



BlueXP Disaster Recovery

NetApp Solutions

NetApp
August 30, 2024

Table of Contents

- BlueXP Disaster Recovery 1
 - 3-2-1 Data Protection for VMware with SnapCenter Plug-in and BlueXP backup and recovery for VMs 1
 - DR using BlueXP DRaaS 42

BlueXP Disaster Recovery

3-2-1 Data Protection for VMware with SnapCenter Plug-in and BlueXP backup and recovery for VMs

The 3-2-1 backup strategy is an industry accepted data protection method, providing a comprehensive approach to safeguarding valuable data. This strategy is reliable and ensures that even if some unexpected disaster strikes, there will still be a copy of the data available.

Author: Josh Powell - NetApp Solutions Engineering

Overview

The strategy is comprised of three fundamental rules:

1. Keep at least three copies of your data. This ensures that even if one copy is lost or corrupted, you still have at least two remaining copies to fall back on.
2. Store two backup copies on different storage media or devices. Diversifying storage media helps protect against device-specific or media-specific failures. If one device gets damaged or one type of media fails, the other backup copy remains unaffected.
3. Finally, ensure that at least one backup copy is offsite. Offsite storage serves as a fail-safe against localized disasters like fires or floods that could render onsite copies unusable.

This solution document covers a 3-2-1 backups solution using SnapCenter Plug-in for VMware vSphere (SCV) to create primary and secondary backups of our on-premises virtual machines and BlueXP backup and recovery for virtual machines to backup a copy of our data to cloud storage or StorageGRID.





Use Cases

This solution addresses the following use cases:

- Backup and restore of on-premises virtual machines and datastores using SnapCenter Plug-in for VMware vSphere.
- Backup and restore of on-premises virtual machines and datastores, hosted on ONTAP clusters, and backed up to object storage using BlueXP backup and recovery for virtual machines.

NetApp ONTAP Data Storage

ONTAP is NetApp's industry leading storage solution that offers unified storage whether you access over SAN or NAS protocols. The 3-2-1 backup strategy ensures on-premises data is protected on more than one media type and NetApp offers platforms ranging from high-speed flash to lower-cost media.

FAS	AFF C-Series	AFF A-Series	ASA A-Series
			
Hybrid flash storage	Capacity all-flash storage	Performance all-flash storage	All-flash SAN storage
Unified (file, block, object)	Unified (file, block, object)	Unified (file, block, object)	Block optimized
Lowest price storage	Balanced price storage	Premium priced storage	Aggressively priced storage
Tier 2 @ 5-10ms latency Backup / Low-cost DR	Refresh of hybrid flash, Tier 1 @ 2-4ms latency Tier 2 workloads VMware datastores	Ideal for Tier 1 business-critical workloads with <1ms latency	Ideal for Tier 1 Block Six Nines Guaranteed

For more information on all of NetApp's hardware platform's check out [NetApp Data Storage](#).

SnapCenter Plug-in for VMware vSphere

The SnapCenter Plugin for VMware vSphere is a data protection offering which is tightly integrated with VMware vSphere and allows easy management of backup and restores for virtual machines. As part of that solution, SnapMirror provides a fast and reliable method to create a second immutable backup copy of virtual machine data on a secondary ONTAP storage cluster. With this architecture in place, virtual machine restore operations can easily be initiated from either the primary or secondary backup locations.

SCV is deployed as a linux virtual appliance using an OVA file. The plug-in now uses a remote plug-in architecture. The remote plug-in runs outside of the vCenter server and is hosted on the SCV virtual appliance.

