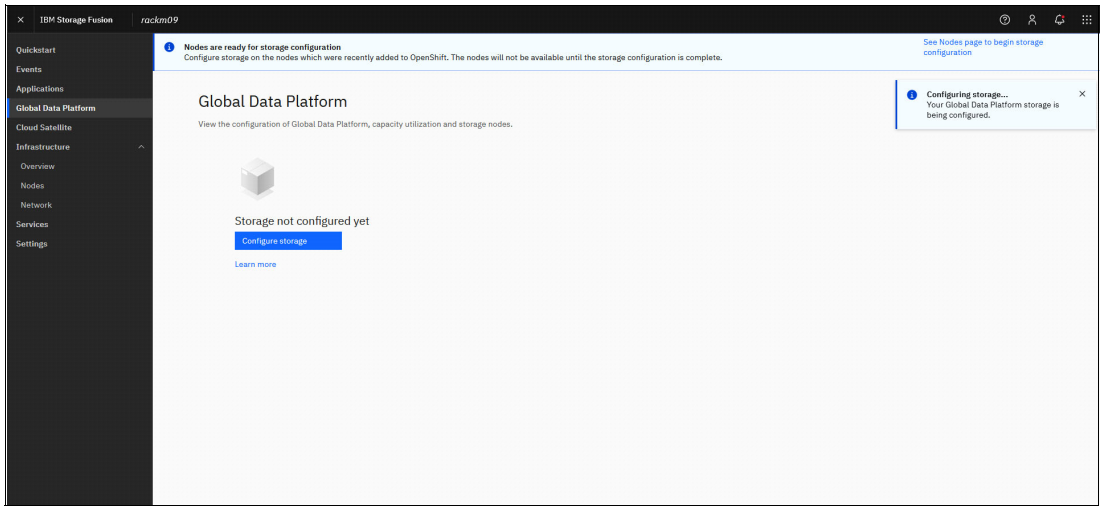

Figure 2-35 Global Data Platform Page

The **Global Data Platform** configuration started as shown in Figure 2-36 on page 24.

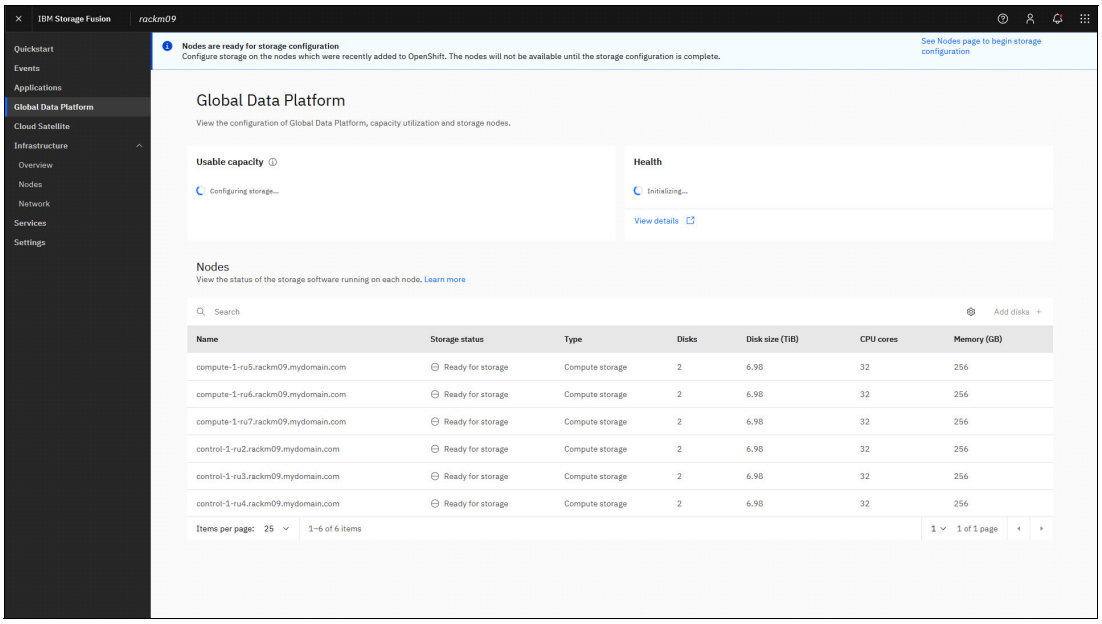

Figure 2-36 Global Data Platform Page - Configure Storage in progress

As soon as the configuration completes, the **Global Data Platform service** page shows the storage cluster health as **OK** and the nodes shows the Storage Status as **Normal** as shown in Figure 2-37.

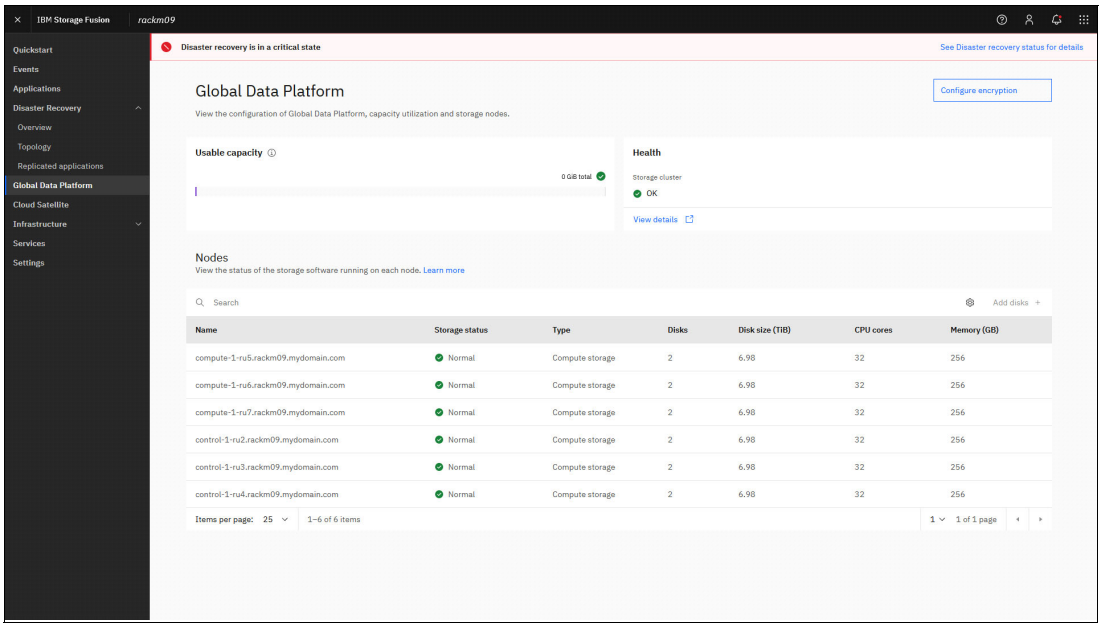

Figure 2-37 Global Data Platform Page - Configure Storage completion

Now, we are ready to use IBM Fusion as shown in Figure 2-38 on page 25.

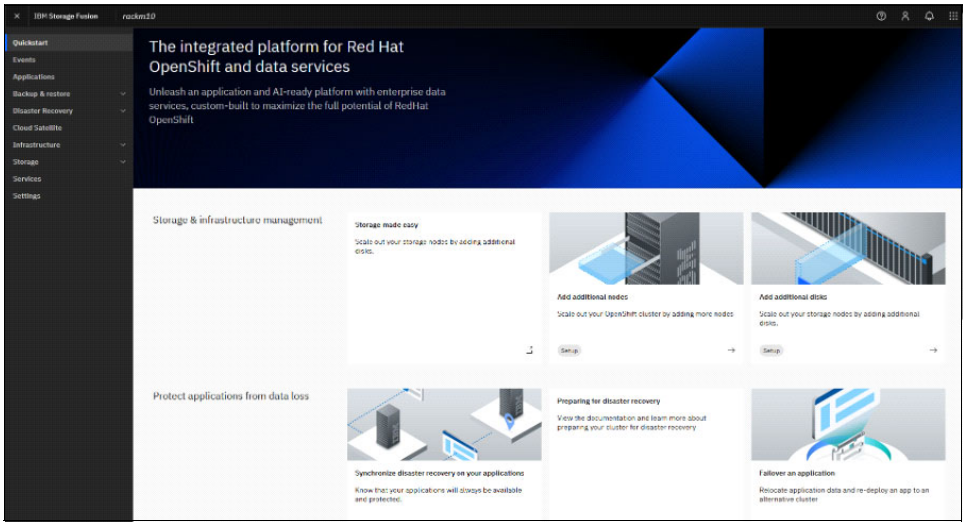

Figure 2-38 IBM Fusion Quick start page

2.1.3 Site2 Installation

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

Network configuration

The network configuration steps for Site2 are similar to the steps for Site1. Follow 2.1.2, “Site1 installation for Deployment type #2” on page 6.

Save the URL given at the end of the network configuration step as shown in Figure 2-39.

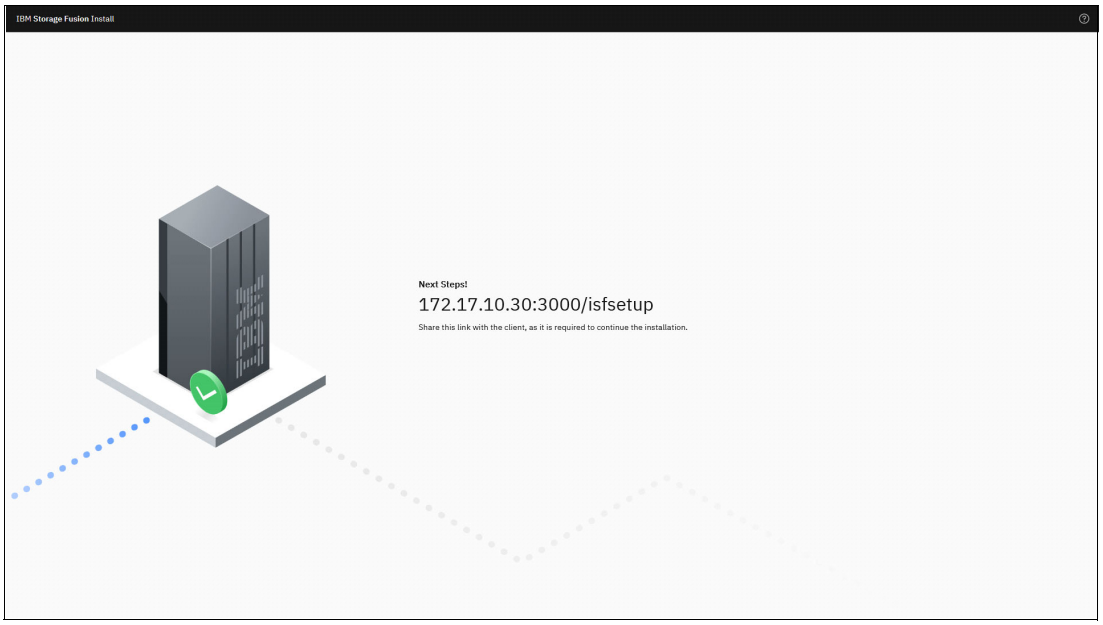

Figure 2-39 Site2 - Network Configuration - Next Steps

IBM Storage Fusion installer

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

1. To proceed with IBM Storage Fusion install, use the URL as shown in Figure 2-39 on page 25 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 **Continue** as shown in Figure 2-40.

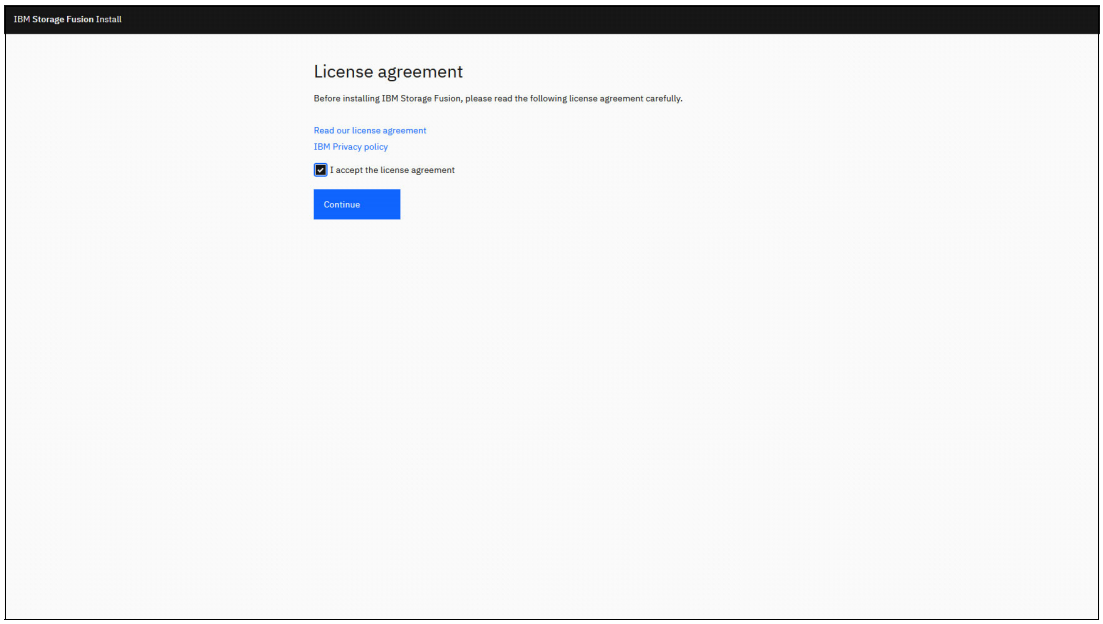

Figure 2-40 License Agreement

Click on **Get Started** in the **Network Precheck** page shown in Figure 2-41.

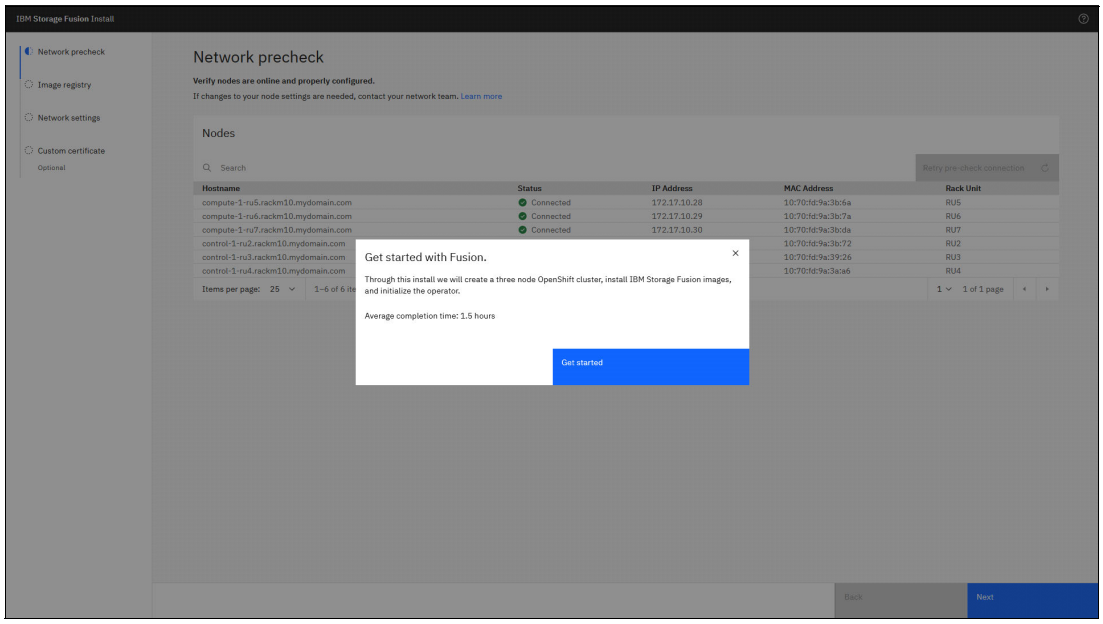

Figure 2-41 Network precheck - Getting Started

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 2-42 on page 27.

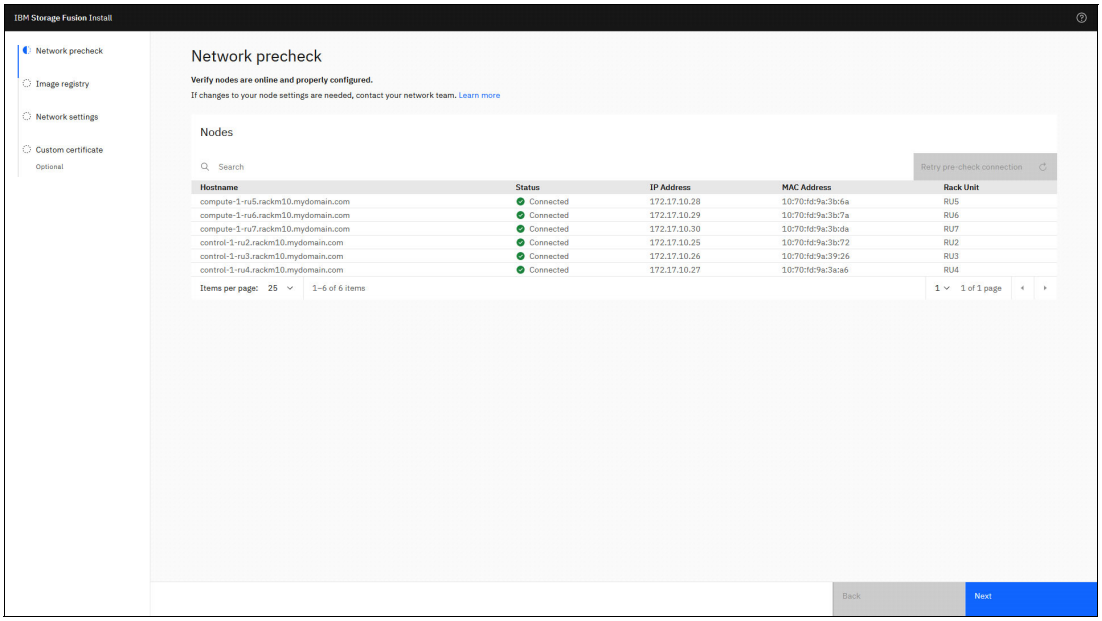

Figure 2-42 Network precheck

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

Click **Next**. Select the image registry as shown in Figure 2-43. Enter the details and click **Next**.

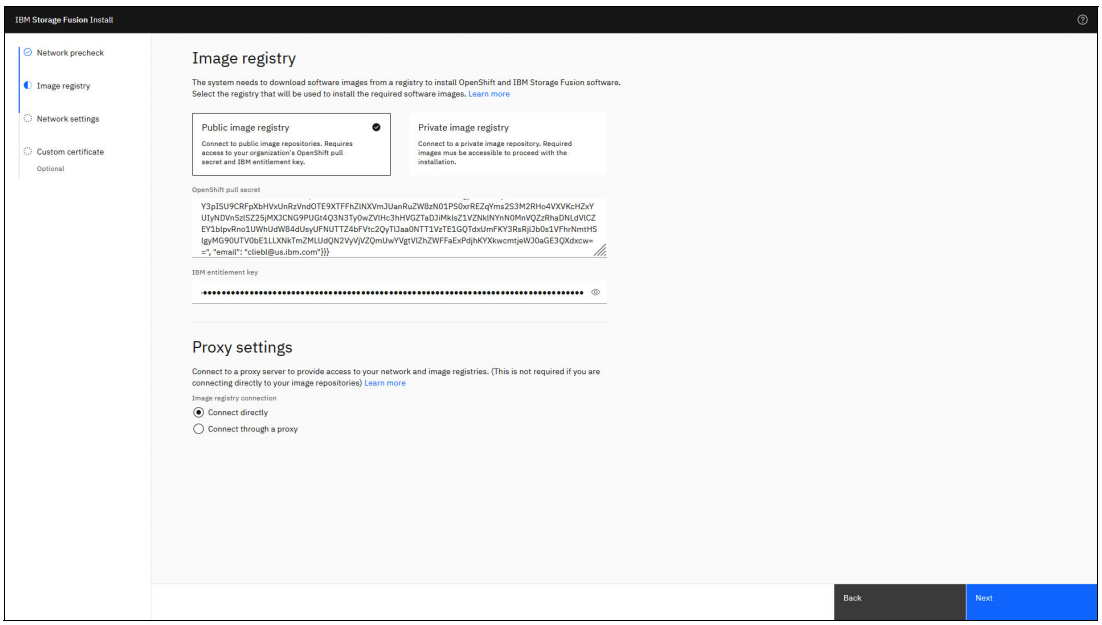

Figure 2-43 Image registry

In the **Network Settings** page as shown in Figure 2-12 on page 11, choose the **Yes** option to use **Global Data Platform** to enable Metro Disaster Recovery. For Site2, choose **Second Cluster** for the question **Is this the first or second cluster in the Metro DR relationship?** as shown in Figure 2-44 on page 28.

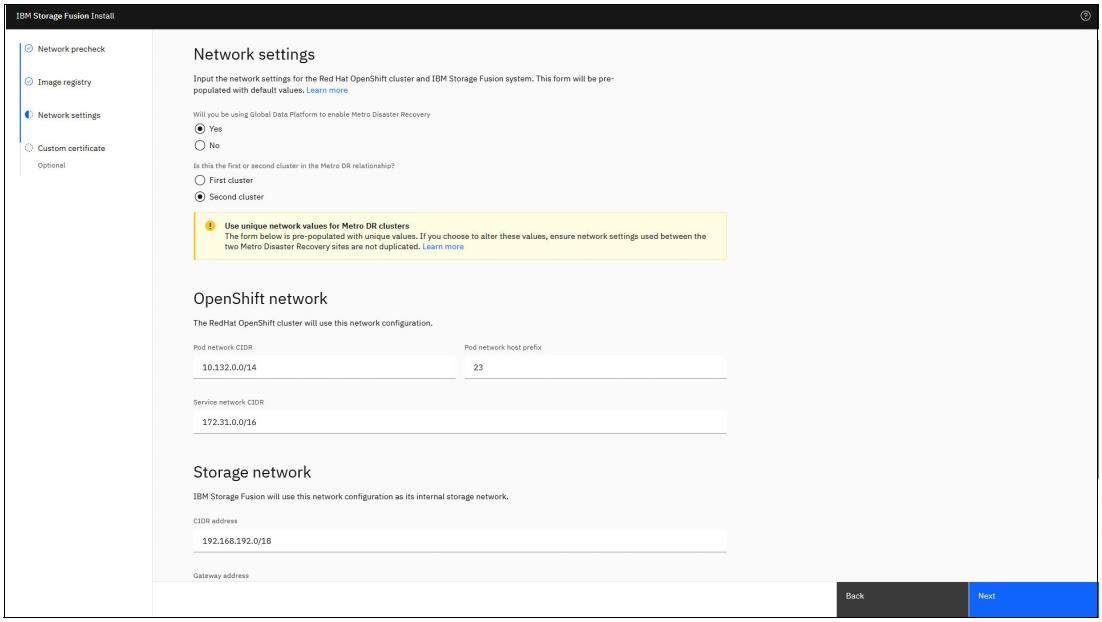

Figure 2-44 Network Settings

Also, as shown in the **Network Settings** page (Figure 2-45) enter the details of the OpenShift network and storage network. Click **Next** to proceed.

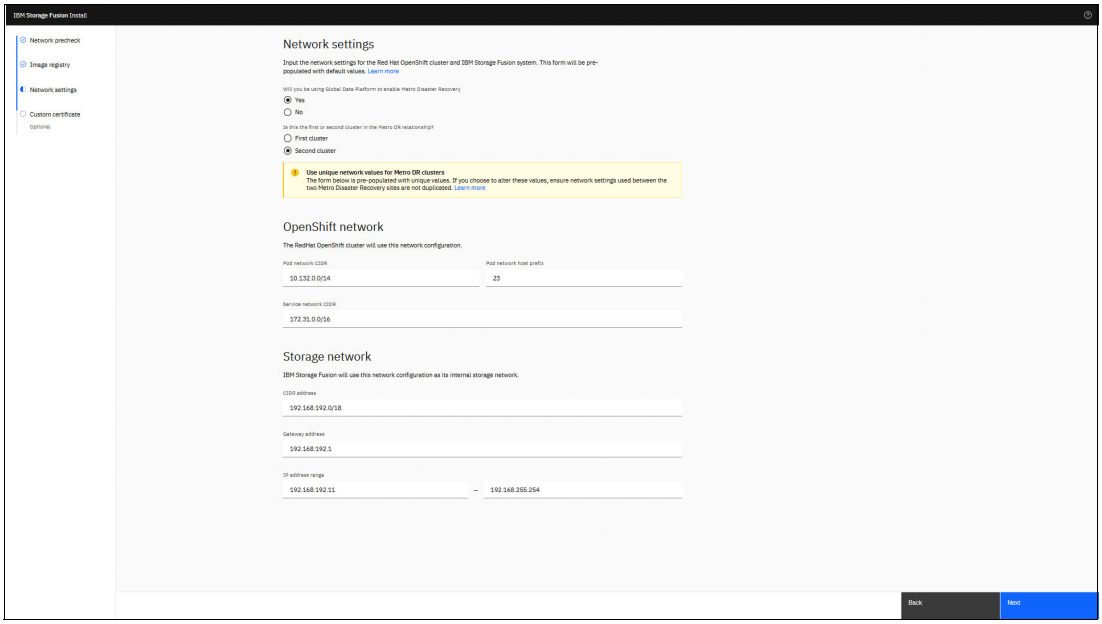

Figure 2-45 Network Settings

On the **Custom certificate** page as shown in Figure 2-46 on page 29, provide the details of your organization certificate, if any. Click **Finish** to start the installation.

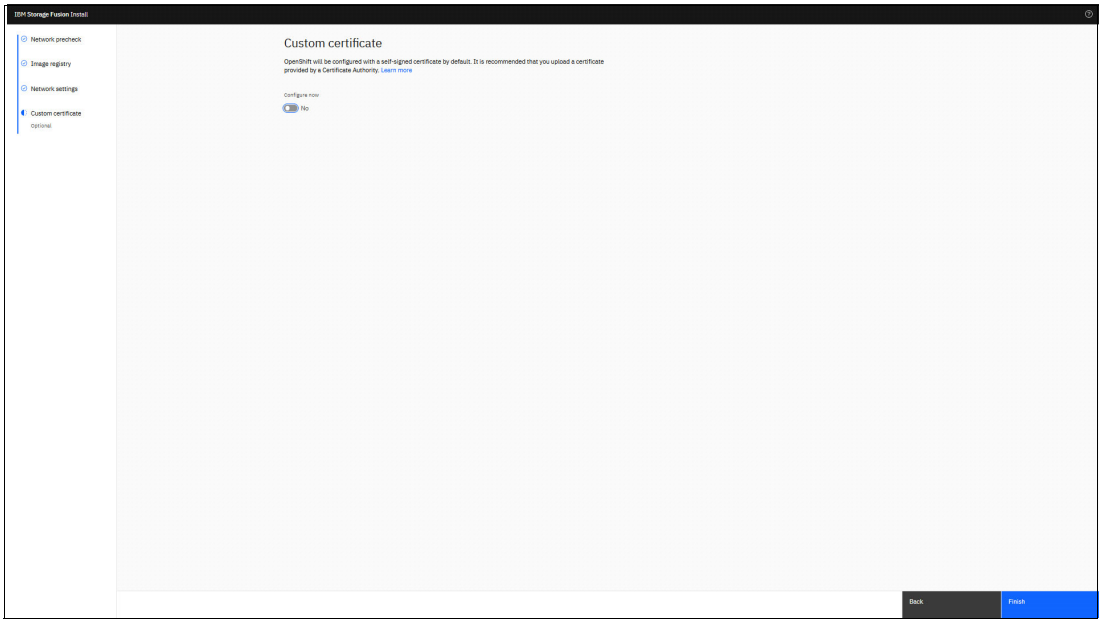

Figure 2-46 Custom certificate

Once the install completes, the message is displayed as shown in Figure 2-47.