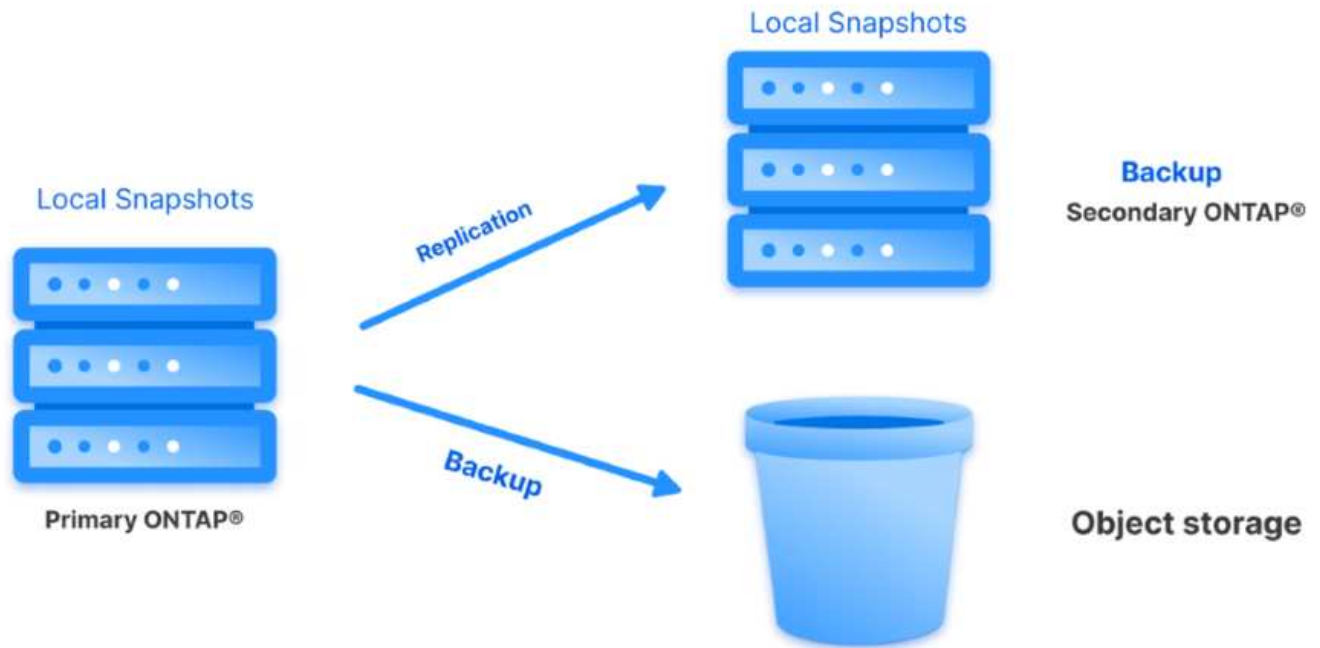
For detailed information on SCV refer to [SnapCenter Plug-in for VMware vSphere documentation](#).

BlueXP backup and recovery for virtual machines

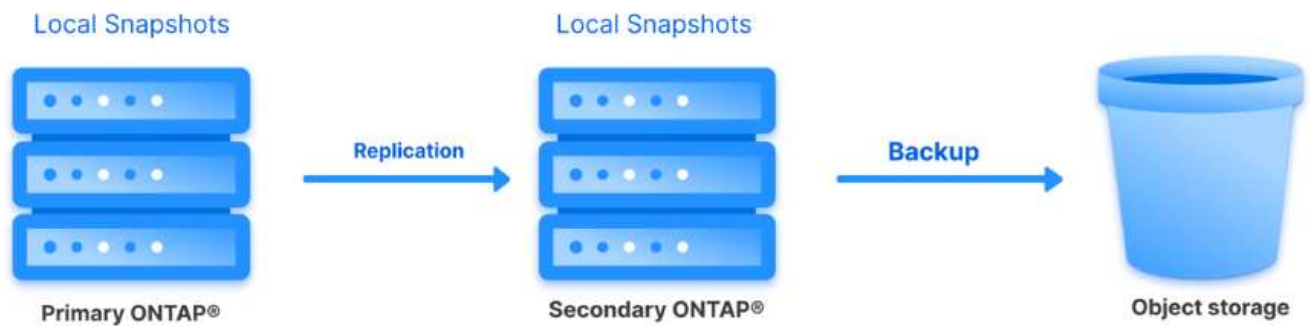
BlueXP backup and recovery is a cloud based tool for data management that provides a single control plane for a wide range of backup and recovery operations across both on-premises and cloud environments. Part of the NetApp BlueXP backup and recovery suite is a feature that integrates with the SnapCenter Plugin for VMware vSphere (on-premises) to extend a copy of the data to object storage in the cloud. This establishes a third copy of the data offsite that is sourced from the primary or secondary storage backups. BlueXP backup and recovery makes it easy to set up storage policies that transfer copies of your data from either of these two on-prem locations.