Note: IBM Storage Fusion installation completed successfully. **Download password and CoreOS key** is enabled.

Click on **Download Password and CoreOS key** button to download OpenShift credentials. Once the credentials are downloaded, the **IBM Storage Fusion** button is enabled.

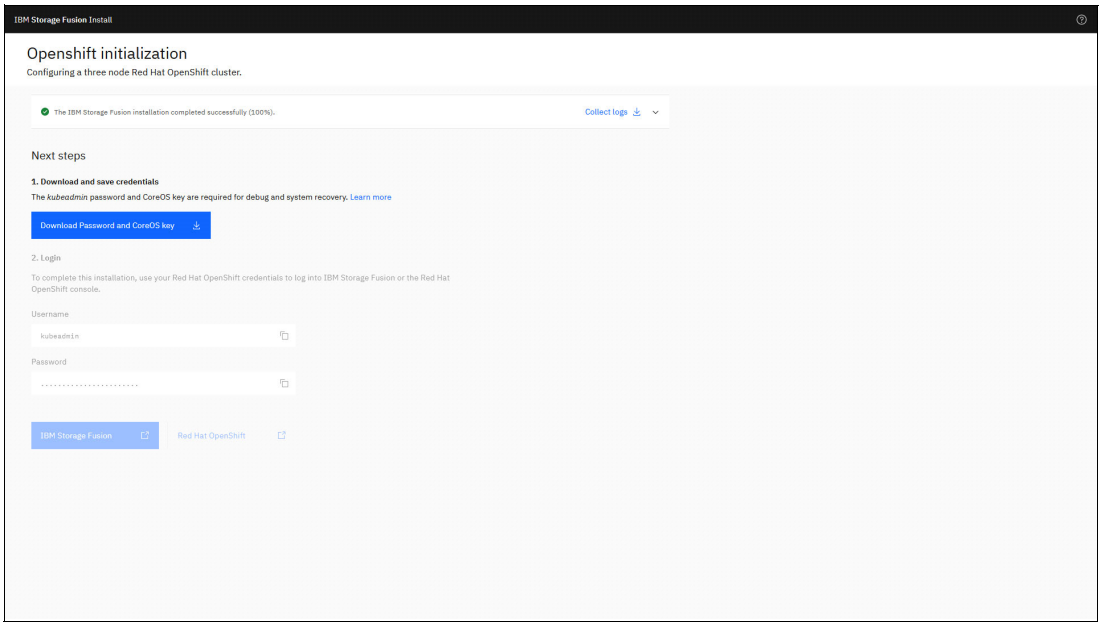

Figure 2-47 OpenShift initialization

Addition of remaining nodes to the OpenShift Cluster

The steps for addition of remaining nodes to the OpenShift cluster for Site1 and Site2 are similar. Follow the steps in “Addition of remaining nodes to the OpenShift Cluster” on page 14.

Global Data Platform Installation

The steps for Global Data Platform Installation for Site1 and Site2 are similar. Follow the steps in “Global Data Platform Installation” on page 17.

Global Data Platform Configuration

This section explains how to set up the Global Data Platform Configuration.

Ensure that the following prerequisites are met.

- IBM Storage Scale on the Site1 is healthy and all IBM Storage Scale core pods are up and running.
- Ensure that the disk count is the same on Site1 **or Standalone appliance** and Site2.
- Ensure that Site2 has the same supported OpenShift Container Platform version as Site1.
- Collection of Connection Snippet from Site1 **or Standalone appliance**.

Log in to IBM Storage Fusion user interface of Site1.

From the left pane menu, navigate to the **Disaster Recovery** → **Overview** page as shown in Figure 2-48.

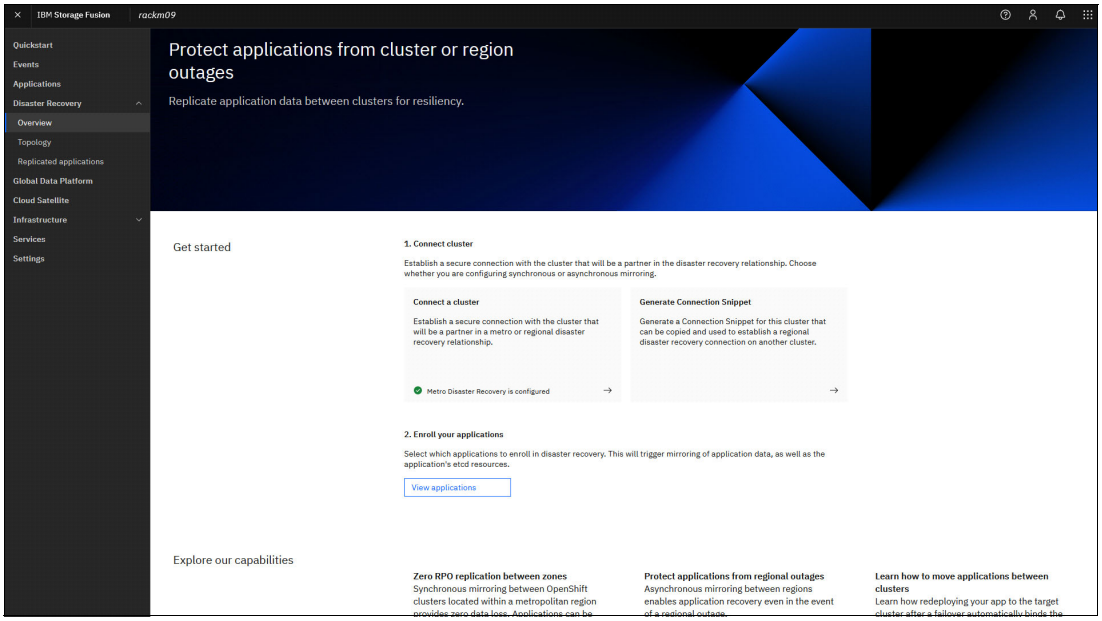

Figure 2-48 Site1 - IBM Storage Fusion User Interface

Click on **Generate Connection Snippet**. It opens up a dialog box as shown in Figure 2-49 on page 31. Click on **Copy snippet** to copy the connection details of Site1.

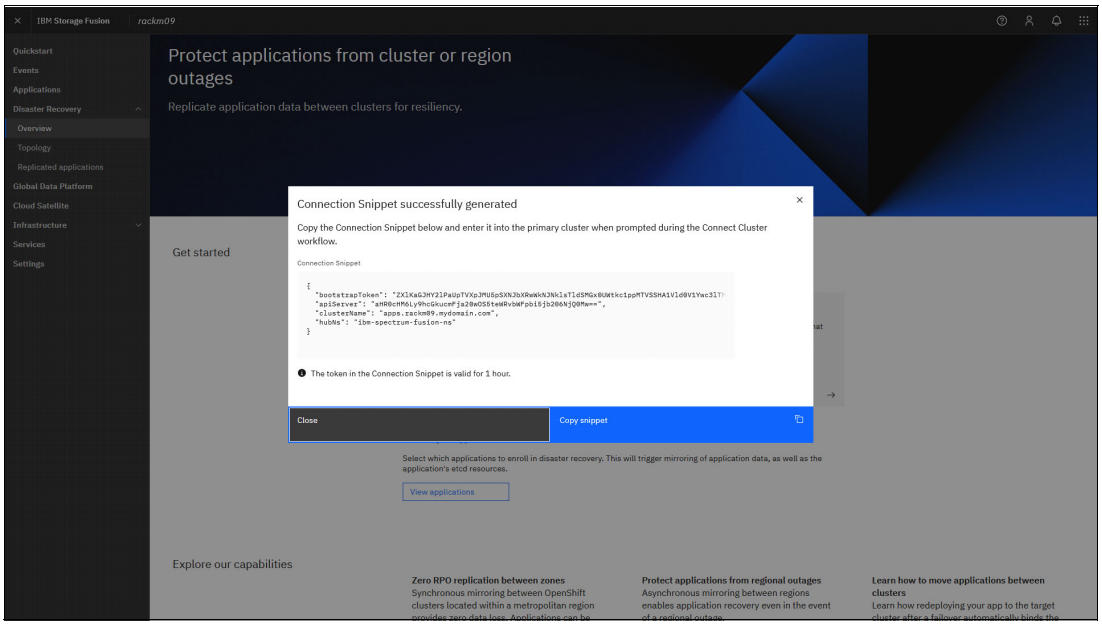

Figure 2-49 Site1 - Disaster Recovery - Overview Page - Connection Snippet

Steps to configure Site2 Storage Configuration

Log in to IBM Storage Fusion user interface of Site2.

In the left pane menu, navigate to **Services**. Click **Configure Now** as shown in Figure 2-50.

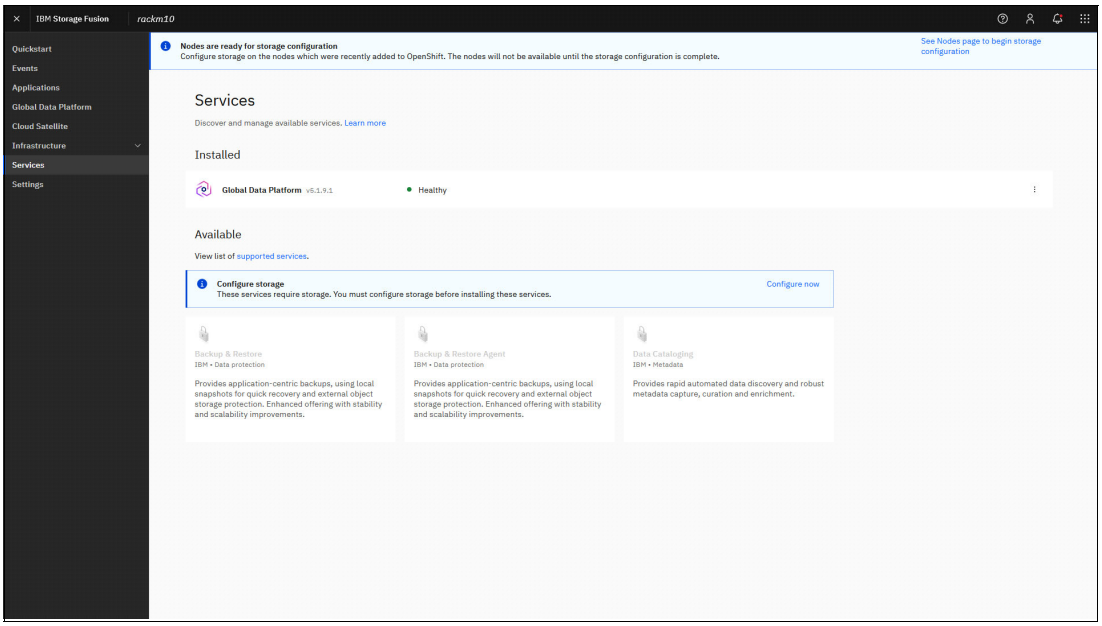

Figure 2-50 Site2 - Services page

The **Global Data Platform** page opens up. Click on **Configure Storage** button as shown in Figure 2-51 on page 32.

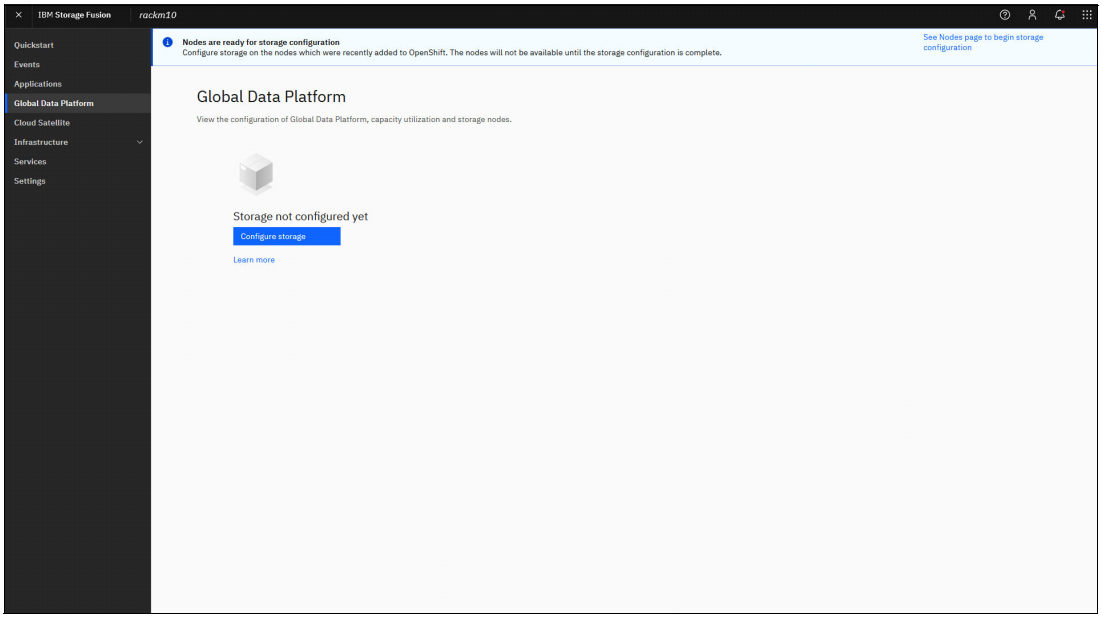

Figure 2-51 Site2 - Global Data Platform

On **Configure storage**, click **Setup Disaster Recovery**. Click the **Next** button as shown in Figure 2-52.

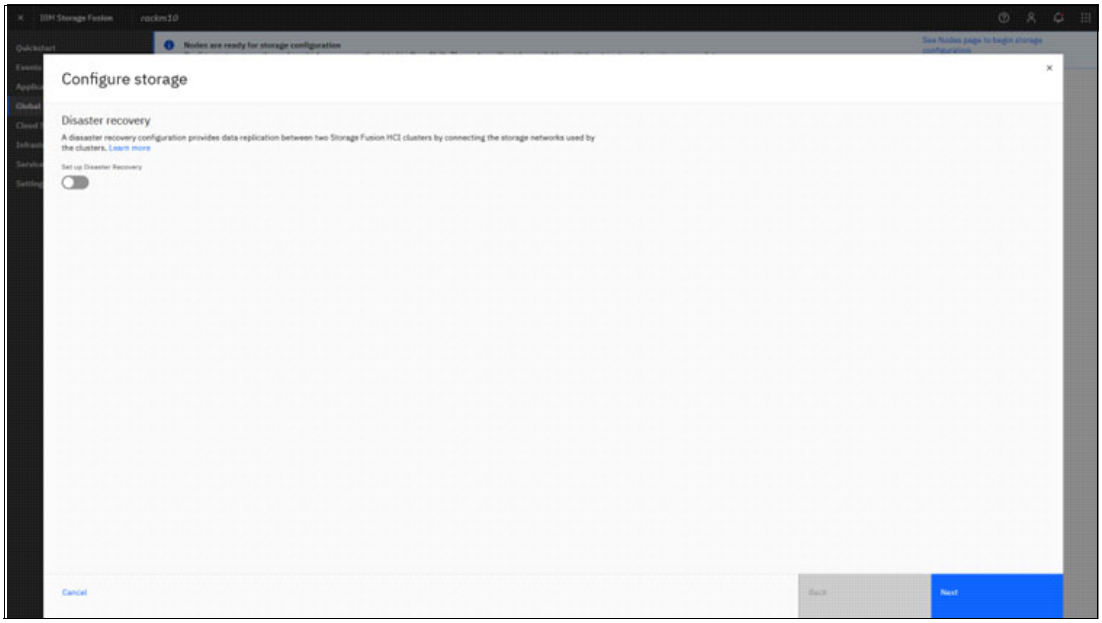

Figure 2-52 Global Data Platform - Configure Storage

On the **Configure storage** page perform the following steps:

- In the **Configure storage** section, select **Metro (**Figure 2-53 on page 33).
- In the **How do you intend to use this cluster** section, select **Second Site**.
- In the **Connect to the first site** section, enter the connection snippet of Site1 as collected in the prerequisites section.
- Enable **Jumbo Frames** as required by your organization.
- Click **Configure** to proceed.

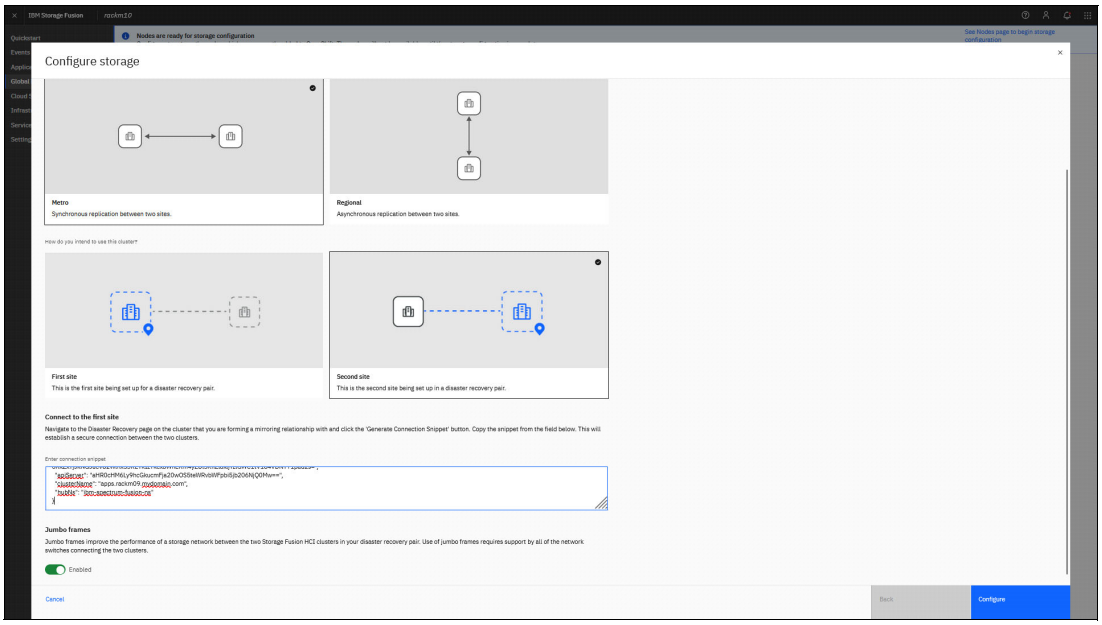

Figure 2-53 Site2 - Global Data Platform - Configure Storage

After selecting **Configure**, the selection box will show the status of **configuring** as shown in Figure 2-54.

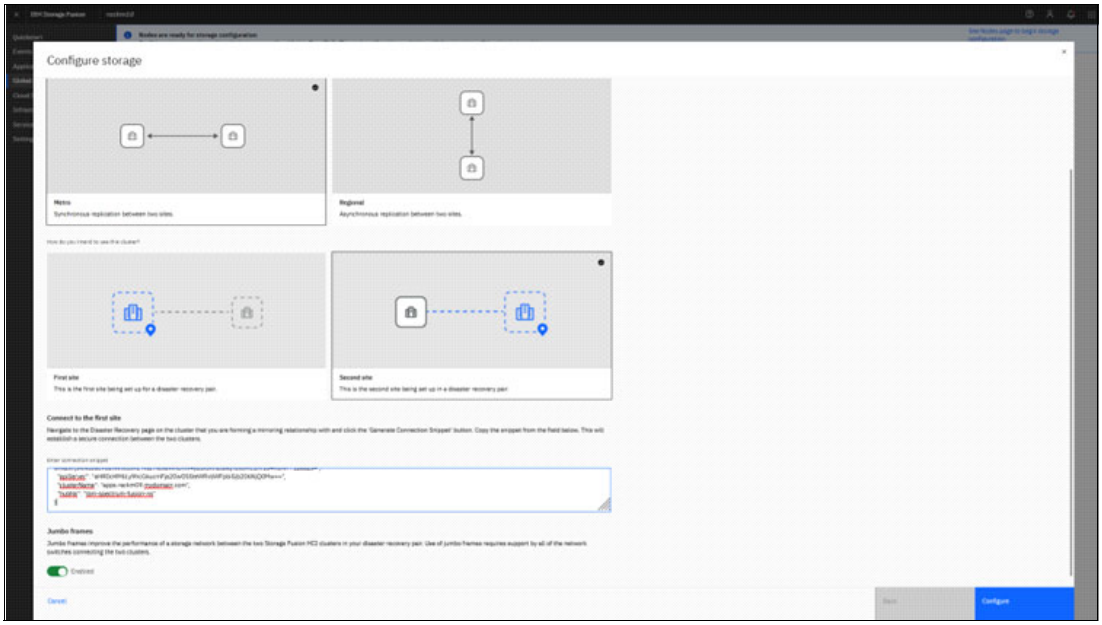

Figure 2-54 Site2 - Global Data Platform – Configure Storage in progress

Once the Global Data Platform configuration is completed successfully, the health of the Storage cluster shows **OK** and the Storage status of nodes shows **Normal** as shown in Figure 2-55 on page 34.

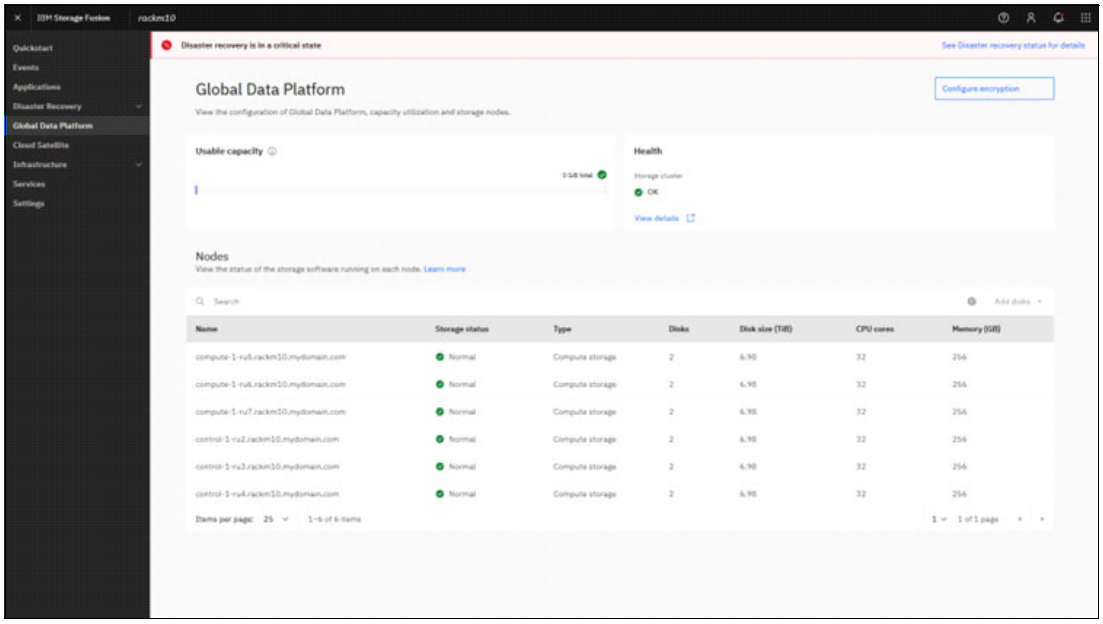

Figure 2-55 Site2 - Global Data Platform - Configure Storage Completed

Disaster recovery

On the IBM Storage Fusion user interface, in the left pane menu, click **Disaster recovery** → **Topology** as shown in Figure 2-56.