Choosing between the primary and secondary backups as the source in BlueXP Backup and Recovery will result in one of two topologies being implemented:

Fan-out Topology – When a backup is initiated by the SnapCenter Plug-in for VMware vSphere, a local snapshot is immediately taken. SCV then initiates a SnapMirror operation that replicates the most recent snapshot to the Secondary ONTAP cluster. In BlueXP Backup and Recovery, a policy specifies the primary ONTAP cluster as the source for a snapshot copy of the data to be transferred to object storage in your cloud provider of choice.



Cascading Topology – Creating the primary and secondary data copies using SCV is identical to the fan-out topology mentioned above. However, this time a policy is created in BlueXP Backup and Recovery specifying that the backup to object storage will originate from the secondary ONTAP cluster.



BlueXP backup and recovery can create backup copies of on-premises ONTAP snapshots to AWS Glacier, Azure Blob, and GCP Archive storage.



AWS Glacier and Deep Glacier



Azure Blob Archive



GCP Archive Storage

In addition, you can use NetApp StorageGRID as the object storage backup target. For more on StorageGRID refer to the [StorageGRID landing page](#).

Solution Deployment Overview

This list provides the high level steps necessary to configure this solution and execute backup and restore operations from SCV and BlueXP backup and recovery:

1. Configure SnapMirror relationship between the ONTAP clusters to be used for primary and secondary data copies.
2. Configure SnapCenter Plug-In for VMware vSphere.
 - a. Add Storage Systems
 - b. Create backup policies
 - c. Create resource groups
 - d. Run backup first backup jobs
3. Configure BlueXP backup and recovery for virtual machines
 - a. Add working environment
 - b. Discover SCV and vCenter appliances
 - c. Create backup policies
 - d. Activate backups
4. Restore virtual machines from primary and secondary storage using SCV.
5. Restore virtual machines from object storage using BlueXP backup and restore.


Prerequisites

The purpose of this solution is to demonstrate data protection of virtual machines running in VMware vSphere and located on NFS Datastores hosted by NetApp ONTAP. This solution assumes the following components are configured and ready for use:

1. ONTAP storage cluster with NFS or VMFS datastores connected to VMware vSphere. Both NFS and VMFS datastores are supported. NFS datastores were utilized for this solution.
2. Secondary ONTAP storage cluster with SnapMirror relationships established for volumes used for NFS datastores.
3. BlueXP connector installed for cloud provider used for object storage backups.
4. Virtual machines to be backed are on NFS datastores residing on the primary ONTAP storage cluster.
5. Network connectivity between the BlueXP connector and on-premises ONTAP storage cluster management interfaces.
6. Network connectivity between the BlueXP connector and on-premises SCV appliance VM and between the BlueXP connector and vCenter.
7. Network connectivity between the on-premises ONTAP intercluster LIFs and the object storage service.
8. DNS configured for management SVM on primary and secondary ONTAP storage clusters. For more information refer to [Configure DNS for host-name resolution](#).

High Level Architecture

The testing / validation of this solution was performed in a lab that may or may not match the final deployment environment.



Solution Deployment

In this solution, we provide detailed instructions for deploying and validating a solution that utilizes SnapCenter Plug-in for VMware vSphere, along with BlueXP backup and recovery, to perform the backup and recovery of Windows and Linux virtual machines within a VMware vSphere cluster located in an on-premises data center. The virtual machines in this setup are stored on NFS datastores hosted by an ONTAP A300 storage cluster. Additionally, a separate ONTAP A300 storage cluster serves as a secondary destination for volumes replicated using SnapMirror. Furthermore, object storage hosted on Amazon Web Services and Azure Blob were employed as targets for a third copy of the data.

We will go over creating SnapMirror relationships for secondary copies of our backups managed by SCV and configuration of backup jobs in both SCV and BlueXP backup and recovery.

For detailed information on SnapCenter Plug-in for VMware vSphere refer to the [SnapCenter Plug-in for VMware vSphere documentation](#).

For detailed information on BlueXP backup and recovery refer to the [BlueXP backup and recovery documentation](#).

Establish SnapMirror relationships between ONTAP Clusters

SnapCenter Plug-in for VMware vSphere uses ONTAP SnapMirror technology to manage the transport of secondary SnapMirror and/or SnapVault copies to a secondary ONTAP Cluster.

SCV backup policies have the option of using SnapMirror or SnapVault relationships. The primary difference is that when using the SnapMirror option, the retention schedule configured for backups in the policy will be the same at the primary and secondary locations. SnapVault is designed for archiving and when using this option a separate retention schedule can be established with the SnapMirror relationship for the snapshot copies on the secondary ONTAP storage cluster.

Setting up SnapMirror relationships can be done in BlueXP where many of the steps are automated, or it can be done using System Manager and the ONTAP CLI. All of these methods are discussed below.

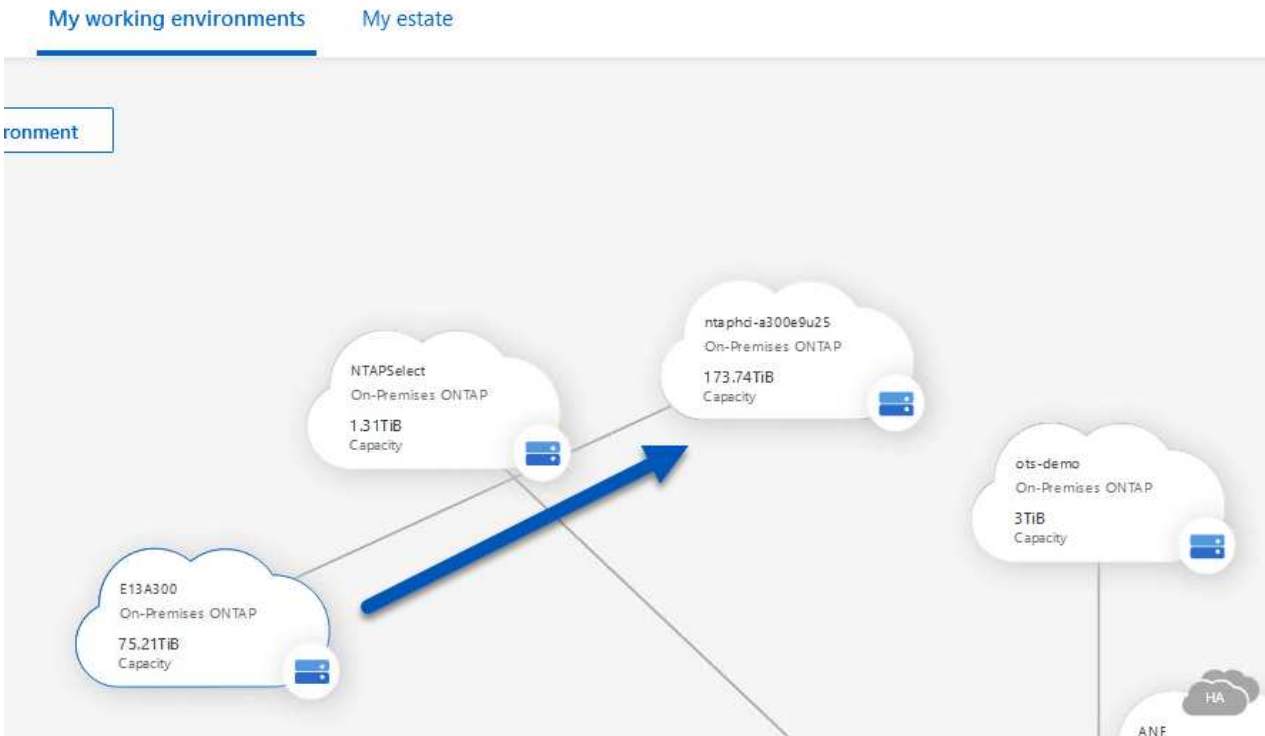
Establish SnapMirror relationships with BlueXP

The following steps must be completed from the BlueXP web console:

Replication setup for primary and secondary ONTAP storage systems

Begin by logging into the BlueXP web console and navigating to the Canvas.