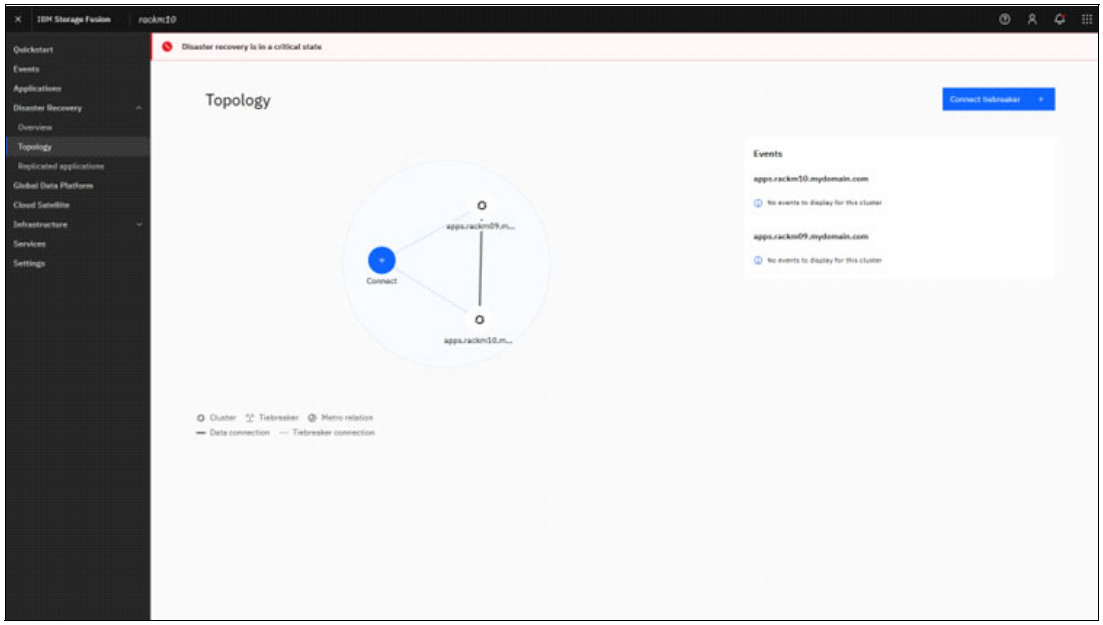

Figure 2-56 Disaster recovery

2.1.4 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 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:
  - Software requirements for 2.7.x.
  - Software requirements for 2.8.0.
  - Software requirements for 2.8.1.
- For tiebreaker allowed ports, see:
  - Securing the IBM Storage Scale system using firewall for 2.7.x.
  - Securing the IBM Storage Scale system using firewall for 2.8.0.
  - Securing the IBM Storage Scale system using firewall for 2.8.1.
- Download the following from the IBM Entitled System Support.
  - Storage_Scale_Data_Management-5.1.9.1-x86_64-Linux.tar.gz for 2.7.x
  - Storage_Scale_Data_Management-5.2.0.0-x86_64-Linux.tar.gz for 2.8.0
  - Storage_Scale_Data_Management-5.2.0.1-x86_64-Linux.tar.gz for 2.8.1

  The following steps show the download and installation:

1. Log in into the Entitled Systems Support (ESS) portal as shown in Figure 2-57.

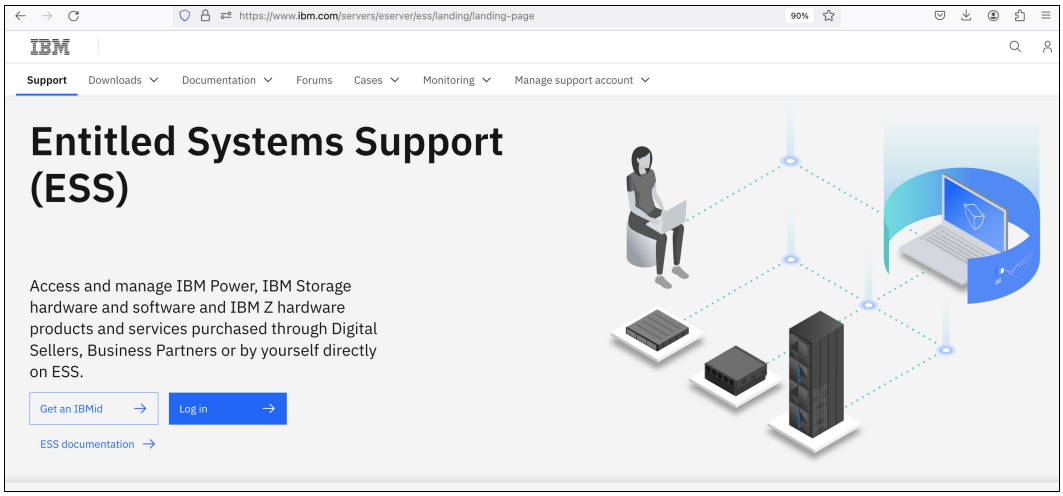

Figure 2-57 Entitled Systems Support (ESS) portal

2. Click on **My Entitled Software** as shown in Figure 2-58 on page 36.

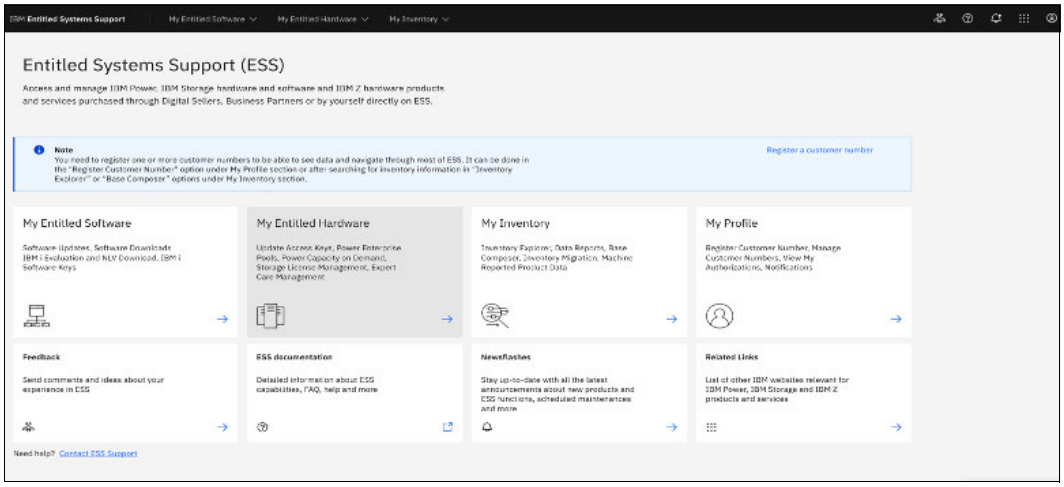

Figure 2-58 Entitled Systems Support website

3. Click on **Software Downloads** as shown in Figure 2-59.

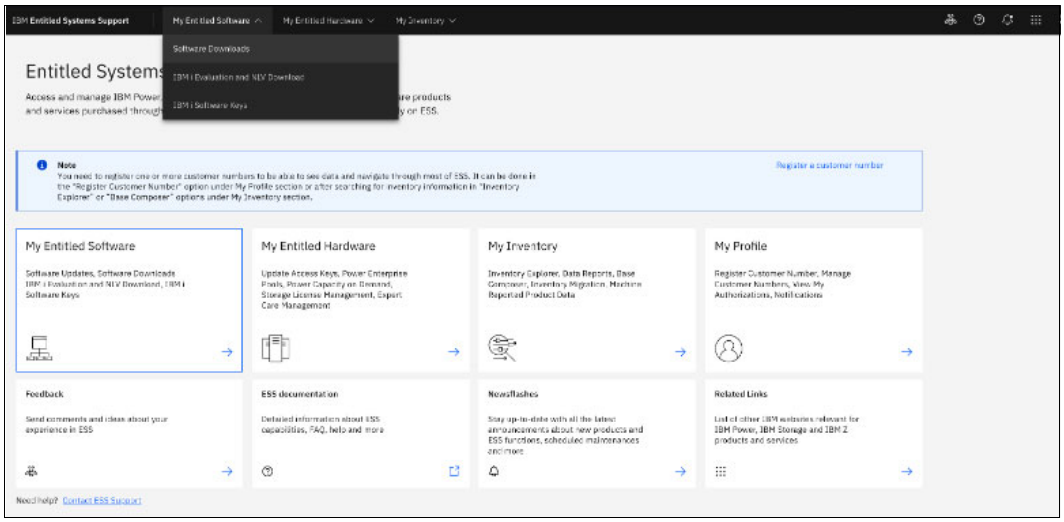

Figure 2-59 Software Downloads

Figure 2-60 provides background for ESS website for **My Entitled Software** downloads.

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 2-60 Software Downloads

4. Search for the product **5771** as shown in Figure 2-61.

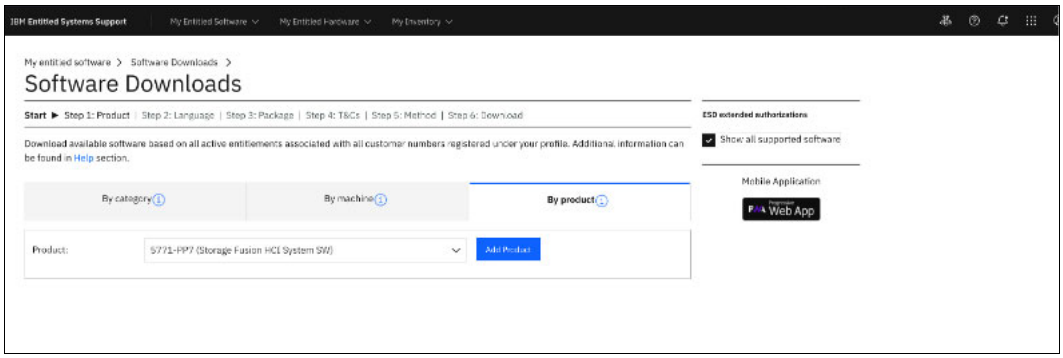

Figure 2-61 Search Product

5. Select the product **5771-PP7 (IBM Storage Fusion HCI)** as shown in Figure 2-62.

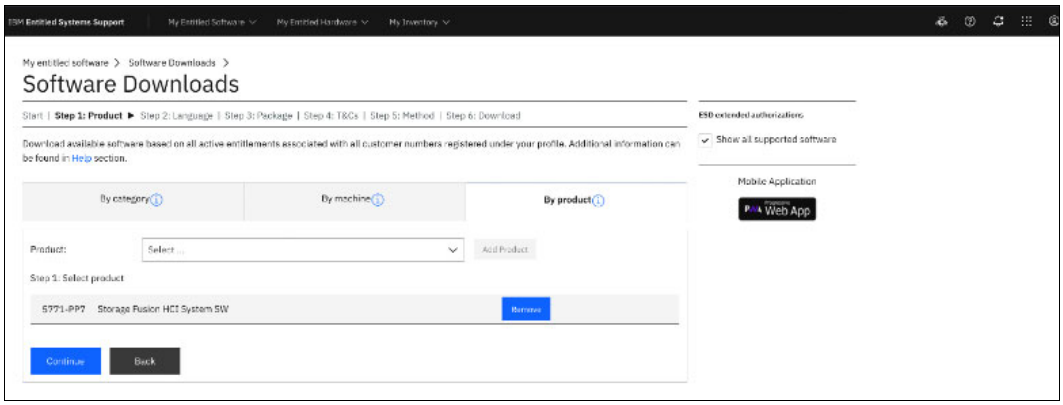

Figure 2-62 Select the product

6. The selected product is displayed as shown in Figure 2-63. Click **Continue**.

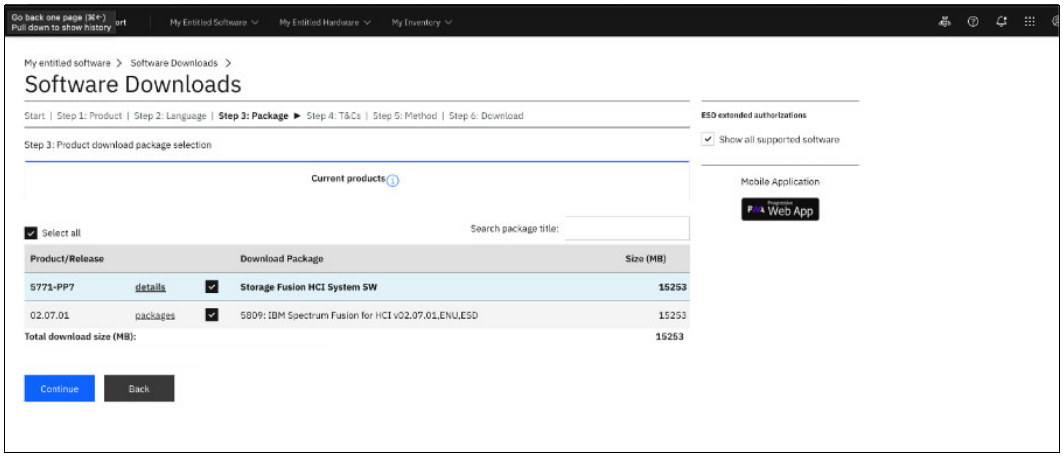

Figure 2-63 Selected product display

7. Click the checkbox as shown in Figure 2-64 on page 38. Click **Continue**.

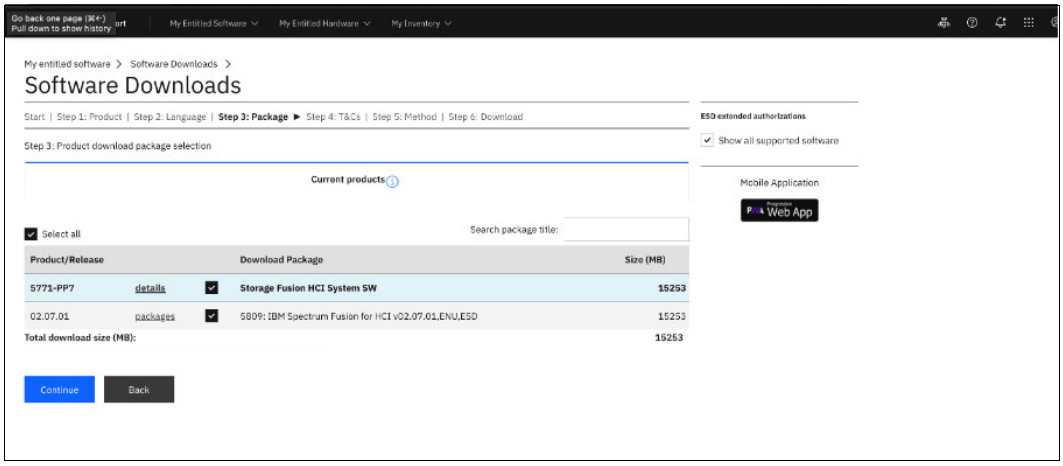

Figure 2-64 Confirm the selected products

8. Read the License Terms and click on **I agree** as shown in Figure 2-64.

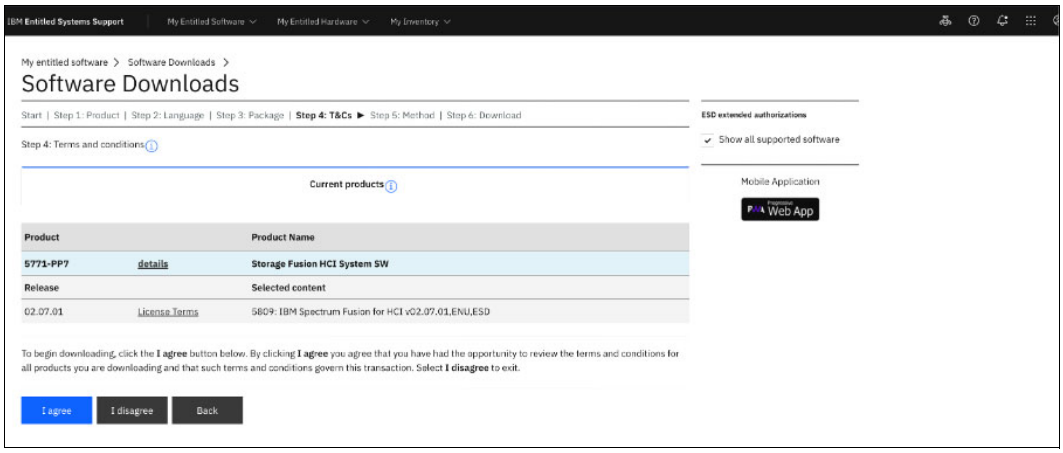

Figure 2-65 License Terms confirmation

9. Select the download method as shown in Figure 2-66 and click **Continue**.

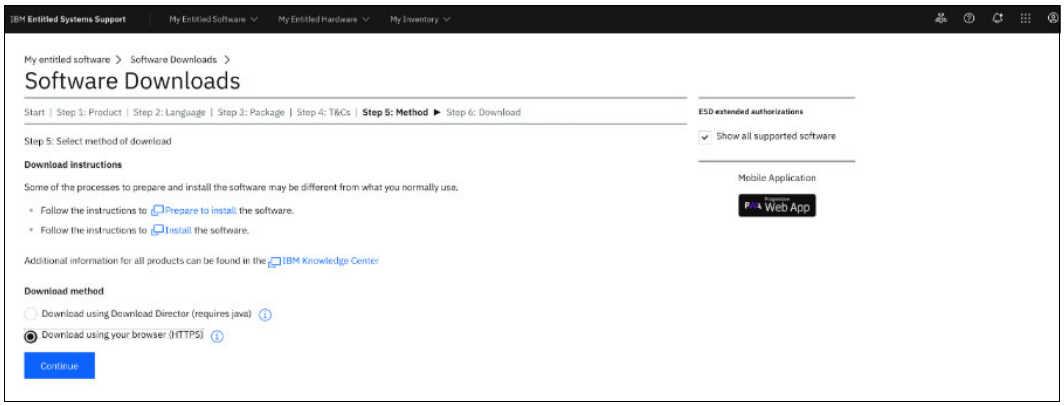

Figure 2-66 Download method

10. Review the download details as shown in Figure 2-67 and click the **Download now** button.

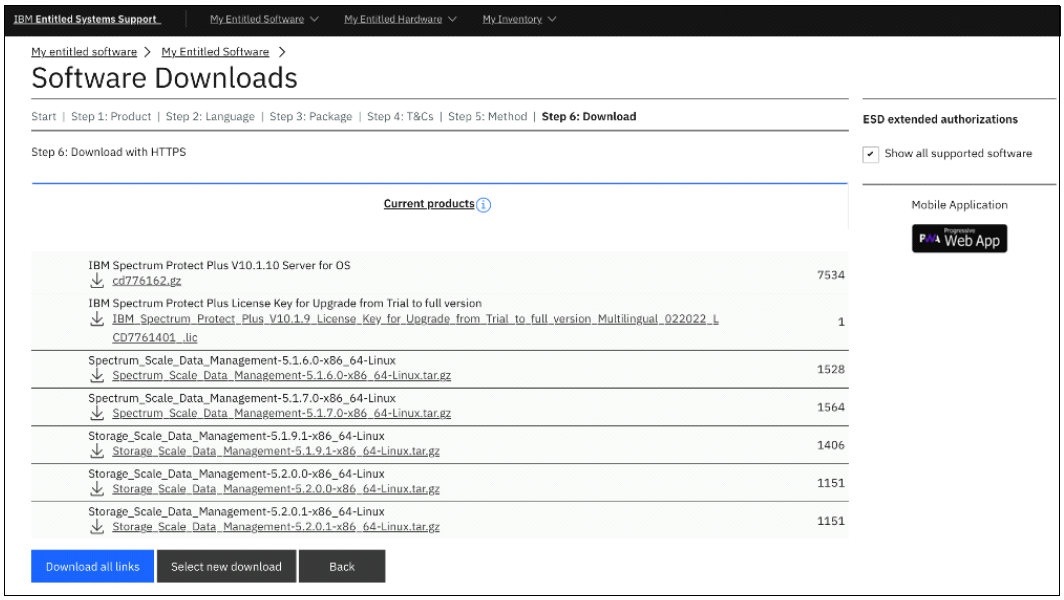

Figure 2-67 Start download

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@isftc11tbm9m10 ~]#
ls Storage_Scale_Data_Management-5.1.9.1-x86_64-Linux.tar.gz
```

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

Example 2-2 Permission change

```
[root@isftc11tbm9m10 ~]# $chmod +x
Storage_Scale_Data_Management-5.1.9.1-x86_64-Linux.tar.gz
$./Storage_Scale_Data_Management-5.1.9.1-x86_64-Linux.tar.gz
```

13. Start the installation as shown in Figure 2-68 on page 40 and Figure 2-69 on page 41. Input the details as requested.

```
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 Storage Scale Data Access Edition 5.1.9.1 (5737-I39)
IBM Storage Scale Data Access Edition 5.1.9.1 (5641-DA1)
IBM Storage Scale Data Access Edition 5.1.9.1 (5641-DA3)
IBM Storage Scale Data Access Edition 5.1.9.1 (5641-DA5)

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/5.1.9.1 ...
tail -n +678 ./cd776372.tar.gz | tar -C /usr/lpp/mmfs/5.1.9.1 --wildcards -xvz  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 smb_debs/ubuntu/ubuntu20 smb_debs/ubuntu/ubuntu22
zimon_debs/ubuntu/ubuntu20 zimon_debs/ubuntu/ubuntu22 ganesha_rpms/rhel7 ganesha_rpms/rhel8 ganesha_rpms/rhel9
ganesha_rpms/sles15 gpfs_rpms/rhel7 gpfs_rpms/rhel8 gpfs_rpms/rhel9 gpfs_rpms/sles15 smb_rpms/rhel7 smb_rpms/rhel8
smb_rpms/rhel9 smb_rpms/sles15 zimon_debs/ubuntu zimon_rpms/rhel7 zimon_rpms/rhel8 zimon_rpms/rhel9
zimon_rpms/sles15 cloudkit gpfs_debs gpfs_rpms scaleapi_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
  - smb_debs/ubuntu/ubuntu20
  - smb_debs/ubuntu/ubuntu22
    ...
  - zimon_rpms/rhel9
  - zimon_rpms/sles15
  - cloudkit
```

Figure 2-68 Installation

Figure 2-69 on page 41 is the continuation of the installation show in Figure 2-68.

```
  - gpfs_debs
  - gpfs_rpms
  - scaleapi_rpms
  - manifest

Removing License Acceptance Process Tool from /usr/lpp/mmfs/5.1.9.1 ...
rm -rf  /usr/lpp/mmfs/5.1.9.1/LAP_HOME /usr/lpp/mmfs/5.1.9.1/LA_HOME

Removing JRE from /usr/lpp/mmfs/5.1.9.1 ...
rm -rf /usr/lpp/mmfs/5.1.9.1/ibm-*gz

===============================================================
Product packages successfully extracted to /usr/lpp/mmfs/5.1.9.1

  Cluster installation and protocol deployment
    To install a cluster or deploy protocols with the IBM Storage Scale Installation Toolkit:
    /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale -h

      To install a cluster on public clouds (AWS, GCP) with the IBM Storage Scale Installation Toolkit:
      /usr/lpp/mmfs/5.1.9.1/cloudkit/cloudkit -h

    To install a cluster manually:  Use the GPFS packages located within /usr/lpp/mmfs/5.1.9.1/gpfs_<rpms/debs>

    To upgrade an existing cluster using the IBM Storage Scale Installation Toolkit:
    1) Review and update the config:  /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale config update
    2) Update the cluster configuration to reflect the current cluster config:
      /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale config populate -N <node>
    3) Use online or offline upgrade depending on your requirements:
      - Run the online rolling upgrade:  /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale upgrade -h
      - Run the offline upgrade: /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale upgrade config offline -N;
          /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale upgrade run
    You can also run the parallel offline upgrade to upgrade all nodes parallely after shutting down GPFS
    and stopping protocol services on all nodes.
    You can run the parallel offline upgrade on all nodes in the cluster, not on a subset of nodes.

    To add nodes to an existing cluster using the IBM Storage Scale Installation Toolkit:
    1) Add nodes to the cluster definition file:  /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale node add -h
    2) Install IBM Storage Scale on the new nodes:  /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale install -h
    3) Deploy protocols on the new nodes:  /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale deploy -h

    To add NSDs or file systems to an existing cluster using the IBM Storage Scale Installation Toolkit:
    1) Add NSDs or file systems to the cluster definition:  /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale nsd add -h
    2) Install the NSDs or file systems:  /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale install -h


    To update the cluster definition to reflect the current cluster config examples:
      /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/spectrumscale config populate -N <node>
    1) Manual updates outside of the installation toolkit
    2) Sync the current cluster state to the installation toolkit prior to upgrade
    3) Switching from a manually managed cluster to the installation toolkit

================================================================================
To get up and running quickly, consult the IBM Storage Scale Protocols Quick Overview:
https://www.ibm.com/docs/en/STXKQY_5.1.9/pdf/scale_povr.pdf
================================================================================
```

Figure 2-69 Install continuation

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

Example 2-3 Ansible toolkit directory

```
[root@isftc11tbm9m10 ~]# cd /usr/lpp/mmfs/5.1.9.1/ansible-toolkit
```

15. View the contents of the directory as shown in Example 2-4. Verify that you see a file named `Spectrumscale`.

Example 2-4 Directory contents of ansible-toolkit directory

```
[root@isftc11tbm9m10 ansible-toolkit]# ls
ansible  cli            documentation  externallibs       license  __pycache__  Spectrumscale
bin      configuration  espylib        installer.snap.py  logs     README       _version.py
```

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@isftc11tbm9m10 ansible-toolkit]# ./Spectrumscale -s 192.168.192.220
[ INFO  ] Installing prerequisites for install node
[ INFO  ] Found existing Ansible installation on system.
[ INFO  ] Install Toolkit setup type is set to IBM Storage 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
192.168.192.220 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
```

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

Example 2-6 Secrets of IBM Storage Scale

```
oc get secret ibm-Storage-scale-core-ssh-key-secret -n ibm-Storage-scale -ojsonpath="{.data.ssh-authorizedkeys}"
oc get secret ibm-Storage-scale-core-ssh-key-secret -n ibm-Storage-scale -ojsonpath="{.data.ssh-privatekey}"
oc get secret ibm-Storage-scale-core-ssh-key-secret -n ibm-Storage-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@isftc11tbm9m10 ansible-toolkit]# cd /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/ansible/vars/
Edit tiebreaker_nodedefinition.json file with secrets collected as in 2.6 and tiebreaker ip
[root@isftc11tbm9m10 vars]# cat tiebreaker_nodedefinition.json
{
"scale_cluster": {
 "scale_version": "5.1.9.1",
 "ssh_authorizedkeys": "c3NoLXJzYSBBQUFBQjNOemFDMXljMkVBQ.."
"ssh_privatekey": "LSOtLS1CRUdJTiBSUOEgUFJJVkFURSBLRVktLSOtLQp.."
"ssh_publickey": "c3NoLXJzYSBBQUFBQjNOemFDMXljMkVBQUFBREFR.."
},
"tiebreaker_node": [
 {
   "fqdn": "192.168.192.220"
 }
]
}
```

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

Example 2-8 Copy the yaml files

```
cp /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/ansible/sample/playbook_tiebreakernode_install.yml
/usr/lpp/mmfs/5.1.9.1/ansible-toolkit/ansible/sample/set_json_variables_tb.yml
/usr/lpp/mmfs/5.1.9.1/ansible-toolkit/ansible/

Run the following command to configure the SSH passwordless:
```

```
sed -i 's/scale_private_public_key_config: false/scale_private_public_key_config: true/g'
/usr/lpp/mmfs/5.1.9.1/ansible-toolkit/ansible/playbook_tiebreakernode_install.yml
```

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

Example 2-9 Run the ansible playbook

```
[root@isftc11tbm9m10 ~]# cd /usr/lpp/mmfs/5.1.9.1/ansible-toolkit/ansible/
[root@isftc11tbm9m10 ansible]# ansible-playbook
playbook_tiebreakernode_install.yml
```

21. Ensure the playbook runs successfully as shown in Example 2-10.

Example 2-10 Run playbook to install Tiebreaker

```
[root@isftc11tbm9m10 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 [Pass all inputs related to creating Spectrum Scale cluster to all nodes]
*********************************************************************************
changed: [localhost] => (item={'fqdn': '192.168.192.220'})

PLAY [tierbreaker_node]
*********************************************************************************
*********************************************

TASK [Gathering Facts]
*********************************************************************************
*********************************************
ok: [192.168.192.220]

TASK [tierbreaker node | install]
*********************************************************************************
**********************************
ok: [192.168.192.220] => {
    "msg": "Tierbreaker node installation is in progress. This may take a while. Please
be patient."
}
```

```
…

TASK [prepare | Authorize all SSH keys]
*********************************************************************************************
*****************************
ok: [192.168.192.220] => (item=192.168.192.220)

TASK [shell]
*********************************************************************************************
*********************************************************
changed: [192.168.192.220]

TASK [shell]
*********************************************************************************************
*********************************************************
changed: [192.168.192.220]

TASK [prepare | Change the Port 12345 in ssh_config]
*********************************************************************************************
*****************
ok: [192.168.192.220]

TASK [prepare | Change the Port 12345 in sshd_config]
*********************************************************************************************
***************
changed: [192.168.192.220]

TASK [prepare | Restart the sshd service]
*********************************************************************************************
****************************
changed: [192.168.192.220]

TASK [Create a file in tiebreaker to avoid mounting any filesystem]
*********************************************************************************************
**
changed: [192.168.192.220] => (item=192.168.192.220)
[WARNING]: Consider using the file module with state=touch rather than running 'touch'.
If you need to use command because file is insufficient you can add
'warn: false' to this command task or set 'command_warnings=False' in ansible.cfg to get
rid of this message.

PLAY RECAP
*********************************************************************************************
*********************************************************
192.168.192.220            : ok=67   changed=6    unreachable=0    failed=0
skipped=46   rescued=0    ignored=0
localhost                  : ok=4    changed=1    unreachable=0    failed=0    skipped=0
rescued=0    ignored=0

Run lsblk to get the device name of the raw disk attached to the tiebreaker VM. The
device name gets used later to add the tiebreaker into the IBM Storage Scale cluster.

[root@isftc11tbm9m10  ansible]# lsblk
NAME          MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda             8:0    0 447.1G  0 disk [root@isftc11tbm9m10 ansible]# ansible-playbook
playbook_tiebreakernode_install
|-sda1          8:1    0   600M  0 part /boot/efi
|-sda2          8:2    0     1G  0 part /boot
`-sda3          8:3    0 444.8G  0 part
```

```
  |-rhel-root 253:0    0   337G  0 lvm  /
  |-rhel-swap 253:1    0   7.8G  0 lvm  [SWAP]
  `-rhel-home 253:2    0   100G  0 lvm  /home
sdb             8:16   0 447.1G  0 disk

Run the following echo command to encode the disk name string:
[root@isftc11tbm9m10  ansible]# echo "/etc/sdb" | base64
```

22. Patch the secret on any of the sites of the IBM Storage Fusion metro-dr pair as shown in Example 2-11.

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

```
oc patch secret isf-metrodr-config-secret -n ibm-Storage-fusion-ns -p
'{"data":{"TieBreakerDevice":"Encoded disk name from previous command"}}'
```

Tiebreaker configuration from IBM Storage Fusion GUI

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

1. Log in to the IBM Storage Fusion GUI.
2. Go to **Disaster Recovery**.
   - Add Tiebreaker IPs and Credentials on the **Disaster recovery** page as shown in Figure 2-70.
   - After saving tiebreaker credentials, click on **Connect** as shown in Figure 2-71 on page 46.