1. Drag and drop the source (primary) ONTAP storage system onto the destination (secondary) ONTAP storage system.



2. From the menu that appears select **Replication**.



3. On the **Destination Peering Setup** page select the destination Intercluster LIFs to be used for the connection between storage systems.

Select the destination LIFs you would like to use for cluster peering setup.
 Replication requires an initial connection between the two working environments which is called a cluster peer relationship.
 For more information about LIF selections, see Cloud Manager documentation.

<input type="checkbox"/> CVO_InterCluster_B ntaphci-a300-02 : a0a-3510 172.21.254.21/24 up	<input type="checkbox"/> CVO_InterCluster_A ntaphci-a300-01 : a0a-3510 172.21.254.21/24 up	<input type="checkbox"/> zoneb-n1 ntaphci-a300-01 : a0a-3484 172.21.228.21/24 up	<input type="checkbox"/> zoneb-n2 ntaphci-a300-02 : a0a-3484 172.21.228.22/24 up	<input checked="" type="checkbox"/> intercluster_node_1 ntaphci-a300-01 : a0a-181 10.61.181.193/24 up	<input checked="" type="checkbox"/> intercluster_node_2 ntaphci-a300-01 : a0a-181 10.61.181.194/24 up
--	--	--	--	---	---

- On the **Destination Volume Name** page, first select the source volume and then fill out the destination volume name and select the destination SVM and aggregate. Click on **Next** to continue.

Select the volume that you want to replicate



288 Volumes

<p>CDM01 ONLINE</p> <p>INFO</p> <table border="0"> <tr><td>Storage VM Name</td><td>FS02</td></tr> <tr><td>Tiering Policy</td><td>None</td></tr> <tr><td>Volume Type</td><td>RW</td></tr> </table> <p>CAPACITY</p> <p>206 GB Allocated</p> <p>53.72 MB Disk Used</p>	Storage VM Name	FS02	Tiering Policy	None	Volume Type	RW	<p>Data ONLINE</p> <p>INFO</p> <table border="0"> <tr><td>Storage VM Name</td><td>FS02</td></tr> <tr><td>Tiering Policy</td><td>None</td></tr> <tr><td>Volume Type</td><td>RW</td></tr> </table> <p>CAPACITY</p> <p>512 GB Allocated</p> <p>0 GB Disk Used</p>	Storage VM Name	FS02	Tiering Policy	None	Volume Type	RW
Storage VM Name	FS02												
Tiering Policy	None												
Volume Type	RW												
Storage VM Name	FS02												
Tiering Policy	None												
Volume Type	RW												
<p>Demo ONLINE</p> <p>INFO</p> <table border="0"> <tr><td>Storage VM Name</td><td>zonea</td></tr> <tr><td>Tiering Policy</td><td>None</td></tr> <tr><td>Volume Type</td><td>RW</td></tr> </table> <p>CAPACITY</p> <p>250 GB Allocated</p> <p>1.79 GB Disk Used</p>	Storage VM Name	zonea	Tiering Policy	None	Volume Type	RW	<p>Demo02_01 ONLINE</p> <p>INFO</p> <table border="0"> <tr><td>Storage VM Name</td><td>Demo</td></tr> <tr><td>Tiering Policy</td><td>None</td></tr> <tr><td>Volume Type</td><td>RW</td></tr> </table> <p>CAPACITY</p> <p>500 GB Allocated</p> <p>34.75 MB Disk Used</p>	Storage VM Name	Demo	Tiering Policy	None	Volume Type	RW
Storage VM Name	zonea												
Tiering Policy	None												
Volume Type	RW												
Storage VM Name	Demo												
Tiering Policy	None												
Volume Type	RW												

Destination Volume Name

Destination Volume Name

Demo_copy

Destination Storage VM

EHC_NFS

Destination Aggregate

EHCaggr01

5. Choose the max transfer rate for replication to occur at.

Max Transfer Rate

You should limit the transfer rate. An unlimited rate might negatively impact the performance of other applications and it might impact your Internet performance.