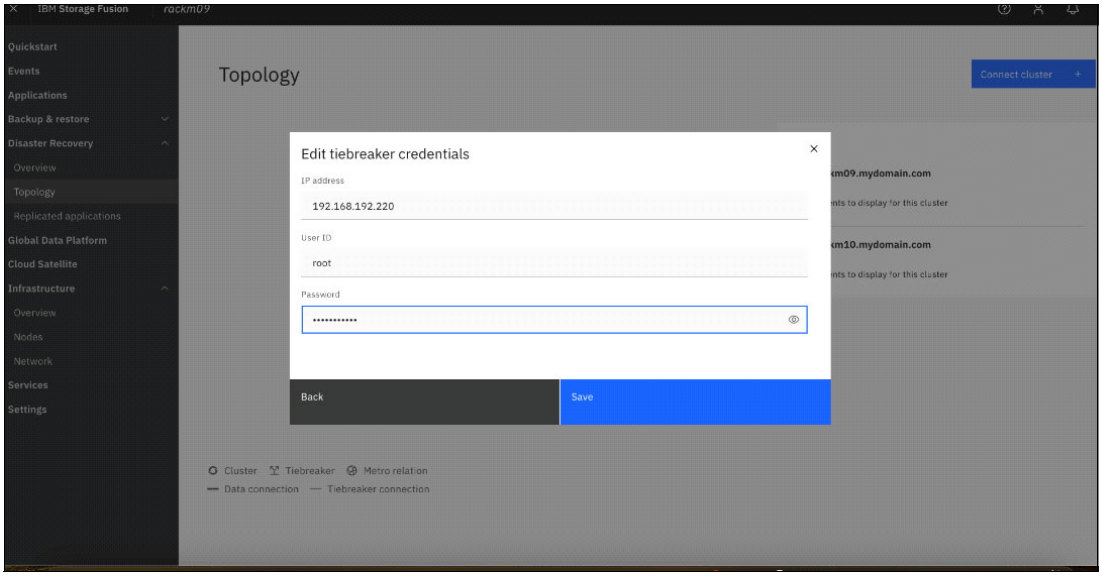

Figure 2-70 Add tiebreaker IPs and Credentials

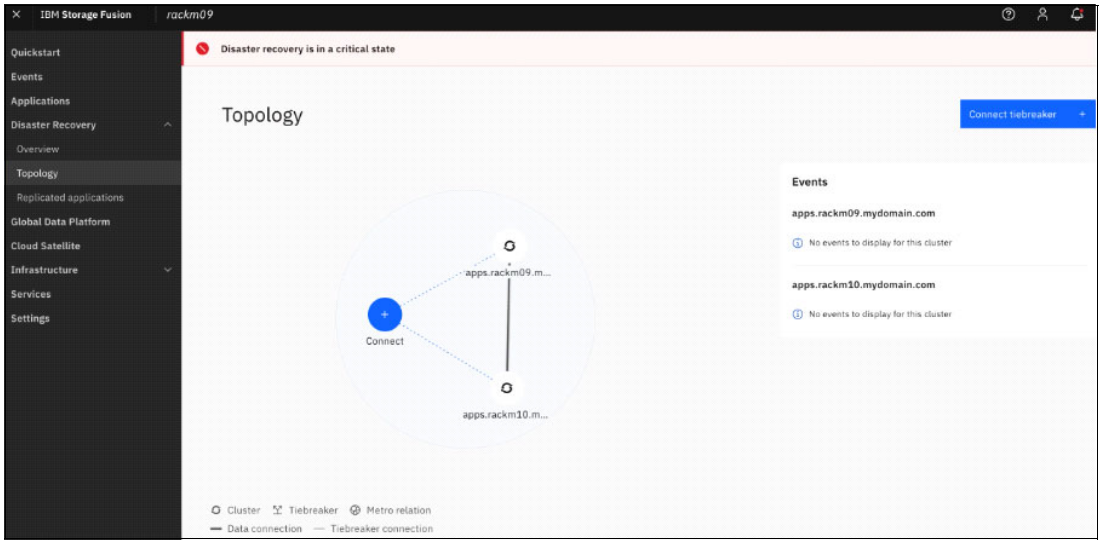

Figure 2-71 Connect Tiebreaker to Metro DR

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

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

```
sh-4.4# mmgetstate -a

 Node number  Node name                                        GPFS state
------------------------------------------------------------------------------
        1  control-1-ru2.daemon.ibm-Storage-scale.stg.rackm09 active
        2  control-1-ru3.daemon.ibm-Storage-scale.stg.rackm09 active
        3  control-1-ru4.daemon.ibm-Storage-scale.stg.rackm09 active
        4  compute-1-ru7.daemon.ibm-Storage-scale.stg.rackm09 active
        5  compute-1-ru5.daemon.ibm-Storage-scale.stg.rackm09 active
        6  compute-1-ru6.daemon.ibm-Storage-scale.stg.rackm09 active
        7  compute-1-ru5.daemon.ibm-Storage-scale.stg.rackm10 active
        8  compute-1-ru6.daemon.ibm-Storage-scale.stg.rackm10 active
        9  compute-1-ru7.daemon.ibm-Storage-scale.stg.rackm10 active
       10 control-1-ru2.daemon.ibm-Storage-scale.stg.rackm10 active
       11 control-1-ru3.daemon.ibm-Storage-scale.stg.rackm10 active
       12 control-1-ru4.daemon.ibm-Storage-scale.stg.rackm10 active
       13 gpfs-tiebreaker                                    active
```

Edit an existing Tiebreaker configuration from IBM Storage Fusion GUI

To Change/Modify (Tiebreaker IP Address and/or user ID and/or password) an existing Metro DR configuration from the IBM Storage Fusion GUI. Follow these steps:

1. Log in to the IBM Storage Fusion GUI of the local or remote site
2. Go to the **Topology** page under the **Disaster recovery** side menu.
   - Right-click on **Tiebreaker** and click **Edit**. Update Tiebreaker IPs and Credentials to the **Topology** page by right-clicking the tiebreaker icon (Figure 2-72 on page 47) and then completing the Edit tiebreaker credentials (Figure 2-73 on page 47).
   - Click on **Save**.

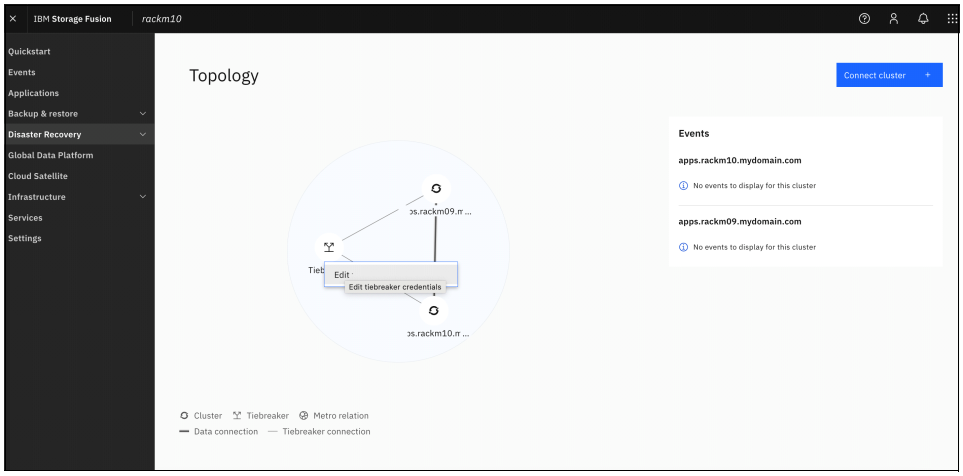

Figure 2-72 Select tiebreaker icon to update credential

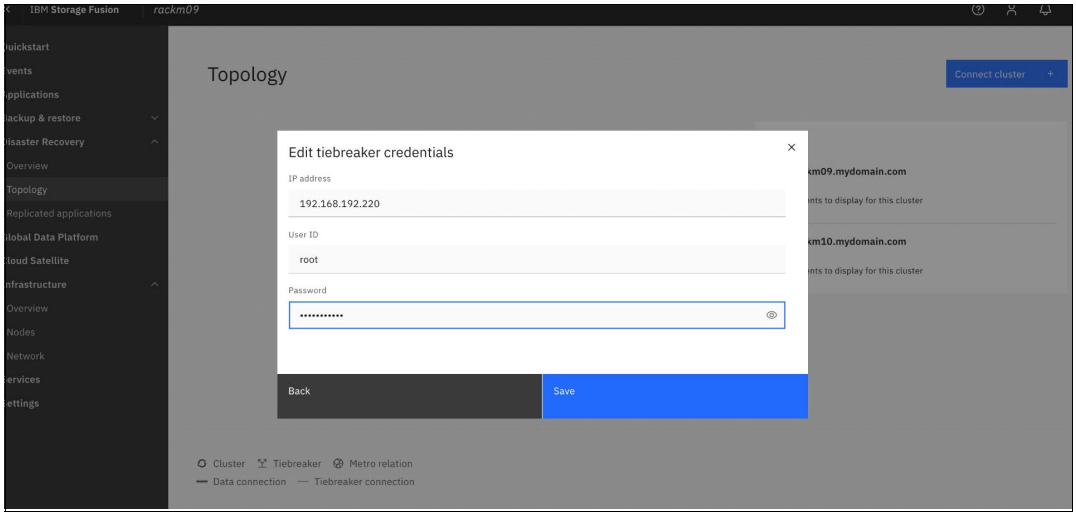

Figure 2-73 Edit tiebreaker IP and credentials and click Save

Verify IBM Storage Scale status with **`mmgetstate -a`** from OCP as shown in Example 2-12 on page 46.

Example 2-13 Verify IBM Storage Scale with mmgetstate -a command

```
sh-5.1# mmgetstate -a

 Node number  Node name                                        GPFS state
------------------------------------------------------------------------------
        14 control-1-ru2.daemon.ibm-Storage-scale.stg.rackm09 active
        15 control-1-ru3.daemon.ibm-Storage-scale.stg.rackm09 active
        16 control-1-ru4.daemon.ibm-Storage-scale.stg.rackm09 active
        17 compute-1-ru7.daemon.ibm-Storage-scale.stg.rackm09 active
        18 compute-1-ru5.daemon.ibm-Storage-scale.stg.rackm09 active
        19 compute-1-ru6.daemon.ibm-Storage-scale.stg.rackm09 active
        20 compute-1-ru5.daemon.ibm-Storage-scale.stg.rackm10 active
        21 compute-1-ru6.daemon.ibm-Storage-scale.stg.rackm10 active
        22 compute-1-ru7.daemon.ibm-Storage-scale.stg.rackm10 active
```

```
23 control-1-ru2.daemon.ibm-Storage-scale.stg.rackm10 active
24 control-1-ru3.daemon.ibm-Storage-scale.stg.rackm10 active
25 control-1-ru4.daemon.ibm-Storage-scale.stg.rackm10 active
26 gpfs-tiebreaker                                     active
```

MetroDR Setup Validations

In the OCP UI, go to **Administration** → **CustomResourceDefinitions** → Search for "MetroDR" → Click on **MetroDR** with group name "metrodr.isf.ibm.com".

Click on **Instances** tab → **Select metrodrsite** → Click on **YAML** tab.

Once the installation is successful and MetroDR clusters are ready, the following values should be seen in the yaml display:

```
overallDRClusterStatus: Healthy
clusterStatusSite1: Healthy
clusterStatusSite2: Healthy
clusterStatusTieBreaker: Healthy
networkStatusSite1: Healthy
networkStatusSite2: Healthy
```

Refer to Figure 2-74, Figure 2-75 on page 49, and Figure 2-76 on page 49.

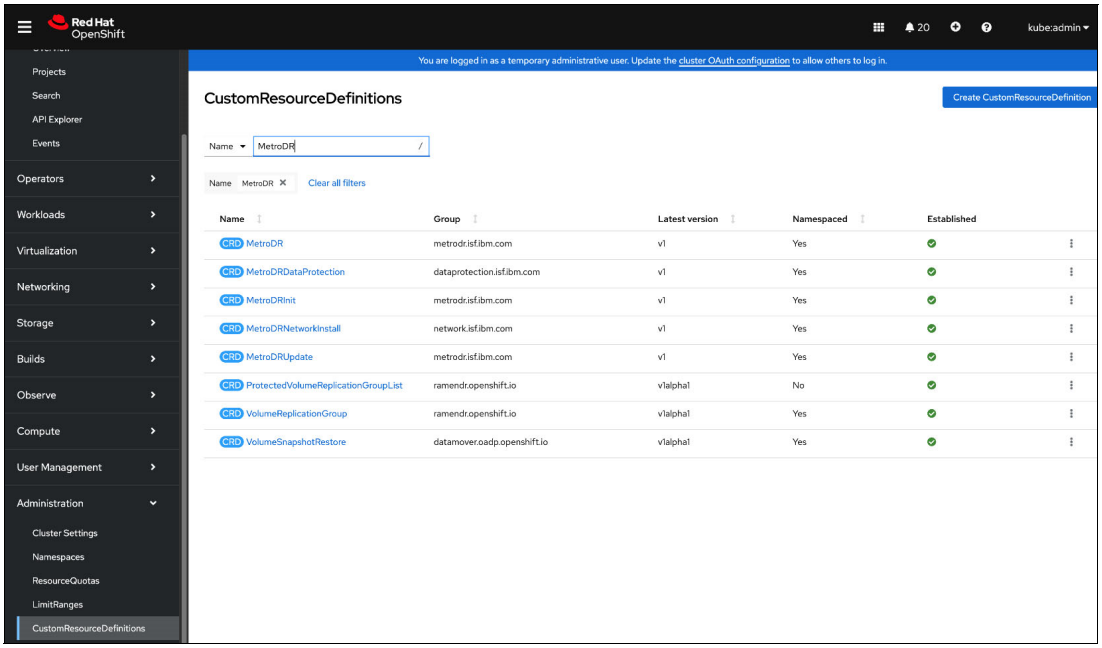

Figure 2-74 CustomResourceDefinitions and search for "MetroDR"

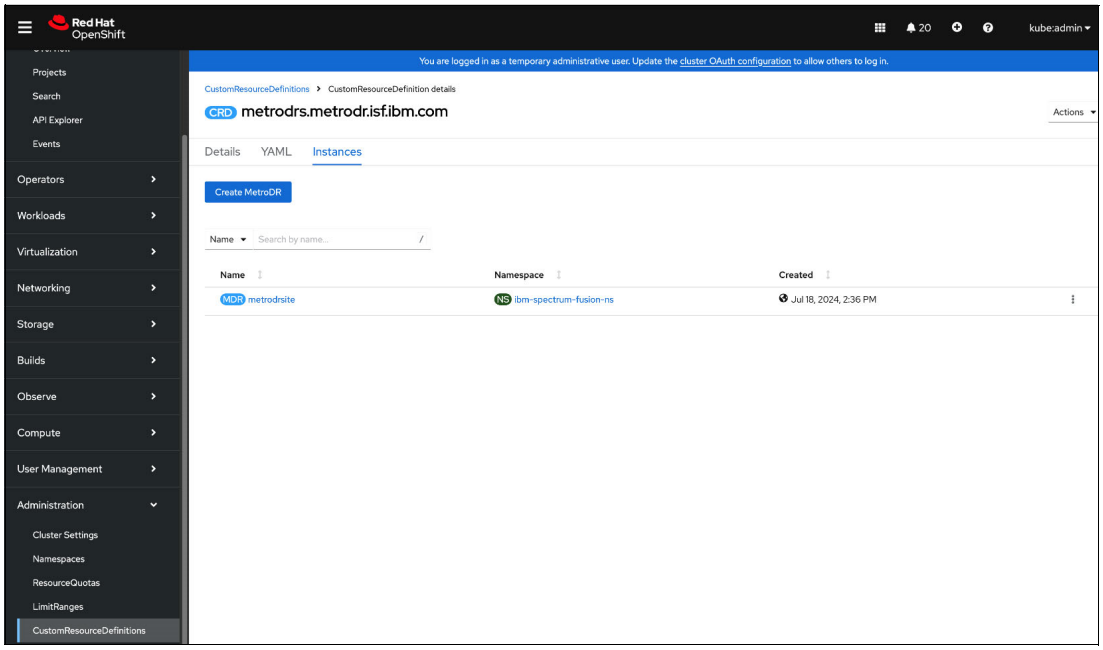

Figure 2-75 Click on MetroDR with group name “metrodr.isf.ibm.com

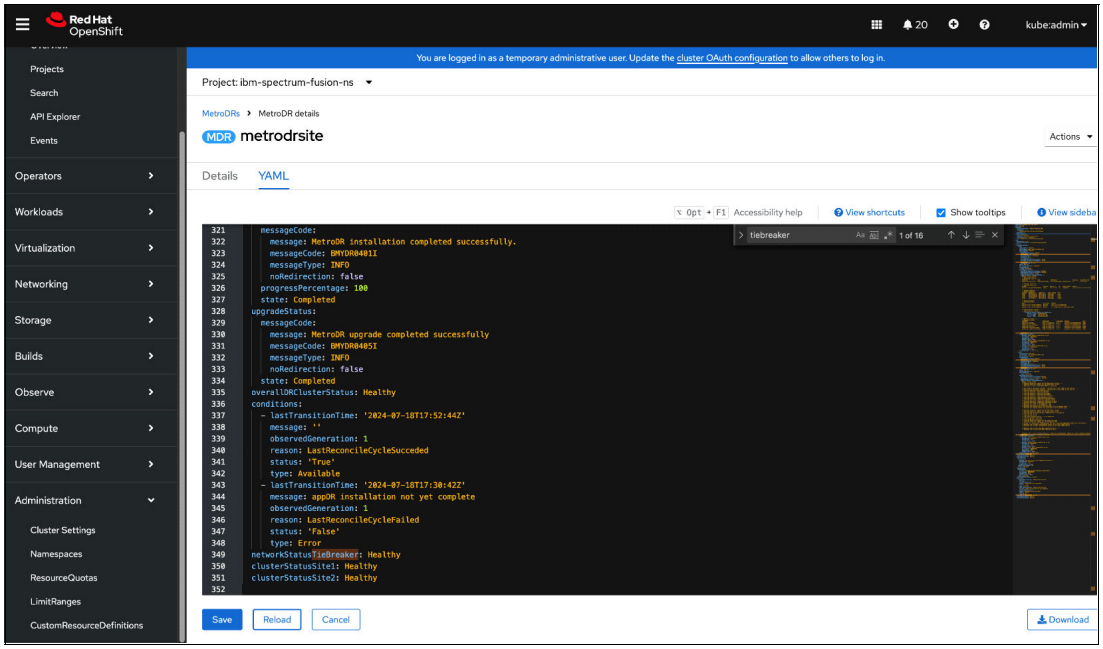

Figure 2-76 Verify metrodrsite component status

3

Setting up application failover and failback

This chapter describes the steps to set up the failover and failback of the file browser application between a local site and remote site with IBM 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 Storage Fusion provides a simple method for setting up the application(s) for disaster recovery (DR).

3.1.1 Prerequisites for application failover/failback

Before the application(s) is setup for DR, complete the following steps.

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

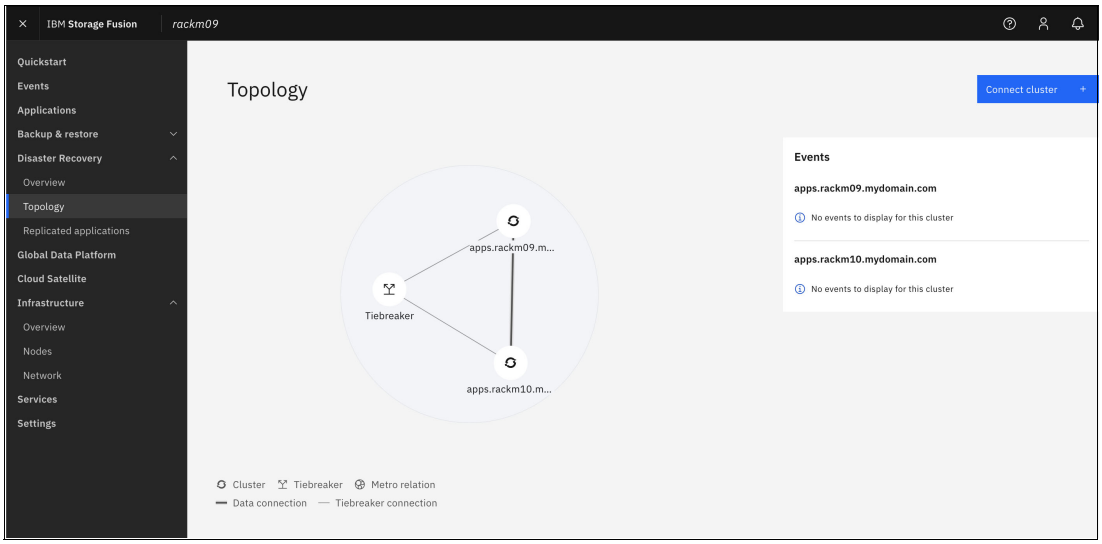

Figure 3-1 Disaster Recovery user interface

2. The application is deployed and displays on the **Applications** page as shown in Figure 3-2 on page 53.

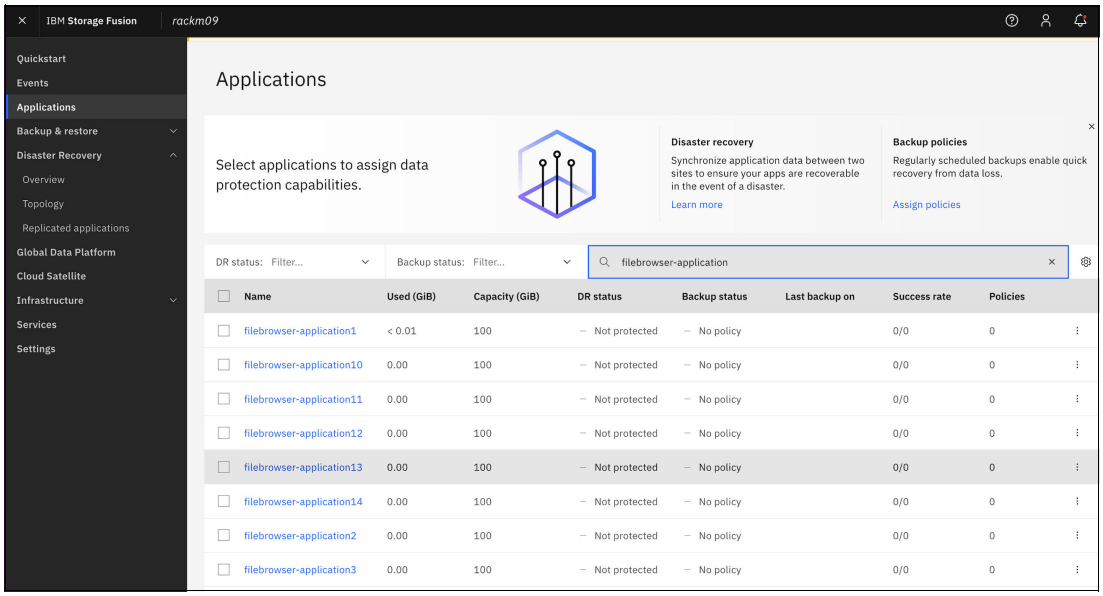

Figure 3-2 User-deployed applications showing up on the Applications page

3.1.2 Setting up applications for DR

You can use the IBM Storage 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 for setting up application(s) for DR.

Setting up a single application for DR

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

We can use any of the following methods or options to enable a single application for DR:

- Applications page
- Individual Application details page
- Disaster Recovery Section of Application Details page
- Replicated applications page

Method 1: Applications user interface.

Here are the steps for setting up the DR using the **Applications** page:

1. Go to the **Applications** page shown in Figure 3-3 on page 54.
2. 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**.
5. **Manage disaster recovery** pops up.
6. Select **Metro** section on **Manage disaster recovery**.
7. Click on **Save** to enable DR.
8. The enrollment process starts and **Adding application(s) to disaster recovery** is displayed as shown in Figure 3-4 on page 54.
9. Applications are showing up in DR status as “Synchronized” on the **applications** page as shown in Figure 3-5 on page 54.

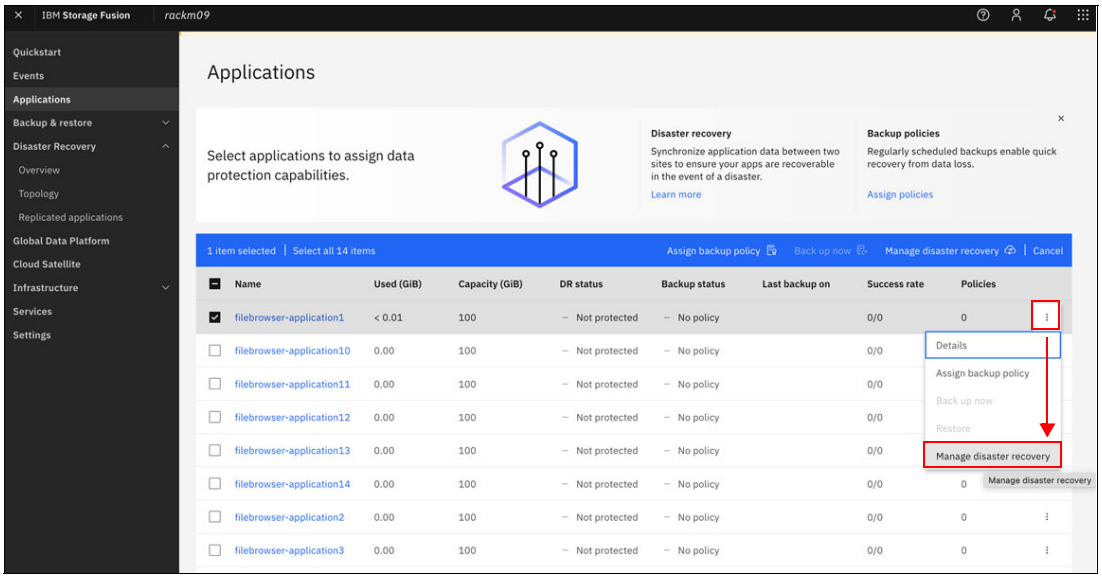

Figure 3-3 Enroll an application for DR from Applications page

The message **Adding disaster recovery** is displayed after enabling disaster recovery as shown in Figure 3-4.

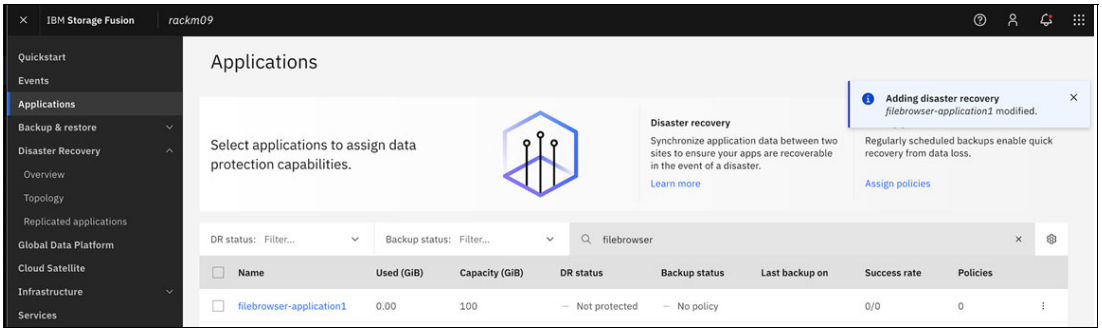

Figure 3-4 Adding application to disaster recovery

Figure 3-5 shows the application status (DR status) changed to “Synchronized” for the disaster recovery-enabled application.

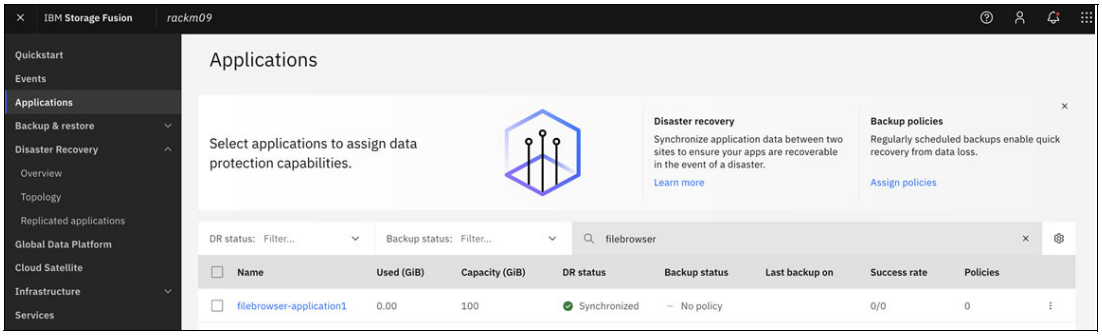

Figure 3-5 Application changed its DR status as “Synchronized”

Method 2: Applications details page

Here are the steps for setting up the DR Method 2 using the **Application details** page.

1. Click on the application from the **Applications** page.
2. The **Application details** page opens up.
3. Click on **Actions** to open the menu.
4. Click on **Manage disaster recovery** from the drop-down menu.
5. **Manage disaster recovery** pops up.
6. Select **Metro** from the popup window.
7. Click on **Save** to enable disaster recovery as shown in Figure 3-6.

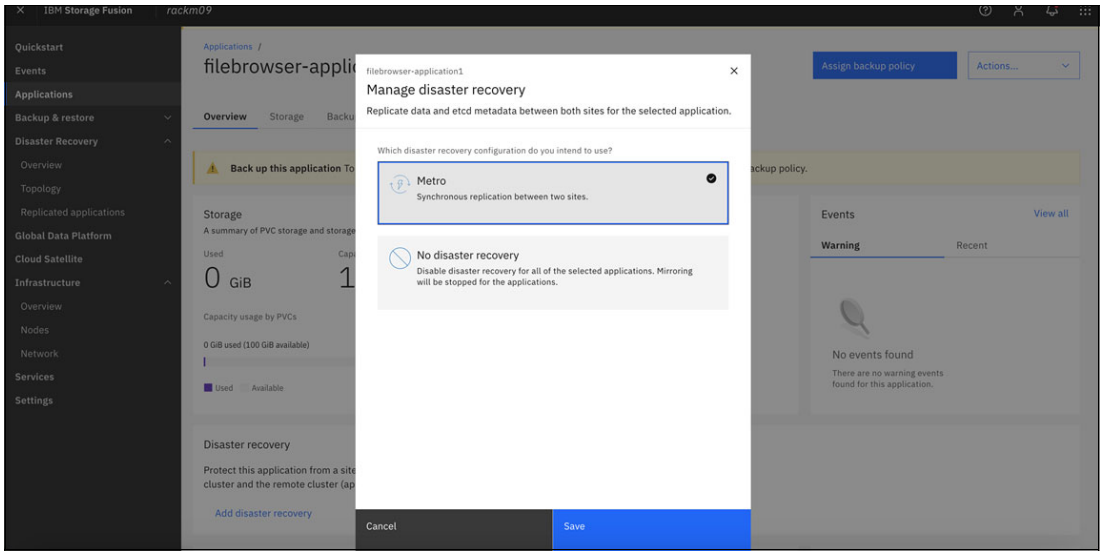

*Figure 3-6 Select **Metro** on Manage disaster recovery popup window and click **Save***

Method 3: Disaster Recovery Section on Application Details page

You can also setup an application for disaster recovery on the **Application details** page section **Disaster recovery**.

Here are the steps for setting up the DR Method 3 using the **Disaster recovery** section on the **application detailed** page.

1. Open the **Application detail** page.
2. Click on **Add disaster recovery** button as shown in Figure 3-7 on page 56.
3. Add the application to the DR displayed as shown in Figure 3-8 on page 56.
4. **Metro status** on the **disaster recovery** section changed to “Synchronized” as shown in Figure 3-9 on page 57

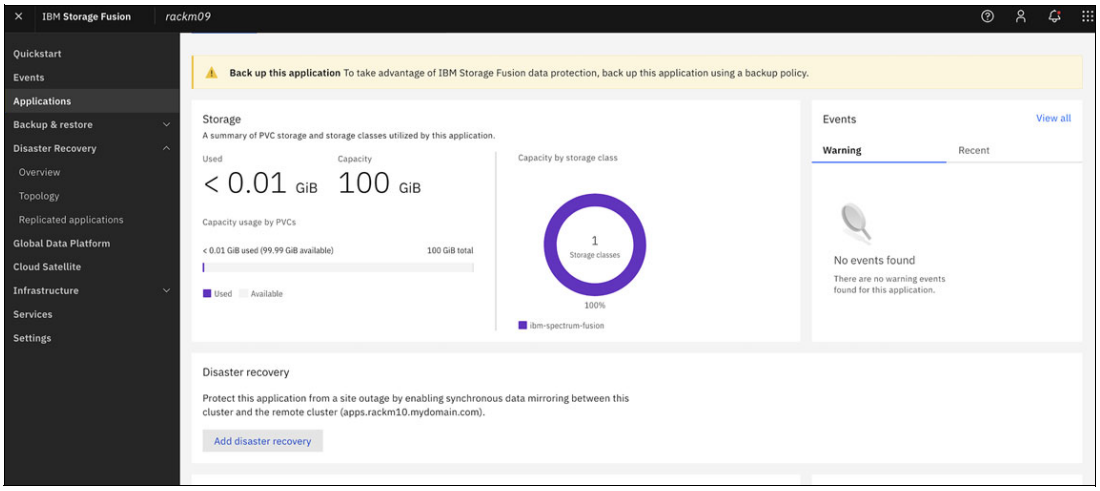

*Figure 3-7 Click on **Add disaster recovery** button on the **Disaster recovery** section*

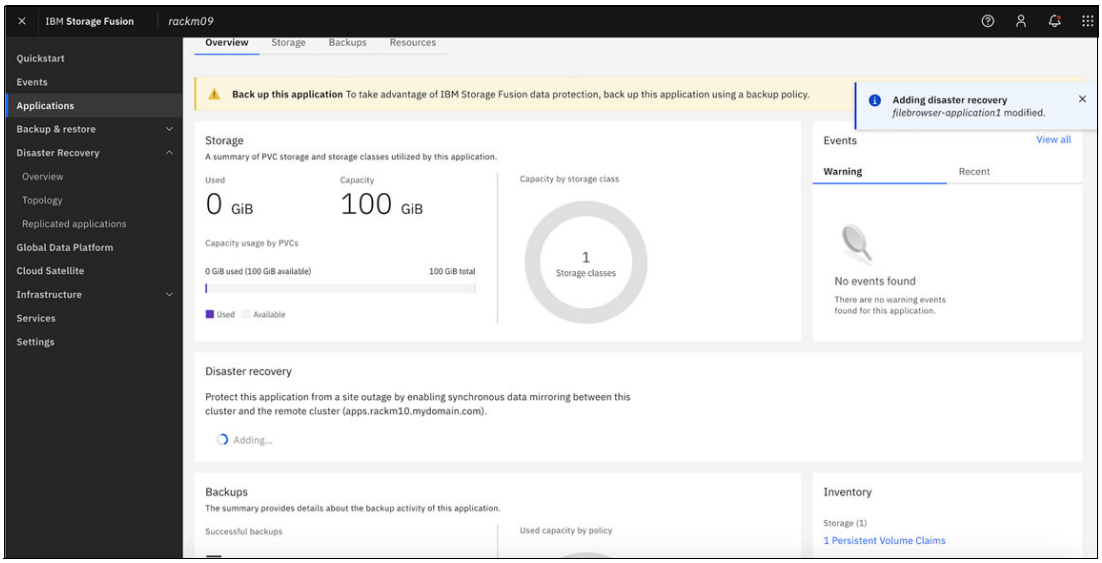

Figure 3-8 Adding disaster recovery is in progress

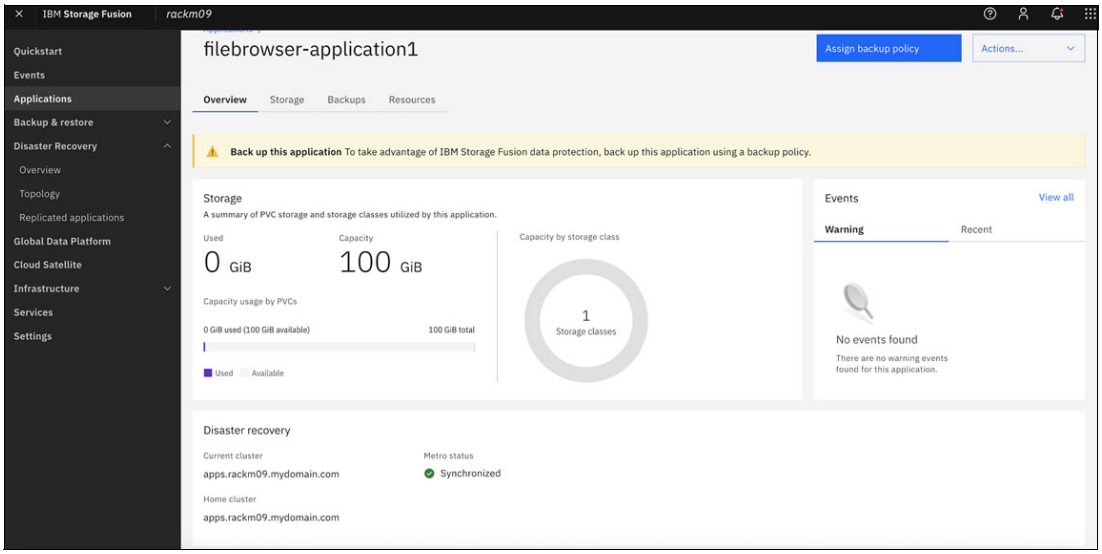

Figure 3-9 Metro status changed to “Synchronized”

Method 4: Replicated applications page

You can also set up an application for DR from the **Replicated applications** page.

1. Click on the disaster Recovery.
2. Open the **Replicated applications** page.
3. Click on the **Enroll applications +** button.
4. Select an application on **Enroll applications in disaster recovery popup window** as shown in Figure 3-10.
5. Click on **check boxes** for the selected applications.
6. Click on **Enroll.**

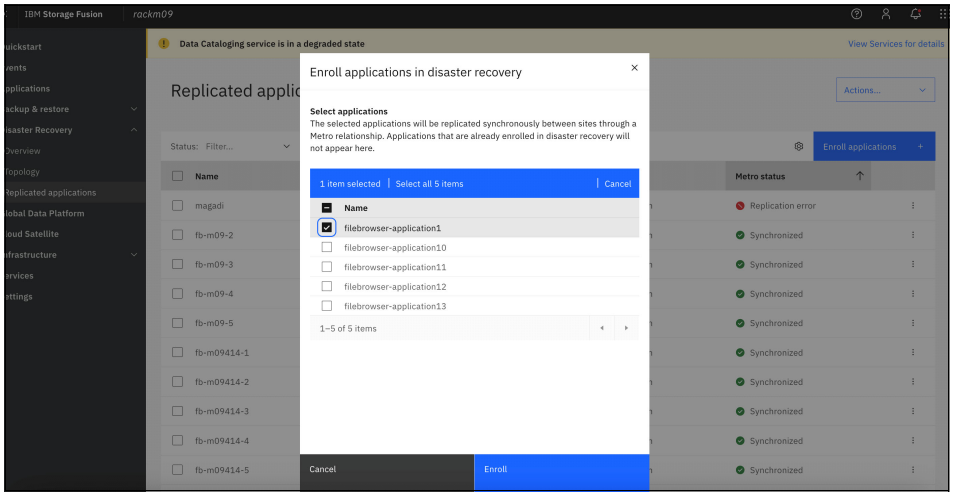

Figure 3-10 Enable disaster recovery from Replicated applications page

Setting up multiple applications for DR

We can use either of the following methods to enable a single application for DR:

- Applications page
- Replicated applications page

Method 1: Applications page

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

1. Open the **Applications** page.
2. Click the **checkbox** next to each Application name as shown in Figure 3-11.
3. Click on the **Manage disaster recovery** button to enroll the selected applications for DR.
4. **Manage disaster recovery** window pops up.
5. Select **Metro** and click on **Save** to enable disaster recovery as shown in Figure 3-12.
6. **Adding disaster recovery** message displayed as shown in Figure 3-13 on page 59.
7. **Disaster recovery enabled** message displayed and applications are showing up as “synchronized”. See Figure 3-14 on page 59.

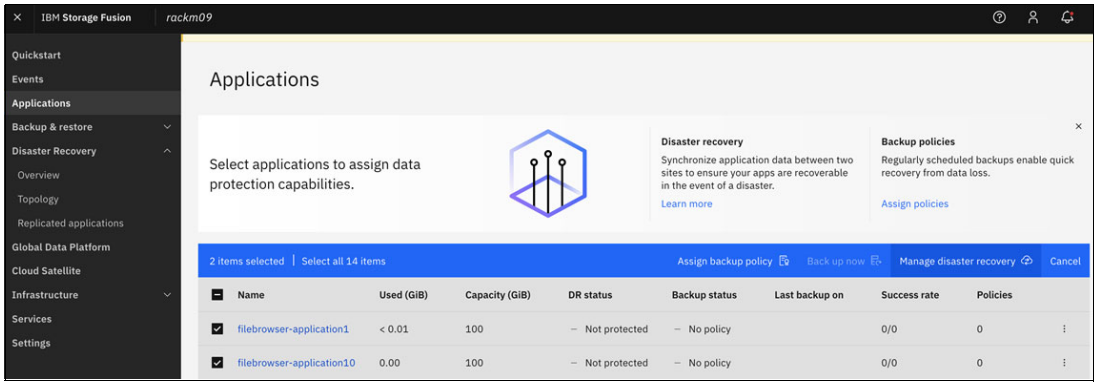

Figure 3-11 Enroll multiple applications for DR

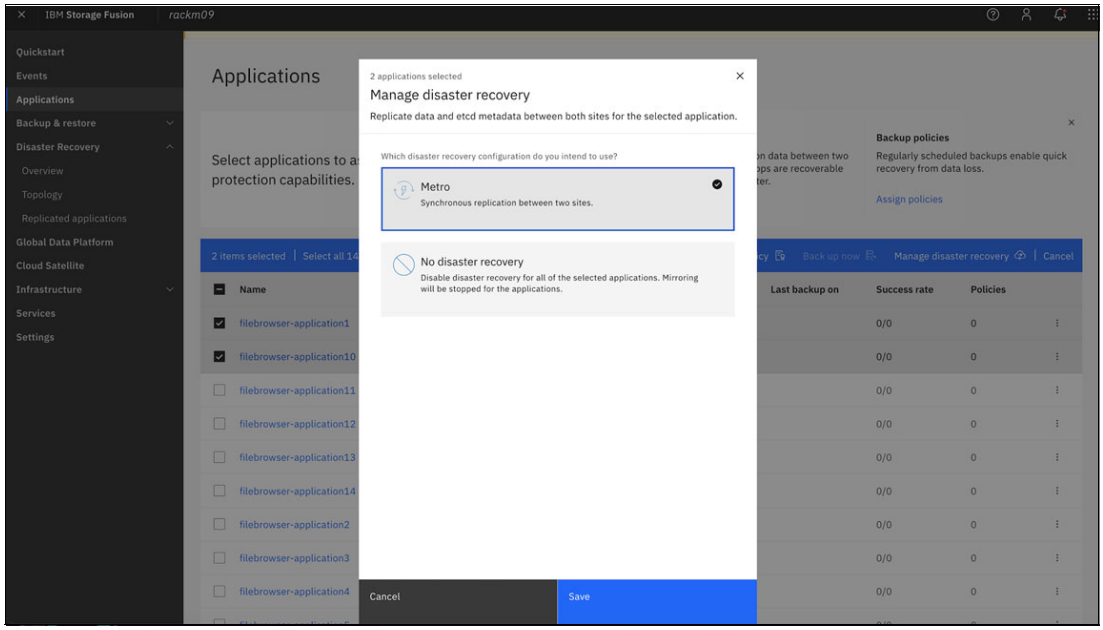

*Figure 3-12 Select **Metro** and click on **Save** to enable DR*

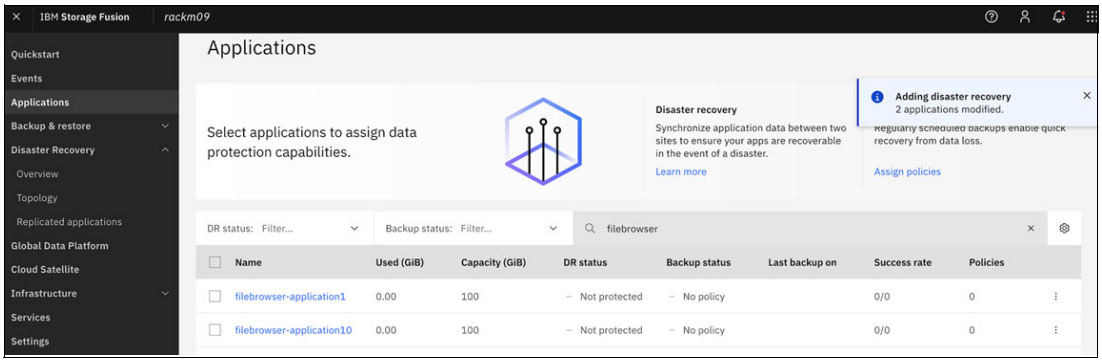

Figure 3-13 Adding DR is in progress

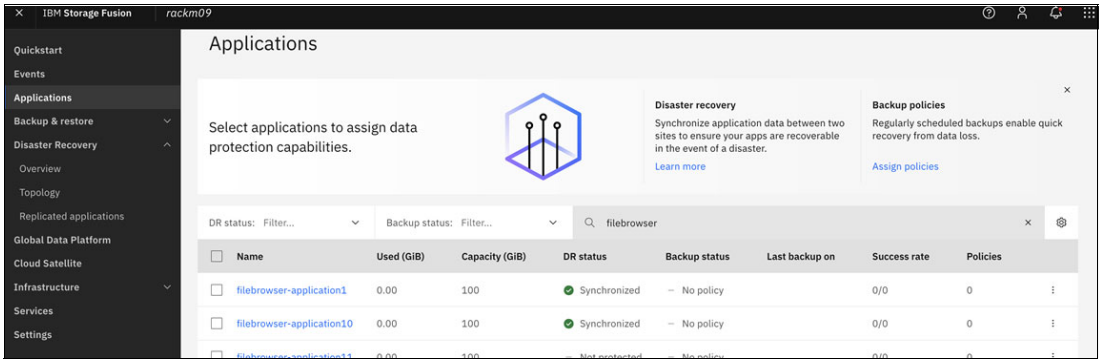

Figure 3-14 DR status changed to “Synchronized”

Method 2: Enroll applications for DR from Replicated Applications page

You can also set up an application for disaster recovery from the **Replicated applications** page.

1. Click on the **Disaster Recovery.**
2. Open the **Replicated Applications** page.
3. Click on the **Enroll applications +** button.
4. Select more than one application on **Enroll applications in disaster recovery** popup window as shown in Figure 3-15 on page 60.
5. Click on **check boxes** for the selected applications.
6. Click on **Enroll.**

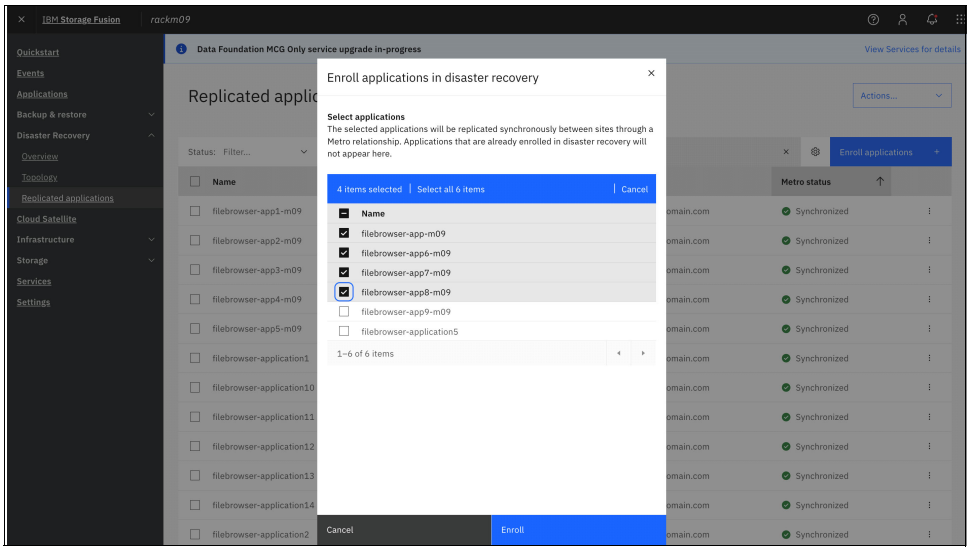

Figure 3-15 Enable disaster recovery from Replicated applications page

3.1.3 Viewing DR-enabled or synchronized applications for Site1 and Site2

View Disaster recovery enabled or synchronized applications on **Replicated applications** page on **Disaster recovery** main pages of Site1 and Site2. Figure 3-16 shows Synchronized applications on Site1 and Figure 3-17 on page 61 shows Synchronized applications on Site2.

Applications Metro status set to “Synchronized” on the **Replicated applications** page on Site1.

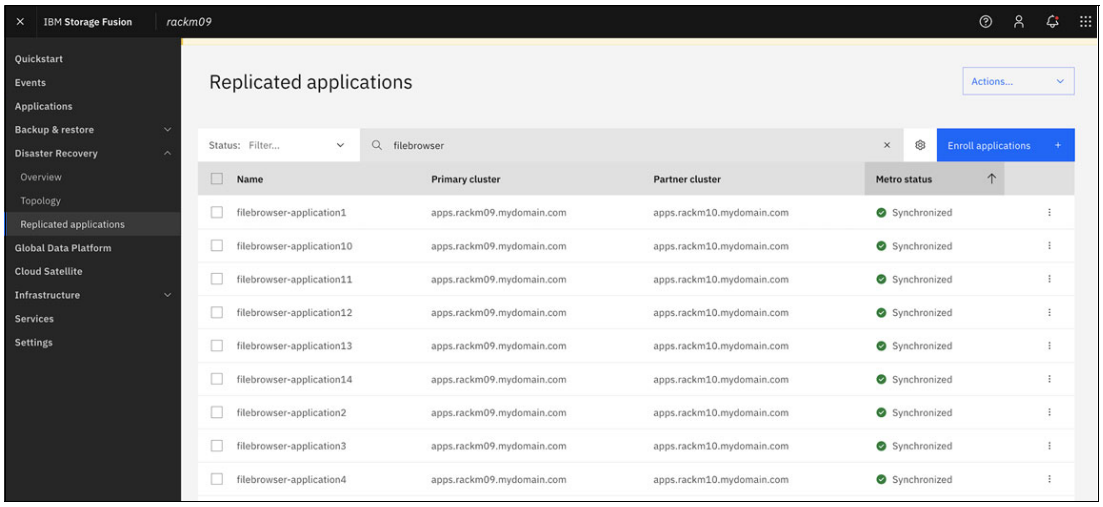

Figure 3-16 Replicated applications page showing “Synchronized” applications on Site1

Applications Metro status set to “Synchronized” on the **Replicated applications** page on Site2.

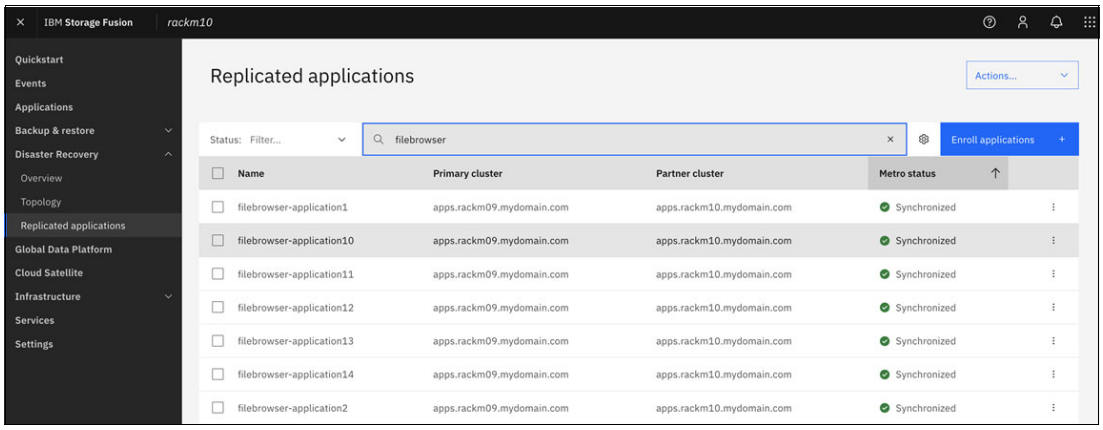

Figure 3-17 Replicated applications page showing "Synchronized" applications on Site2

3.1.4 Removing application(s) from DR

If you do not want to use any application for disaster recovery, you can disable the application(s) from disaster recovery.