- Limited to: MB/s
- Unlimited (recommended for DR only machines)

6. Choose the policy that will determine the retention schedule for secondary backups. This policy can be created beforehand (see the manual process below in the **Create a snapshot retention policy** step) or can be changed after the fact if desired.

↑ Previous Step

Default Policies

Additional Policies

CloudBackupService-1674046623282

Original Policy Name: CloudBackupService-1674046623282

Creates a SnapVault relationship which replicates Snapshot copies with the following labels to the destination volume:
hourly (12), daily (15), weekly (6)
(# of retained Snapshot copies in parenthesis)

CloudBackupService-1674047424679

Custom Policy - No Comment

More info

CloudBackupService-1674047718637


Custom Policy - No Comment

More info

7. Finally, review all information and click on the **Go** button to start the replication setup process.

↑ Previous Step

Review your selection and start the replication process



Source Volume Allocated Size:	250 GB	Destination Aggregate:	EHCaggr01
Source Volume Used Size:	1.79 GB	Destination Storage VM:	EHC_NFS
Source Thin Provisioning:	Yes	Max Transfer Rate:	100 MB/s
Destination Volume Allocated Size:	250 GB	SnapMirror Policy:	Mirror
Destination Thin Provisioning:	No	Replication Schedule:	One-time copy

Establish SnapMirror relationships with System Manager and ONTAP CLI

All required steps for establishing SnapMirror relationships can be accomplished with System Manager or the ONTAP CLI. The following section provides detailed information for both methods:

Record the source and destination Intercluster logical interfaces

For the source and destination ONTAP clusters, you can retrieve the inter-cluster LIF information from System Manager or from the CLI.

1. In ONTAP System Manager, navigate to the Network Overview page and retrieve the IP addresses of Type: Intercluster that are configured to communicate with the AWS VPC where FSx is installed.

Name	Status	Storage VM	IPspace	Address	Current Node	Current Port	Portset	Protocols	Type	Thru
vveam_repo	✓	Backup	Default	10.61.181.179	E13A300_1	a0a-181		SMB/CIFS, NFS, S3	Data	0
CM01	✓		Default	10.61.181.180	E13A300_1	a0a-181			Cluster/Node Mgmt	0
HC_N1	✓		Default	10.61.181.183	E13A300_1	a0a-181			Intercluster,Cluster/Node Mgmt	0
HC_N2	✓		Default	10.61.181.184	E13A300_2	a0a-181			Intercluster,Cluster/Node Mgmt	0
SE_ora_vvm_614	✓	ora_vvm	Default	10.61.181.185	E13A300_1	a0a-181		SMB/CIFS, NFS, FL...	Data	0

2. To retrieve the Intercluster IP addresses using the CLI run the following command:

```
ONTAP-Dest::> network interface show -role intercluster
```

Establish cluster peering between ONTAP clusters

To establish cluster peering between ONTAP clusters, a unique passphrase entered at the initiating ONTAP cluster must be confirmed in the other peer cluster.

1. Set up peering on the destination ONTAP cluster using the `cluster peer create` command. When prompted, enter a unique passphrase that is used later on the source cluster to finalize the creation process.

```
ONTAP-Dest::> cluster peer create -address-family ipv4 -peer-addr  
source_intercluster_1, source_intercluster_2  
Enter the passphrase:  
Confirm the passphrase:
```

2. At the source cluster, you can establish the cluster peer relationship using either ONTAP System Manager or the CLI. From ONTAP System Manager, navigate to Protection > Overview and select Peer Cluster.