Disabling DR for a Single application

We can use any of the following methods or options to disable a single application for disaster recovery.

- Applications page
- Individual Applications details page
- Replicated applications page

Method 1: Disable DR from the Applications page

You can disable an application from disaster recover from the **Applications** page with the following steps:

1. Click on the **Applications** page.
2. Select an application and click on 3 dots and click on **Manage disaster recovery** as shown in Figure 3-18 on page 62.
3. Select **Metro** and click on **No disaster recovery** and save it as shown in Figure 3-19 on page 62.
4. Confirm by clicking **Remove** on the **Remove disaster Recovery** popup window shown in Figure 3-20 on page 63.
5. Figure 3-21 on page 63 shows the **Removing disaster recovery** message.
6. Go back to the **Applications** page, Figure 3-22 on page 63, and see the DR status changed to **Not protected.**

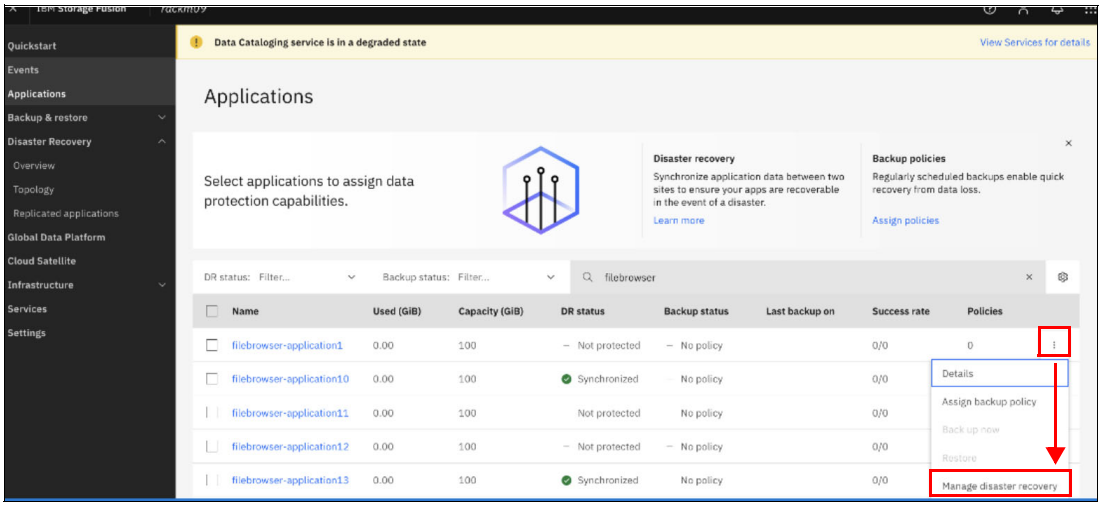

*Figure 3-18 Select an application and click on 3 dots and click on **manage disaster recovery***

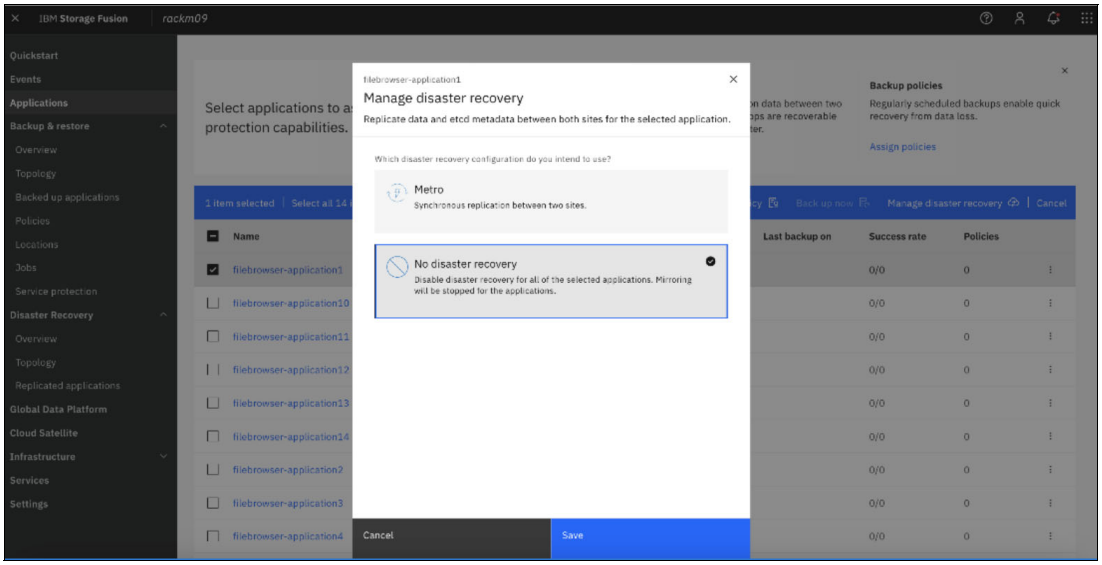

*Figure 3-19 Select **Metro** and click on **No disaster recovery** and save it*

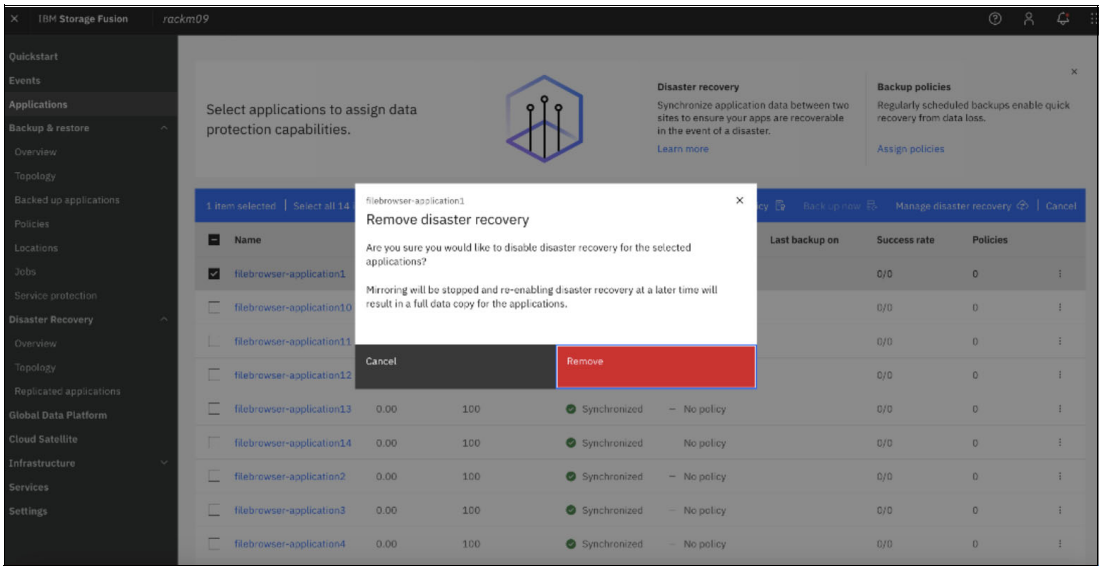

*Figure 3-20 Click on **Remove** to confirm*

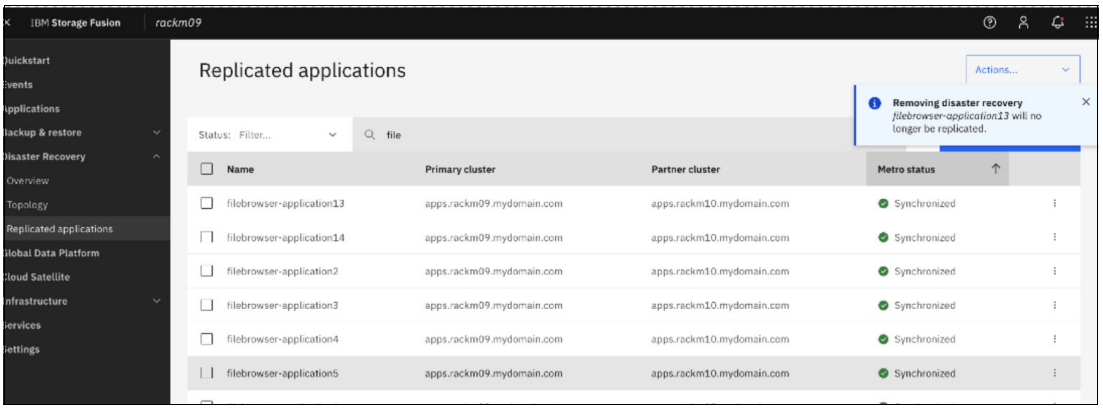

*Figure 3-21 **Removing disaster recovery** message is displayed*

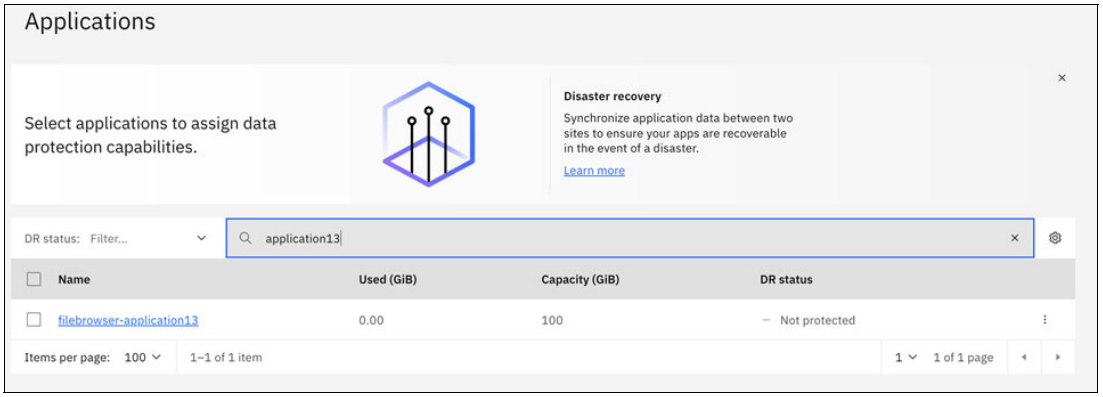

*Figure 3-22 DR status changed to **Not protected***

Method 2: Disable disaster recovery from the Application detail page

You can disable an application from disaster recovery from the **Applications detail** page with the following steps.

1. Click on the **Application** and open detail page.
2. Click on **Manage disaster recovery** in Figure 3-23.
3. Confirm by clicking **Remove** on the **Remove disaster Recovery** popup window shown in Figure 3-24.
4. **Removing disaster recovery** message displayed in Figure 3-25 on page 65.
5. Metro status “Synchronization” removed on the **Disaster recovery** section. See Figure 3-26 on page 65.
6. Go back to the **application** page and see DR status changed to **Not protected** in Figure 3-27 on page 65.

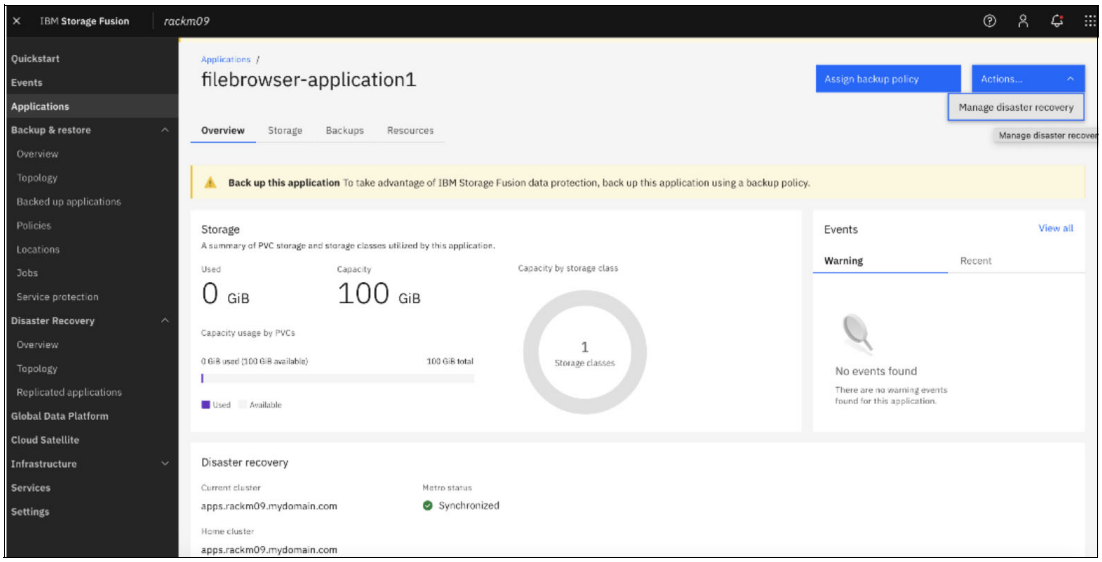

*Figure 3-23 Click on **Manage disaster recovery** on the **application details** page*

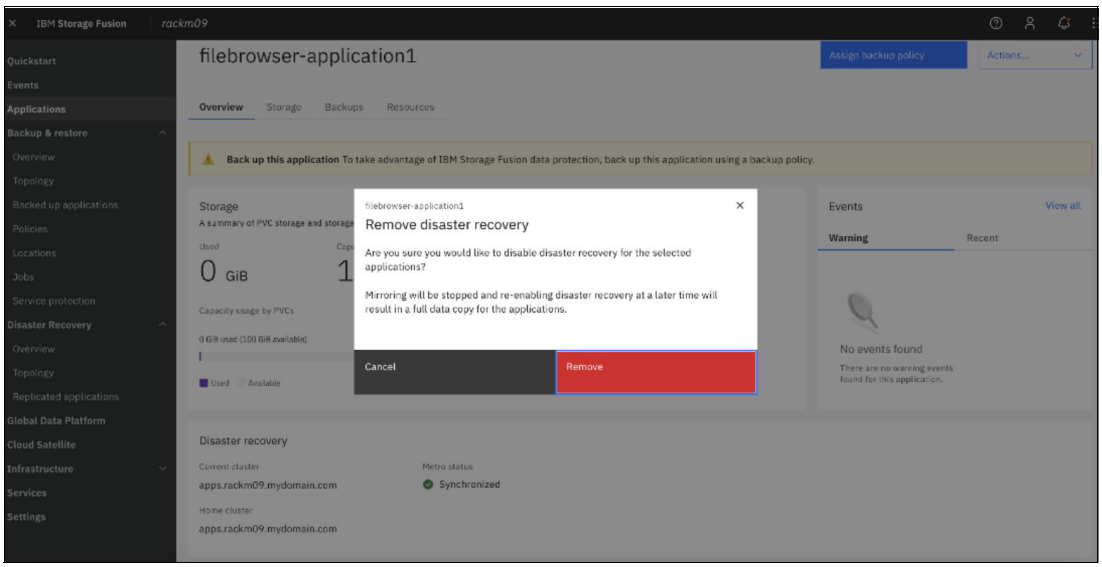

*Figure 3-24 Click on **remove** on the **Remove disaster recovery** popup window*

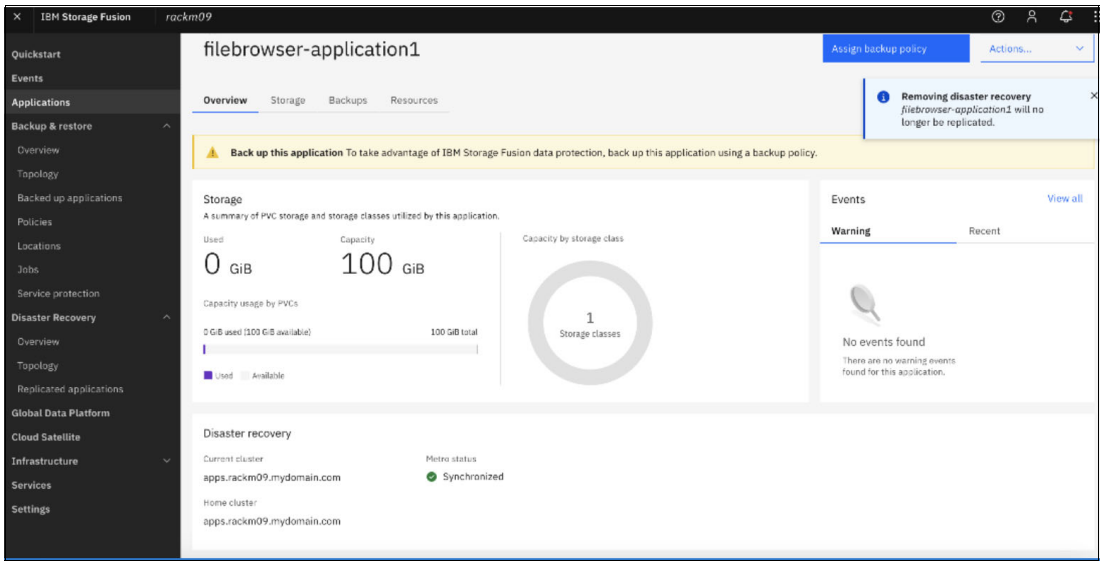

Figure 3-25 ***Removing disaster recovery*** *message displayed*

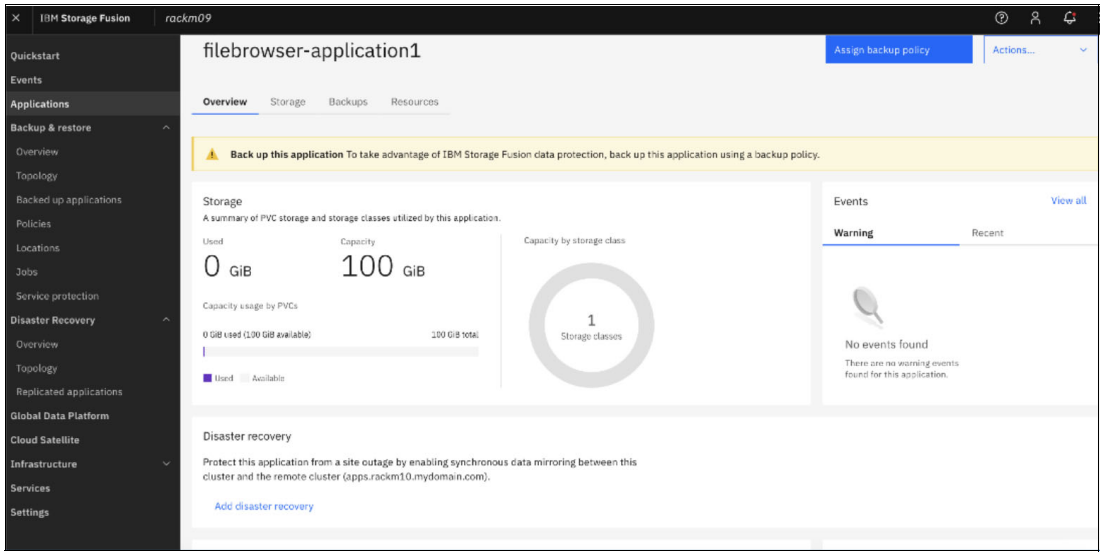

Figure 3-26 DR status removed from the ***Disaster Recovery*** *section*

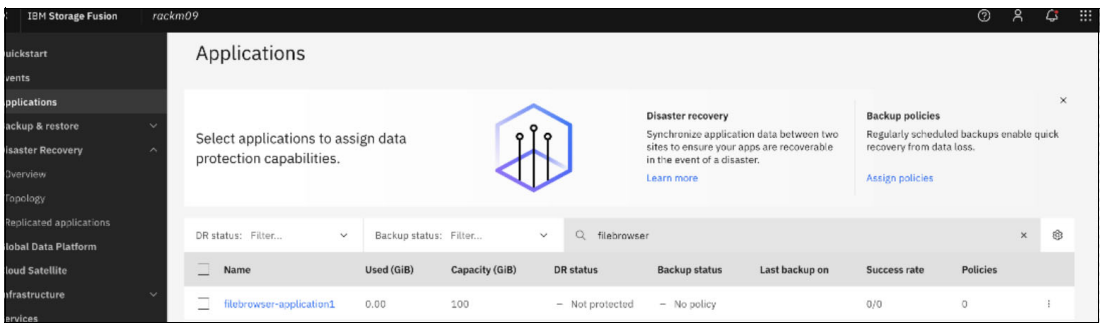

Figure 3-27 DR status changed to ***Not protected*** *on the* ***application*** *page*

Method 3: Remove disaster recovery from the Replicated application page

You can also disable an application from disaster recovery from the **Replicated applications** page with the following steps.

1. Click on the disaster Recovery.
2. Open the **Replicated applications** page.
3. Select an application by clicking a check box and click on **Manage disaster recovery** on the blue bar or click on the 3 dots and click on **Manage disaster recovery** on the drop down box as shown in Figure 3-28 and Figure 3-29.
4. Click on **No disaster recovery on the Manage disaster recovery** popup page and click **Save** as shown in Figure 3-30 on page 67.
5. Click on **Remove** on the **Remove disaster recovery** popup page as shown in Figure 3-31 on page 67.
6. You can see a message **Removing disaster recovery** displayed as shown in Figure 3-31 on page 67.
7. The application will be removed from the **Replicated applications** page as shown on Figure 3-32 on page 67.
8. DR status of the application will be changed to **Not protected** on the **Applications** page as shown in Figure 3-33 on page 68.

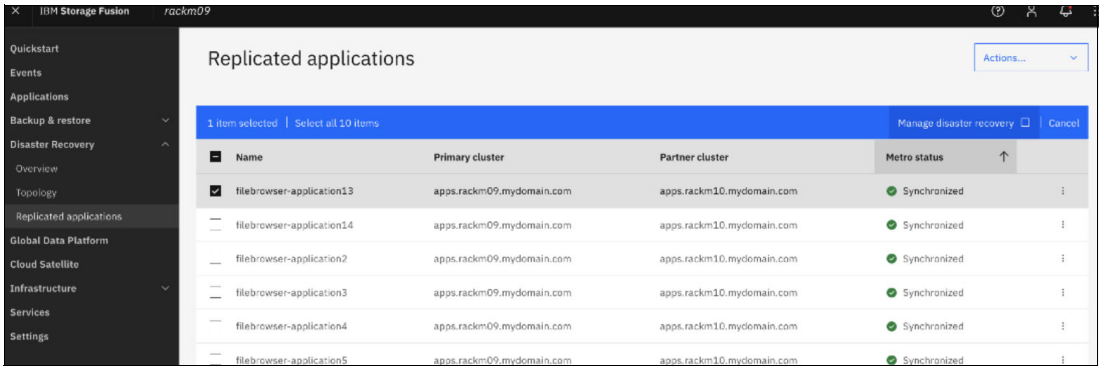

*Figure 3-28 Select an application and click on **Manage disaster recovery** on the blue bar*

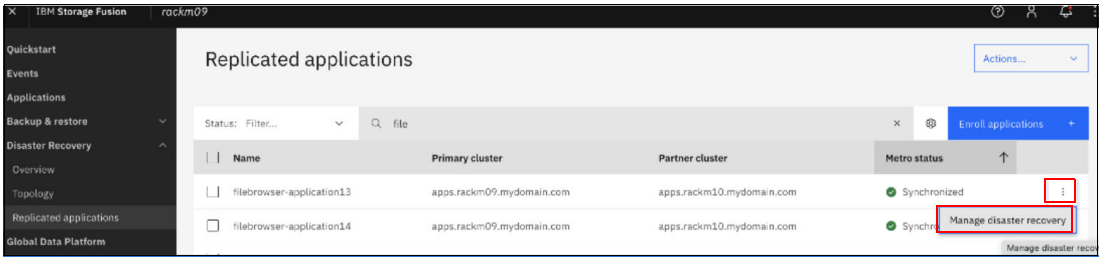

*Figure 3-29 Click on 3 dots and **manage disaster recovery***

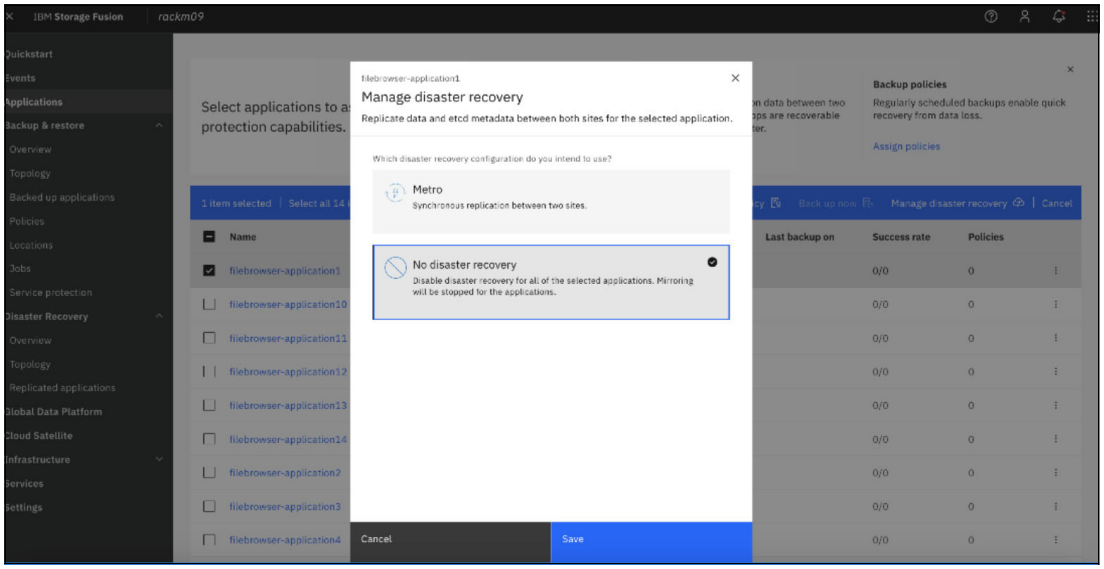

*Figure 3-30 Select **No disaster recovery** and click **Save***

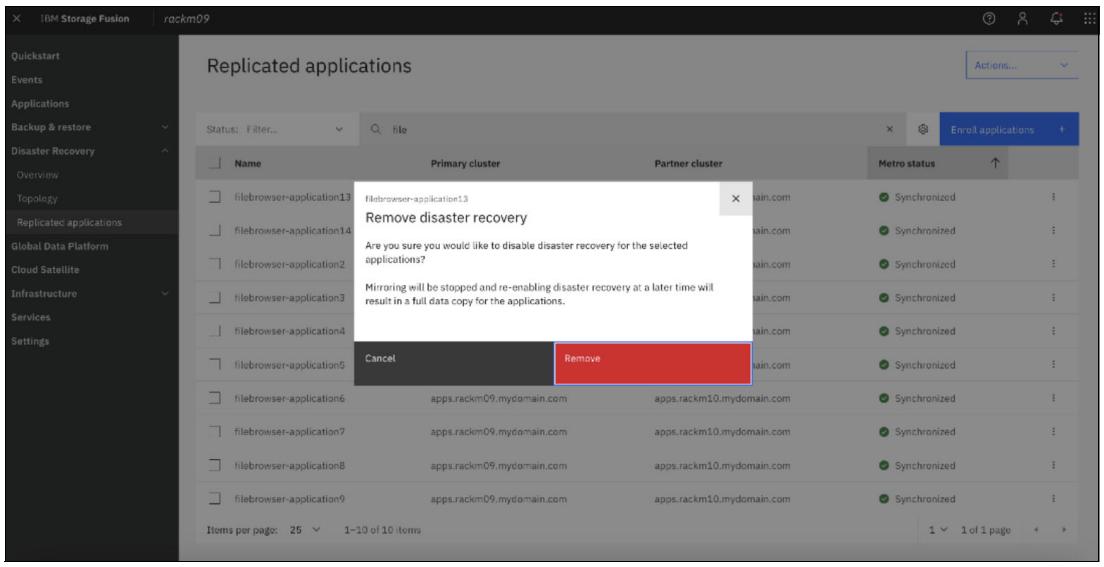

*Figure 3-31 Click on **Remove** to confirm*

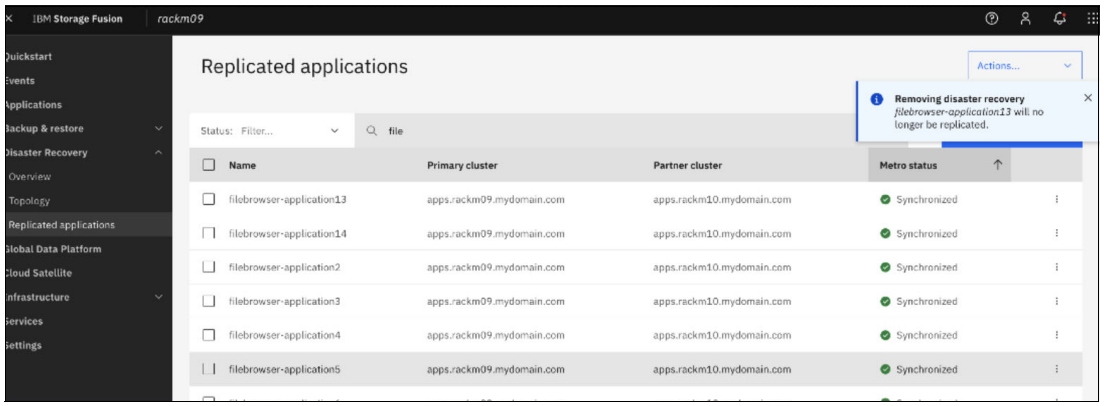

*Figure 3-32 **Removing disaster recovery** message displayed*

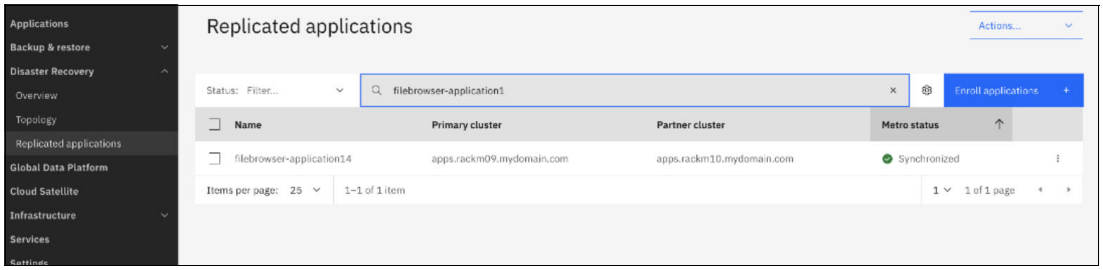

*Figure 3-33 Disaster recovery disabled application removed from **Replicated applications** page*

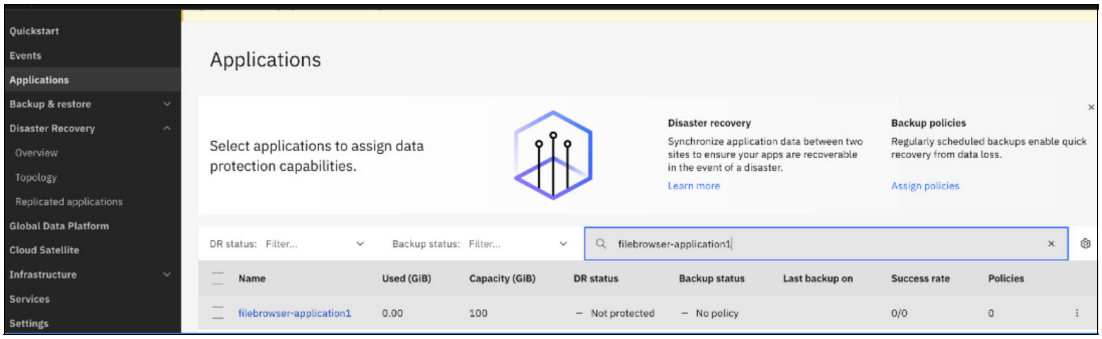

*Figure 3-34 DR status of the application moved to **Not protected***

Disabling Disaster recovery from multiple applications

We can use any of the following methods or options to disable multiple applications from DR:

- Applications page
- Replicated applications page

Method 1: Disable DR from the applications page

You can remove disaster recovery from an application from the **Applications** page with the following steps.

1. Click on the **Applications** page.
2. Select multiple applications as shown in Figure 3-35 on page 69.
3. Click on **Manage disaster recovery** on the blue bar as shown in Figure 3-36 on page 69.
4. Click on **No disaster recovery** and save it as shown in Figure 3-37 on page 69.
5. Confirm by clicking **Remove** on the **Remove disaster Recovery** popup window shown in Figure 3-38 on page 70.
6. Figure 3-39 on page 70 shows the **Removing disaster recovery** message.
7. Go back to the **Applications** page and see the DR status changed to **Not protected** as shown in Figure 3-40 on page 71.

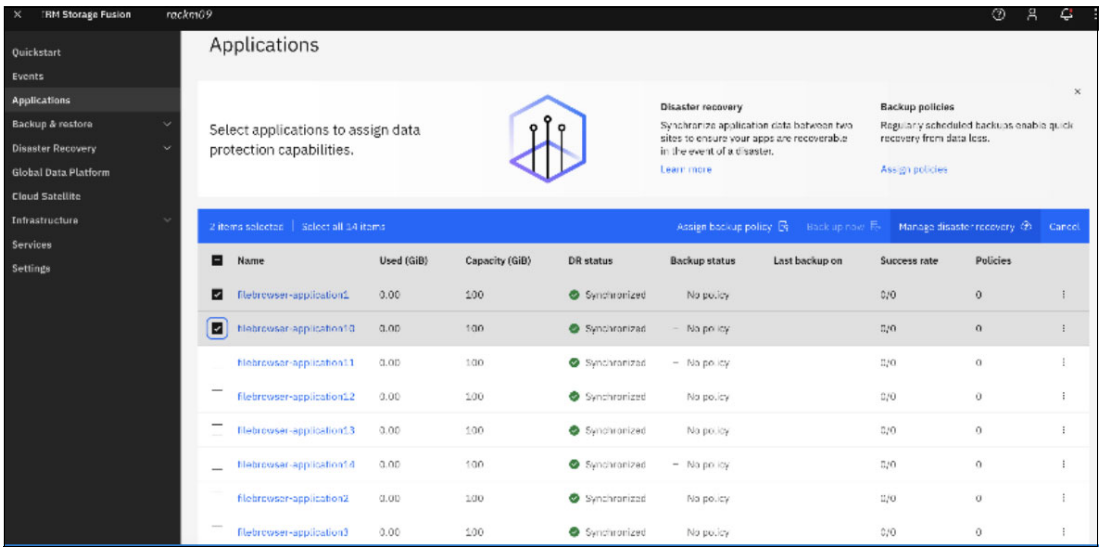

Figure 3-35 Select multiple applications by clicking check boxes

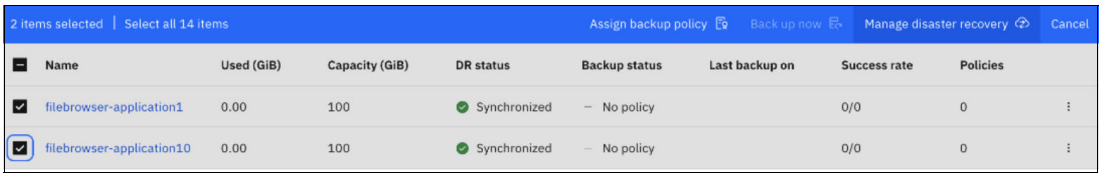

Figure 3-36 ***Click Manage disaster recovery*** *on blue bar*

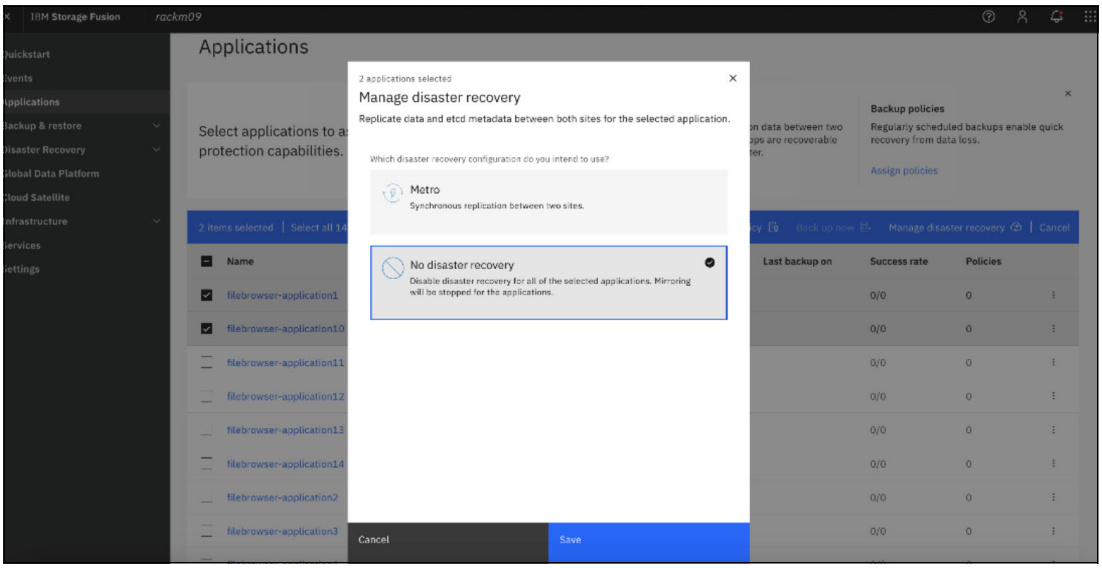

Figure 3-37 Select ***No disaster recovery*** *on the* ***Manage disaster recovery*** *popup window and click* ***save***

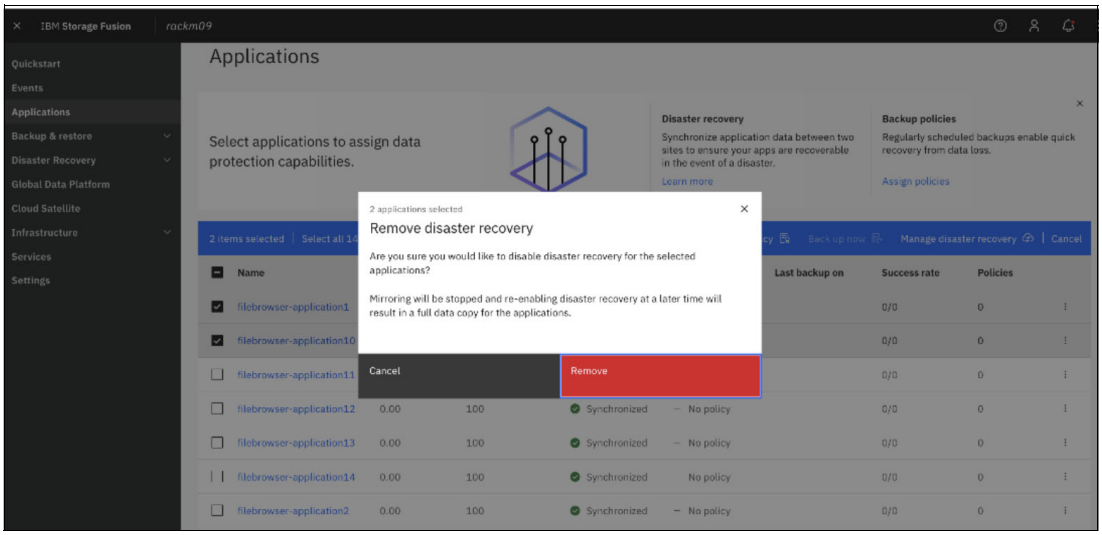

Figure 3-38 Click **Remove** to confirm

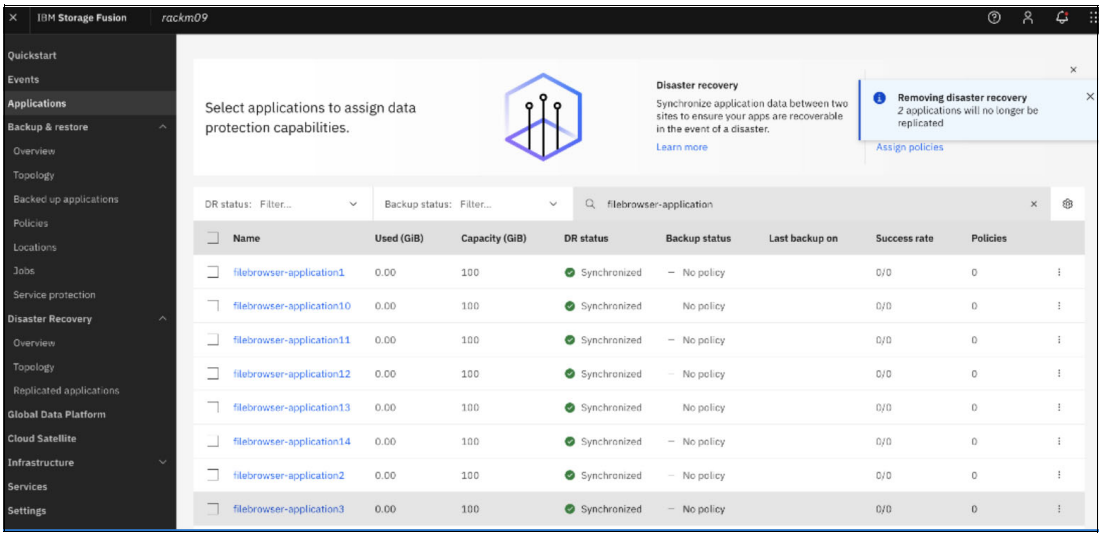

Figure 3-39 ***Removing disaster recovery*** *message displayed*

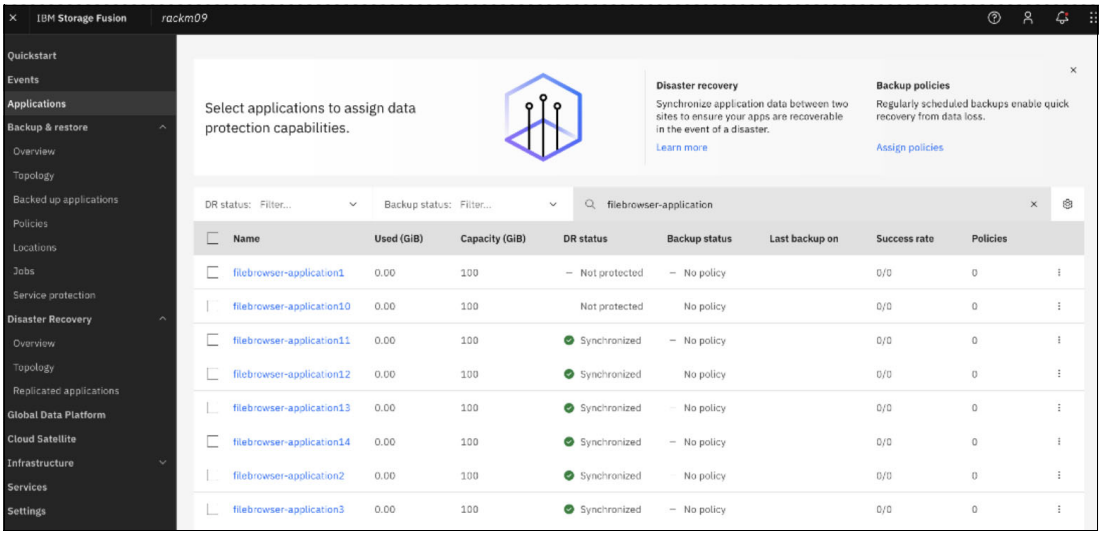

*Figure 3-40 DR status of applications changed to **Not protected***

Method 2: Remove disaster recovery from Replicated applications page

You can remove disaster recovery from multiple applications from the **Replicated applications** page with the by following the steps:

1. Click on the **Replicated applications** page.
2. Select multiple applications and click on **Manage disaster recovery** button on the blue bar as shown in figure Figure 3-41.
3. Click on **No disaster recovery** on the **Manage disaster recovery** popup window and click **Save** as shown in Figure 3-42 on page 72.
4. Click on **Remove** to disable the disaster recovery enrollment on the selected applications on the **Remove disaster recovery** popup window as shown in Figure 3-43 on page 72.
5. Disaster recovery disabled applications will be removed from the **Replicated applications** page as shown in figure Figure 3-44 on page 72.
6. Applications DR status changed to **Not protected** on the Applications page as shown in Figure 3-45 on page 73.

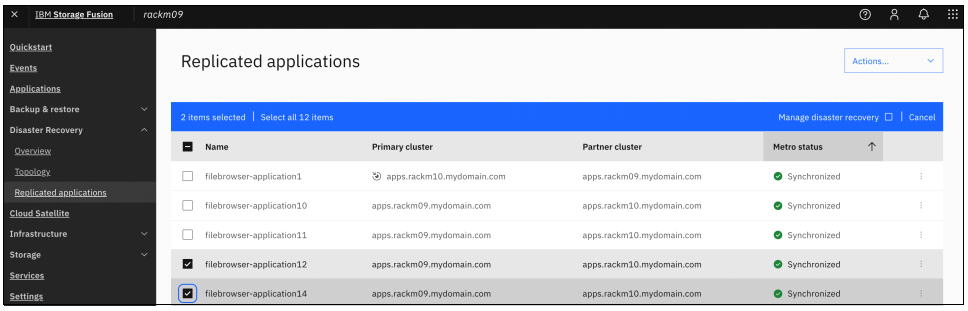

Figure 3-41 Select multiple applications and click on Manage disaster recovery button on the blue bar

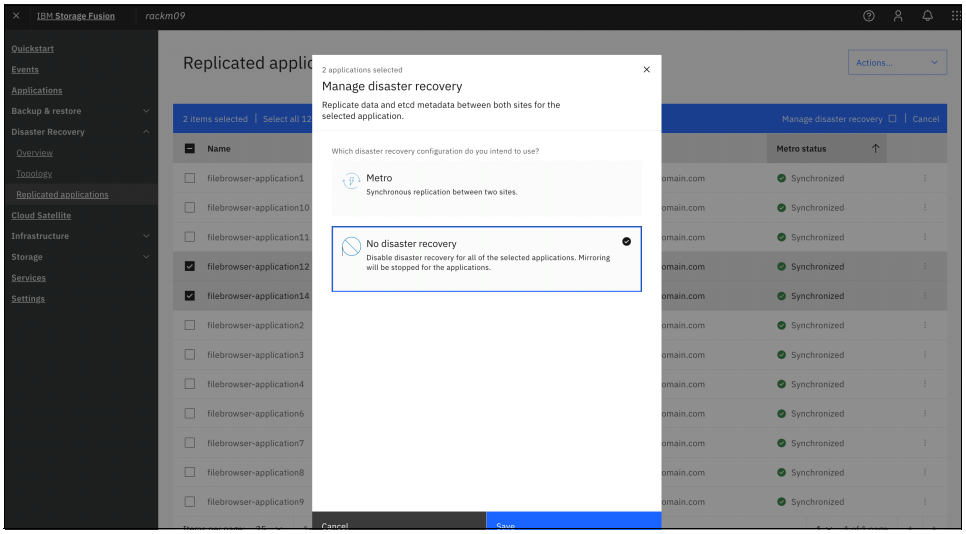

*Figure 3-42 Click on **No disaster recovery** and then **save***

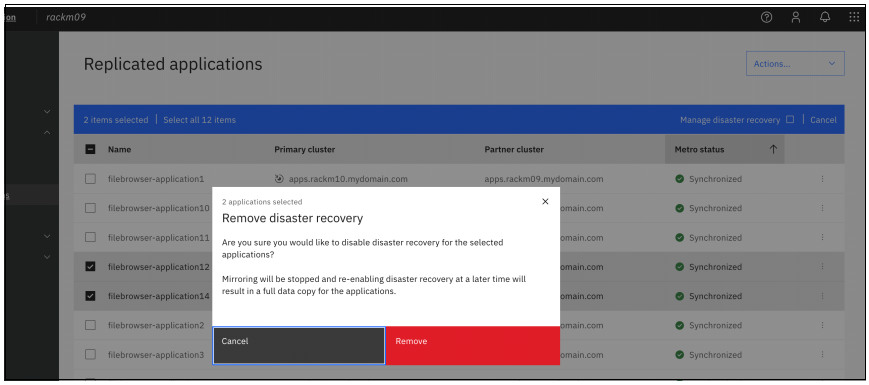

*Figure 3-43 Click on **Remove** to confirm*

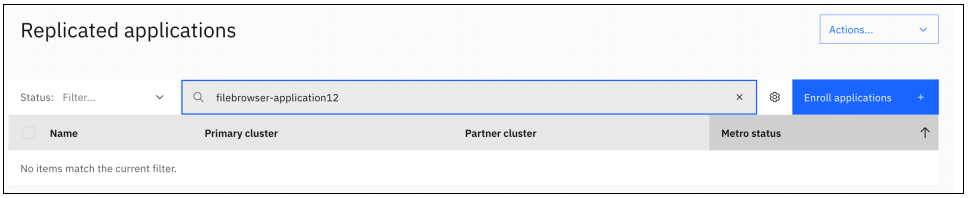

*Figure 3-44 Selected applications removed from the **Replicated applications** page*

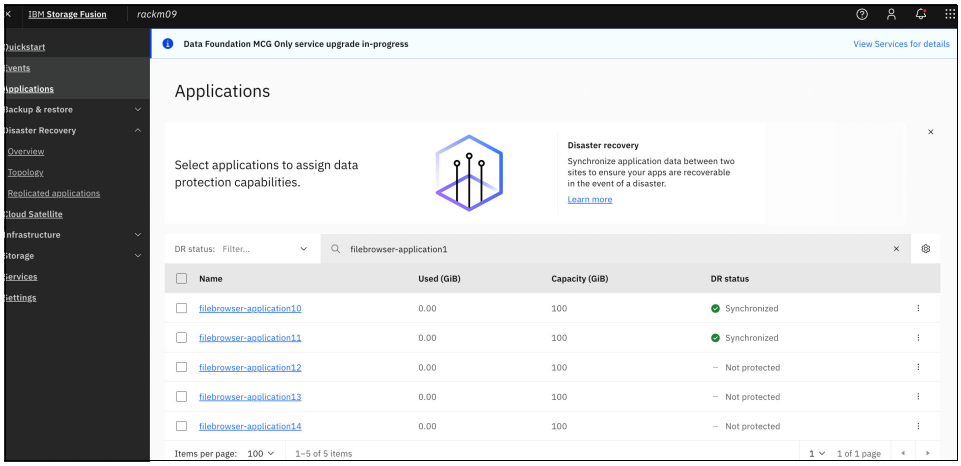

*Figure 3-45 DR status moved to **Not protected** for the DR disabled applications*

4

Application failover and failback

This chapter describes the steps to demonstrate the failover and failback of the file browser application between a local site and remote site with IBM Storage Fusion HCI.

4.1 Failover

The applications are enrolled for disaster recovery (DR) and the persistent volumes are shared across the sites. Just 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. Log in to IBM Storage Fusion of the remote site.
2. Go to the **Replicated applications** page.
3. Click on the **Actions** button.
4. Go to the **Failover** window.
5. Browse and select the applications as shown in Figure 4-1 on page 76.