- DASHBOARD
- STORAGE ^
 - Overview
 - Volumes
 - LUNs
 - Consistency Groups
 - NVMe Namespaces
 - Shares
 - Buckets
 - Qtrees
 - Quotas
 - Storage VMs
 - Tiers
- NETWORK ^
 - Overview
 - Ethernet Ports
 - FC Ports
- EVENTS & JOBS ∨
- PROTECTION ^
 - Overview 1
 - Relationships
- HOSTS ∨

Overview

< Intercluster Settings

Network Interfaces

- IP ADDRESS
- ✓ 10.61.181.184
 - ✓ 172.21.146.217
 - ✓ 10.61.181.183
 - ✓ 172.21.146.216

Cluster Peers

- PEERED CLUSTER NAME
- ✓ FsxId0ae40e08acc0dea67
 - ✓ OTS02

Peer Cluster 2

Generate Passphrase

Manage Cluster Peers

3

Mediator ?

Not configured.

Configure

Storage VM Peers ⋮

- PEERED STORAGE VMS
- ✓ 3

- In the Peer Cluster dialog box, fill out the required information:
 - Enter the passphrase that was used to establish the peer cluster relationship on the destination ONTAP cluster.

- b. Select **Yes** to establish an encrypted relationship.
- c. Enter the intercluster LIF IP address(es) of the destination ONTAP cluster.
- d. Click **Initiate Cluster Peering** to finalize the process.

Peer Cluster ×

Local Remote

STORAGE VM PERMISSIONS

All storage VMs (incl... X

Storage VMs created in the future also will be given permissions.

PASSPHRASE ?

.....

It cannot be determined from the passphrase whether this relationship was encrypted. Is the relationship encrypted?

Yes No

To generate passphrase, [Launch Remote Cluster](#)

Intercluster Network Interfaces IP Addresses

172.30.15.42

172.30.14.28

Cancel

+ Add

Initiate Cluster Peering Cancel

4. Verify the status of the cluster peer relationship from the destination ONTAP cluster with the following command:

```
ONTAP-Dest::> cluster peer show
```


Establish SVM peering relationship

The next step is to set up an SVM relationship between the destination and source storage virtual machines that contain the volumes that will be in SnapMirror relationships.

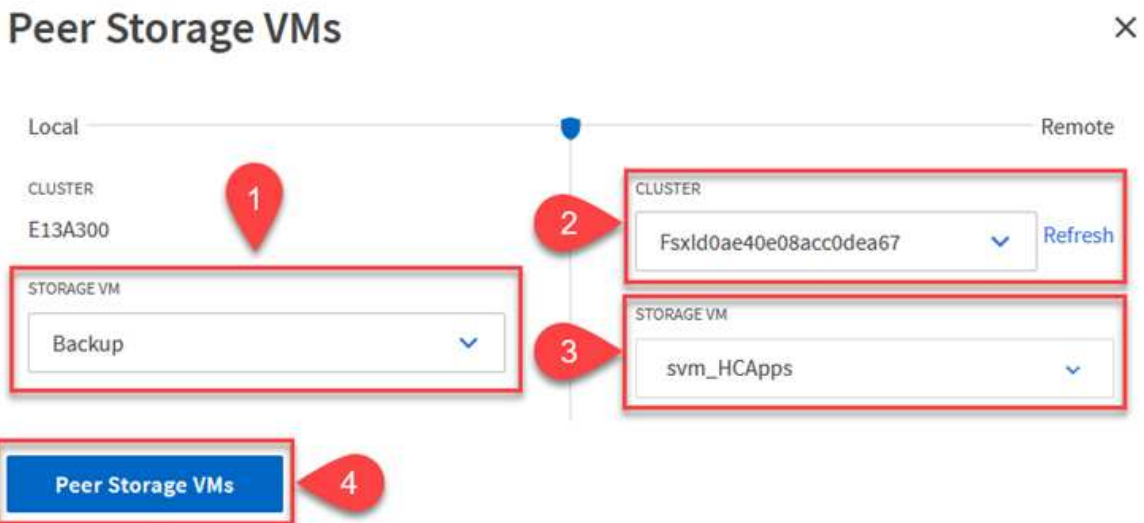
1. From the destination ONTAP cluster, use the following command from the CLI to create the SVM peer relationship:

```
ONTAP-Dest::> vserver peer create -vserver DestSVM -peer-vserver Backup -peer-cluster OnPremSourceSVM -applications snapmirror
```