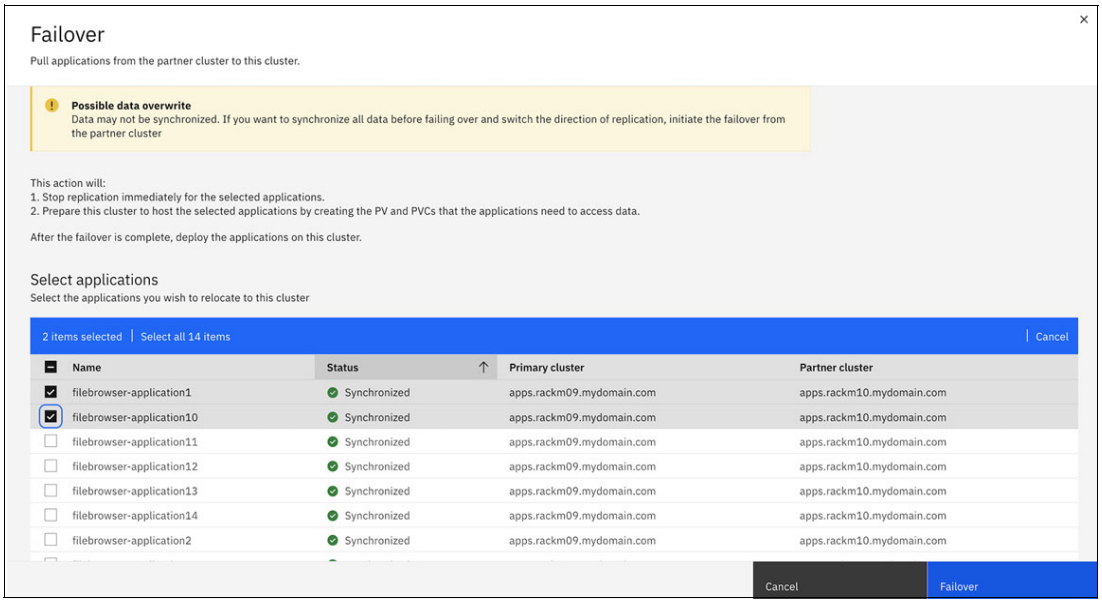

Figure 4-1 Select applications to be failed over on the failover window

Initiate the failover process

This section describes the steps of the failover process:

1. Go to **Disaster Recovery** and click on **Replicated application** page as shown in Figure 4-2.
2. Click on **Actions** and **Failover.**

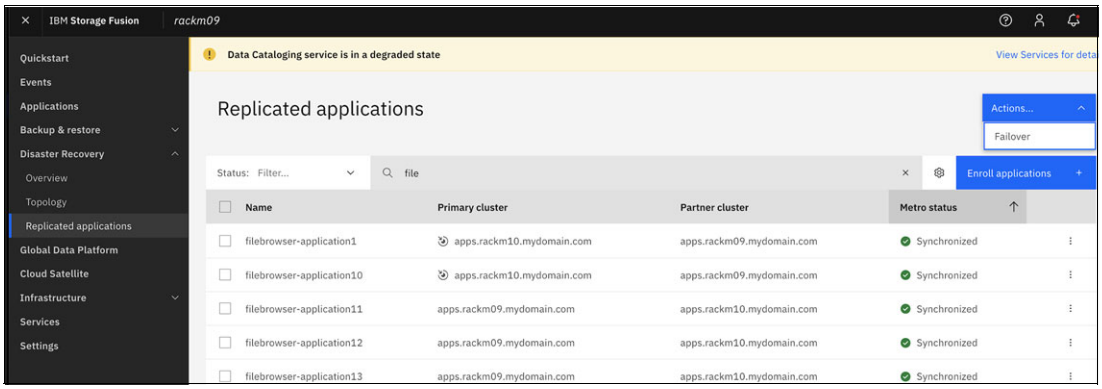

Figure 4-2 Failover process

3. Click on the **Failover** button to initiate failover. The dialog box will appear as shown in Figure 4-3 on page 77.

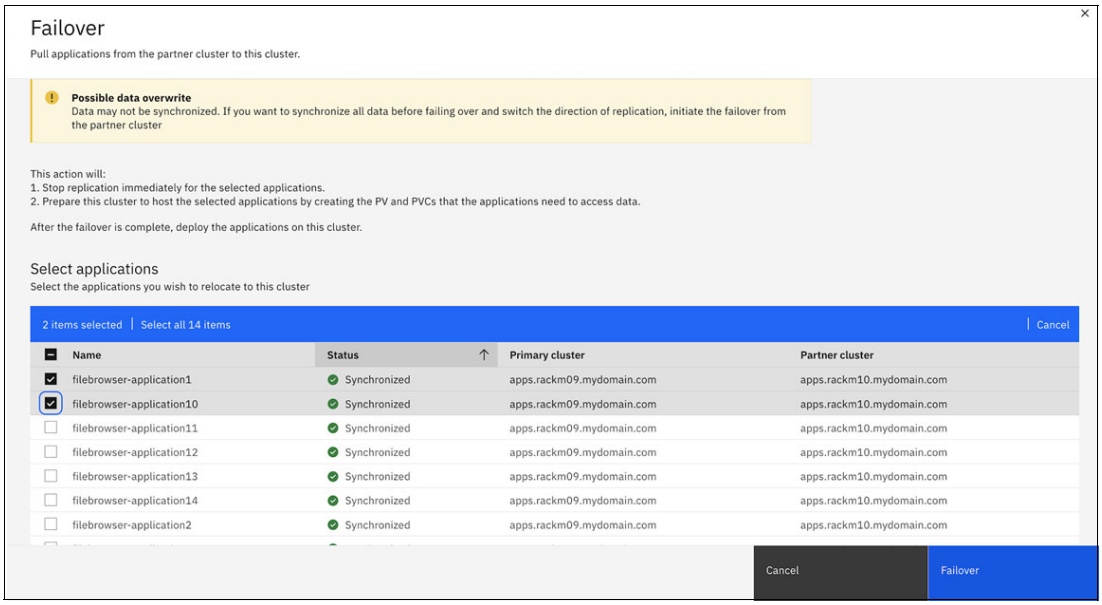

*Figure 4-3 Click on **Failover** to initiate failover*

4. **Failover initiated** message displayed on the **Replicated applications** page for the selected applications as shown in Figure 4-4.

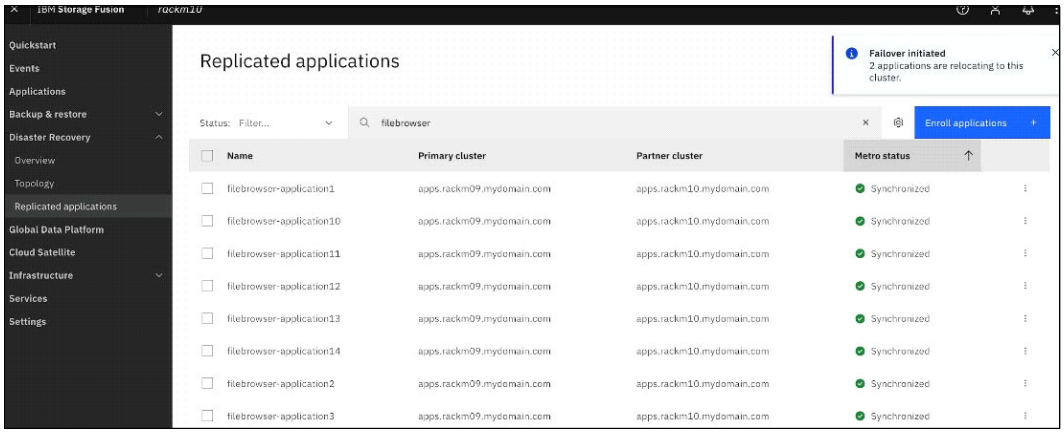

Figure 4-4 Confirmation message for failover process initiated

5. Once the failover is complete for the application, the **Failover complete** message is displayed as shown in Figure 4-5.

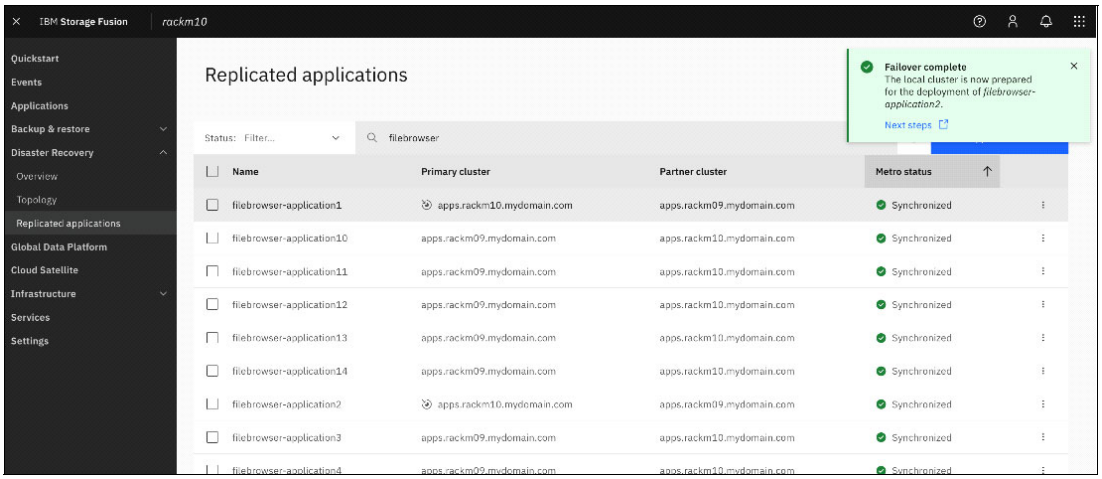

Figure 4-5 Failover completion message

6. The application status icon is updated under primary and partner cluster columns as shown in Figure 4-6.

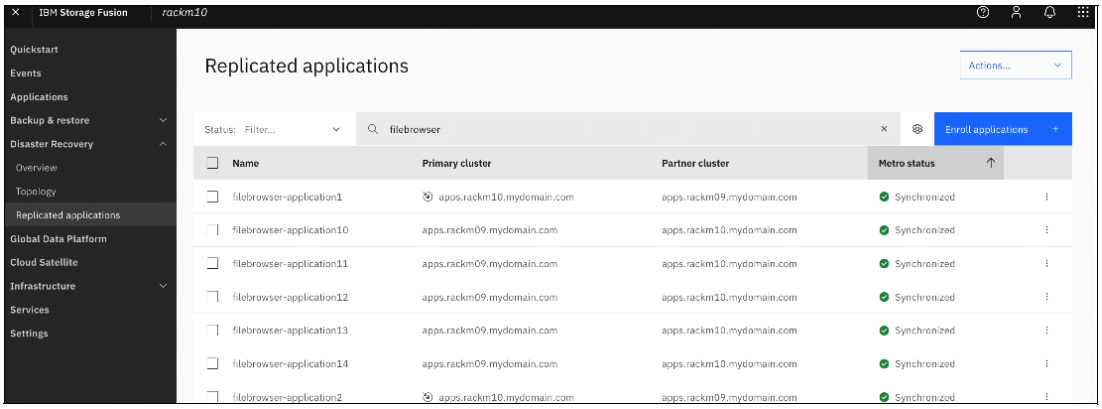

Figure 4-6 Applications status with icon updated under primary and partner cluster column

7. Now, redeploy the failed over application(s) on remote site from the OpenShift console as shown in Figure 4-7 on page 79.

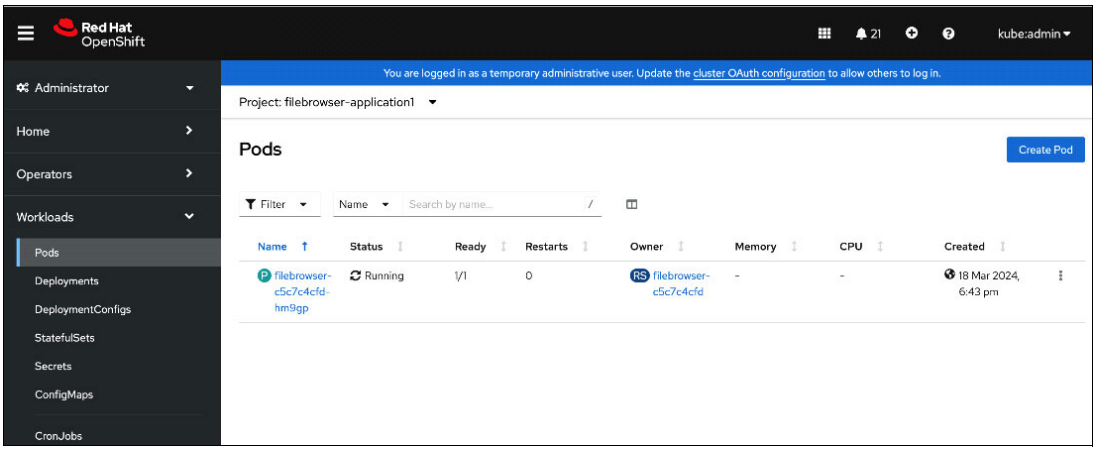

Figure 4-7 Remote site - Application redeployment from OpenShift console

8. Access redeployed applications from a remote site as shown in Figure 4-8.

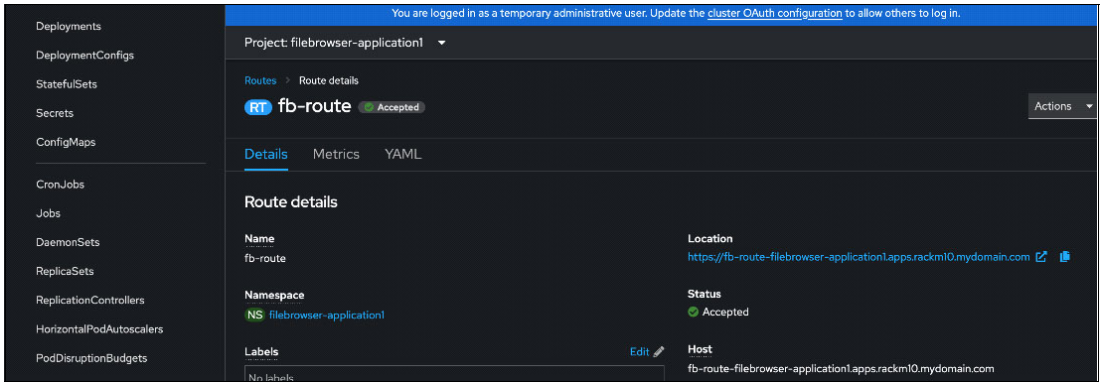

Figure 4-8 Remote site - Click on location on OpenShift Console to access the redeployed application

9. View failed over applications on a remote site as shown in Figure 4-9 on page 80.

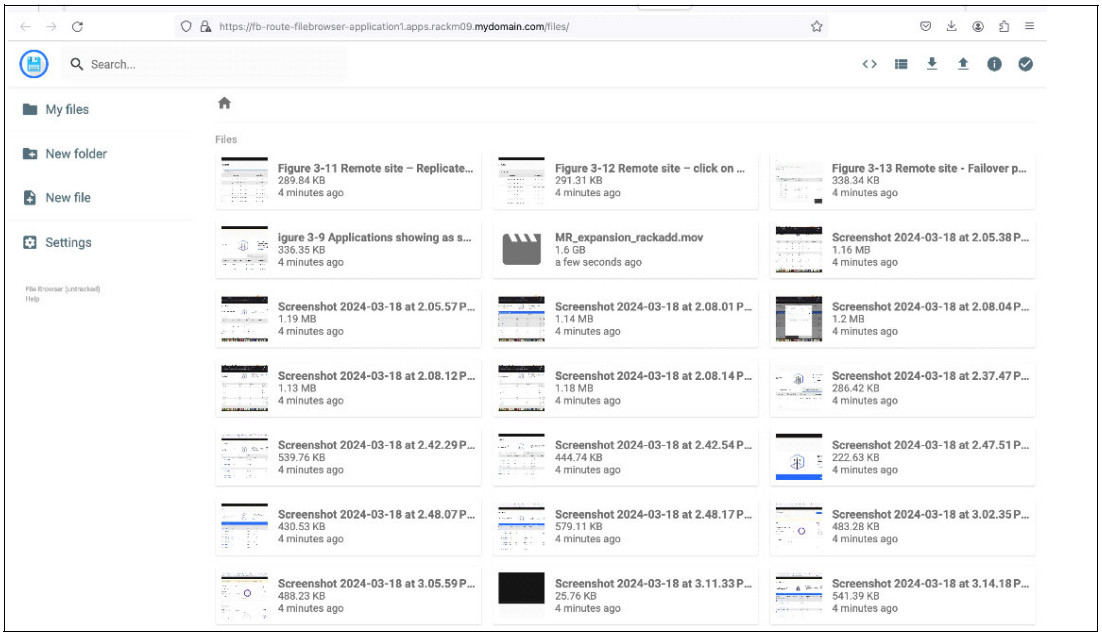

Figure 4-9 Application access from remote site

4.2 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 these steps:

1. Log in to IBM Storage Fusion of local site (Site1).
2. Go to **Replicated applications** page.
3. Click on **Action** → **Failover**, select the failed over applications and click **failover** in Figure 4-10.

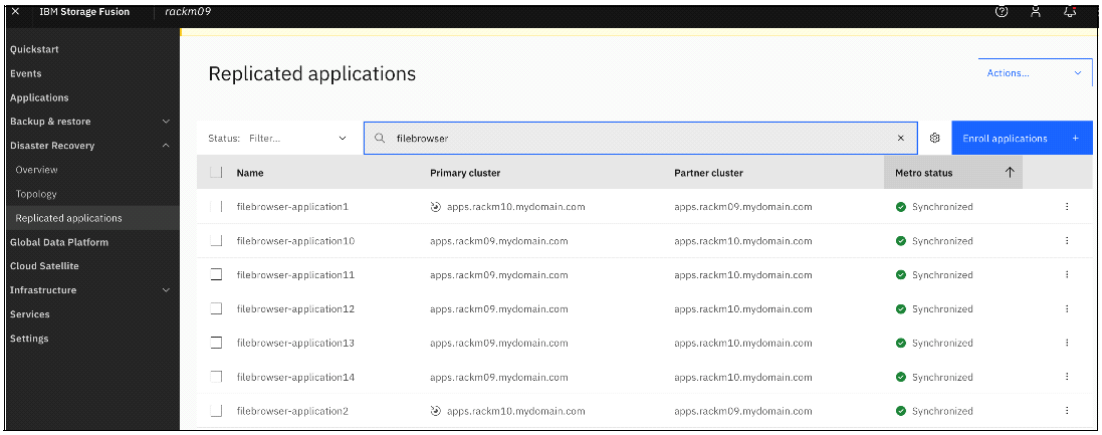

Figure 4-10 View failed over applications

Initiate failback process

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

1. Log in to IBM Storage Fusion of local site (Site1).
2. Go to the **Replicated applications** page.
3. Click **Actions** → **Failover**.
4. Select failover applications.
5. Click **Failover** button as shown in Figure 4-11.

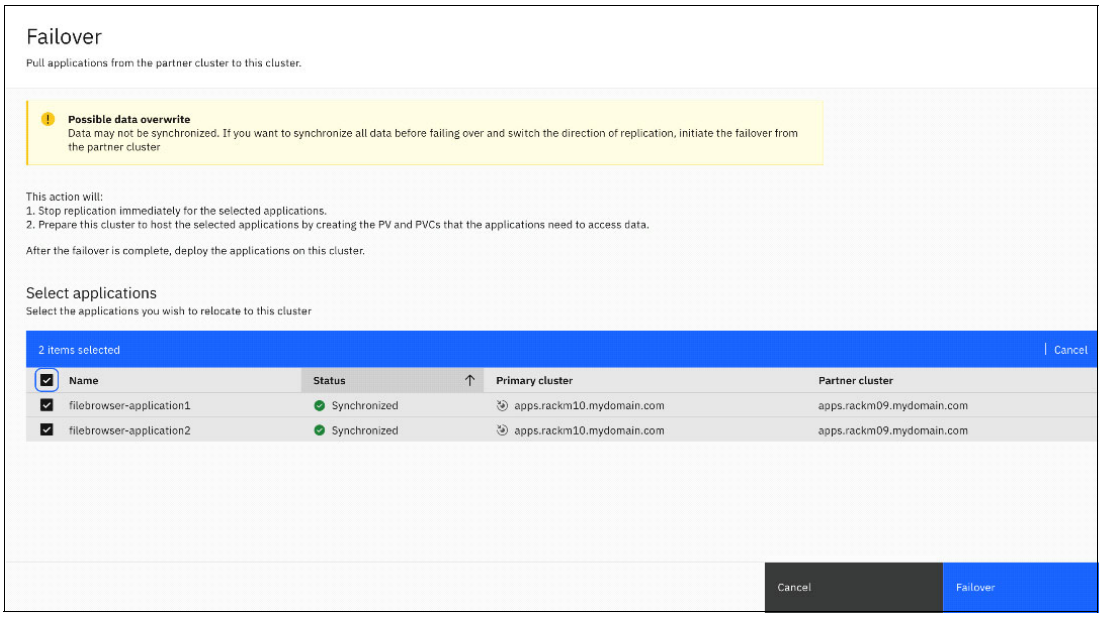

*Figure 4-11 Select applications and click on **Failover** to initiate failback of failover applications*

6. The message **failover initiated** is displayed for the failback initiated application(s) as shown in Figure 4-12.

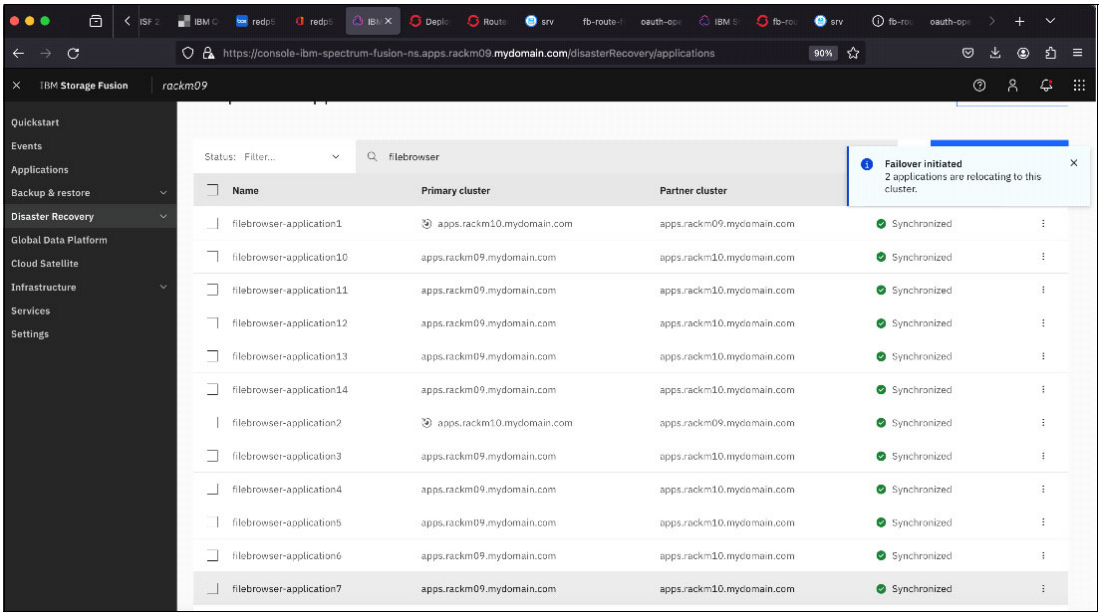

Figure 4-12 Failover initiated message displayed after initiating failback

7. After the completion of the failover for the selected applications, a **Failover complete** message is displayed as shown in Figure 4-13.

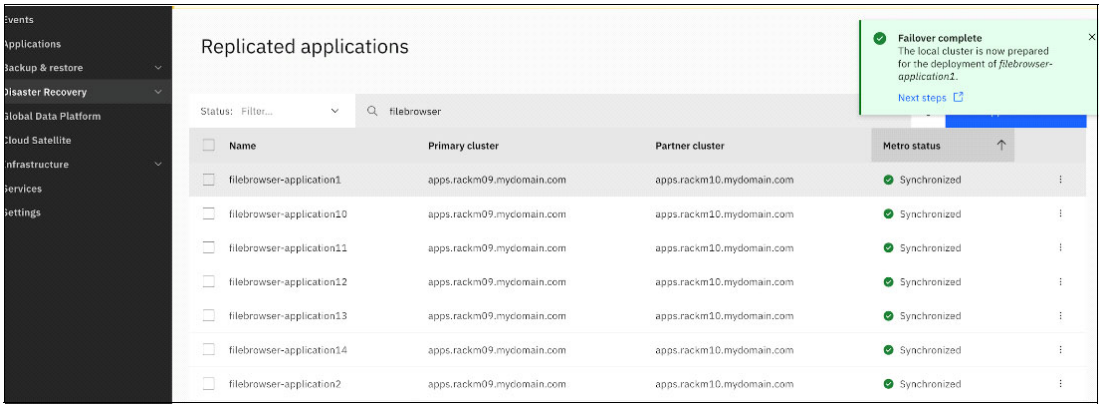

Figure 4-13 Failover complete message displayed for the selected applications

8. When the failback completes, the application status for primary and partner cluster is reflected back to the original status as shown in Figure 4-14.

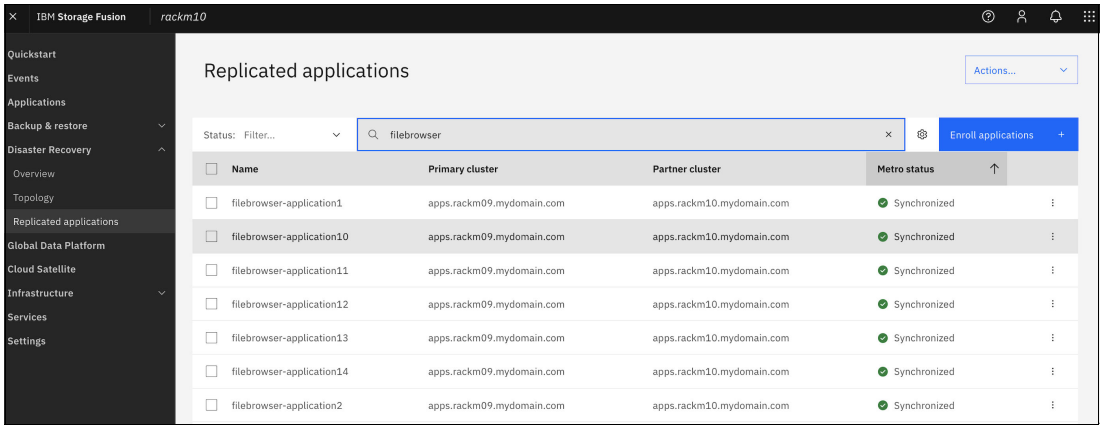

Figure 4-14 Failed back application(s) status under primary and partner cluster columns

9. View failed back applications on the OpenShift console for the failed back applications as shown in Figure 4-15 on page 83.

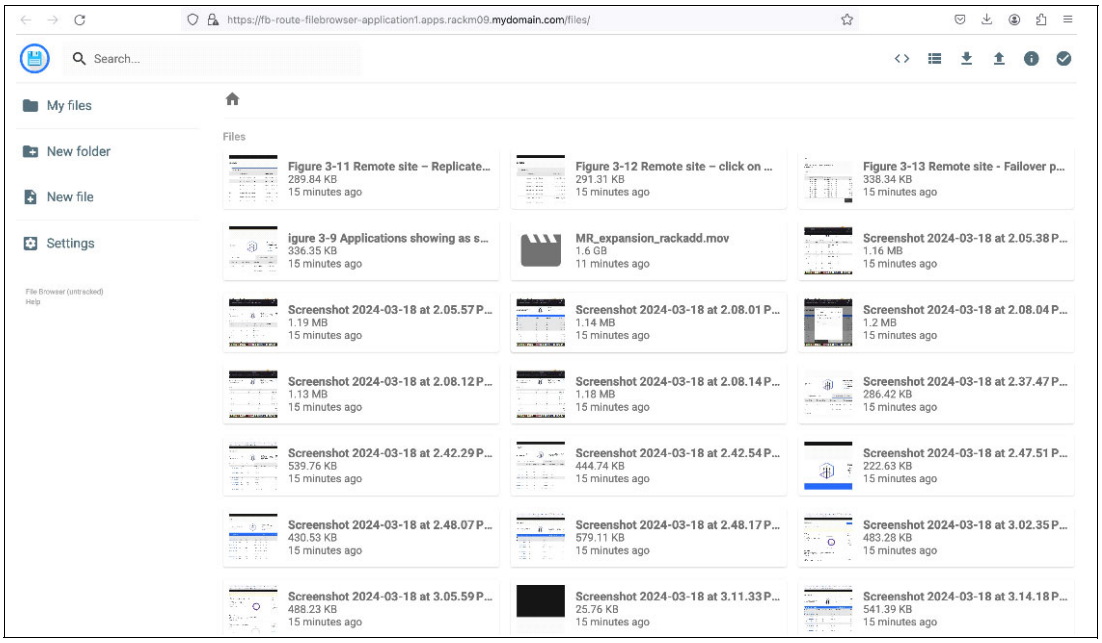

Figure 4-15 View failed back application(s) on local site

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.

- *Accelerating IBM watsonx.data with IBM Fusion HCI,* REDP-5720
- *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:

- IBM Storage

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

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

  https://www.ibm.com/products/storage-fusion
- IBM Storage Fusion HCI 2.8.x Metro-DR

  https://www.ibm.com/docs/en/sfhs/2.8.x?topic=disaster-recovery

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

REDP-5708-01

ISBN 0738461709

Printed in U.S.A.