2. From the source ONTAP cluster, accept the peering relationship with either ONTAP System Manager or the CLI.
3. From ONTAP System Manager, go to Protection > Overview and select Peer Storage VMs under Storage VM Peers.



4. In the Peer Storage VM's dialog box, fill out the required fields:
 - The source storage VM
 - The destination cluster
 - The destination storage VM



5. Click Peer Storage VMs to complete the SVM peering process.

Create a snapshot retention policy

SnapCenter manages retention schedules for backups that exist as snapshot copies on the primary storage system. This is established when creating a policy in SnapCenter. SnapCenter does not manage retention policies for backups that are retained on secondary storage systems. These policies are managed separately through a SnapMirror policy created on the secondary FSx cluster and associated with the destination volumes that are in a SnapMirror relationship with the source volume.

When creating a SnapCenter policy, you have the option to specify a secondary policy label that is added to the SnapMirror label of each snapshot generated when a SnapCenter backup is taken.



On the secondary storage, these labels are matched to policy rules associated with the destination volume for the purpose of enforcing retention of snapshots.

The following example shows a SnapMirror label that is present on all snapshots generated as part of a policy used for daily backups of our SQL Server database and log volumes.

Select secondary replication options

Update SnapMirror after creating a local Snapshot copy.

Update SnapVault after creating a local Snapshot copy.

Secondary policy label

sql-daily

Error retry count

For more information on creating SnapCenter policies for a SQL Server database, see the [SnapCenter documentation](#).

You must first create a SnapMirror policy with rules that dictate the number of snapshot copies to retain.

1. Create the SnapMirror Policy on the FSx cluster.

```
ONTAP-Dest::> snapmirror policy create -vserver DestSVM -policy  
PolicyName -type mirror-vault -restart always
```

2. Add rules to the policy with SnapMirror labels that match the secondary policy labels specified in the SnapCenter policies.

```
ONTAP-Dest::> snapmirror policy add-rule -vserver DestSVM -policy  
PolicyName -snapmirror-label SnapMirrorLabelName -keep  
#ofSnapshotsToRetain
```

The following script provides an example of a rule that could be added to a policy:

```
ONTAP-Dest::> snapmirror policy add-rule -vserver sql_svm_dest
-policy Async_SnapCenter_SQL -snapmirror-label sql-ondemand -keep 15
```



Create additional rules for each SnapMirror label and the number of snapshots to be retained (retention period).

Create destination volumes

To create a destination volume on ONTAP that will be the recipient of snapshot copies from our source volumes, run the following command on the destination ONTAP cluster:

```
ONTAP-Dest::> volume create -vserver DestSVM -volume DestVolName
-aggregate DestAggrName -size VolSize -type DP
```

Create the SnapMirror relationships between source and destination volumes

To create a SnapMirror relationship between a source and destination volume, run the following command on the destination ONTAP cluster:

```
ONTAP-Dest::> snapmirror create -source-path
OnPremSourceSVM:OnPremSourceVol -destination-path DestSVM:DestVol -type
XDP -policy PolicyName
```

Initialize the SnapMirror relationships

Initialize the SnapMirror relationship. This process initiates a new snapshot generated from the source volume and copies it to the destination volume.

To create a volume, run the following command on the destination ONTAP cluster:

```
ONTAP-Dest::> snapmirror initialize -destination-path DestSVM:DestVol
```

Configure the SnapCenter Plug-in for VMware vSphere

Once installed, the SnapCenter Plug-in for VMware vSphere can be accessed from the vCenter Server Appliance Management interface. SCV will manage backups for the NFS datastores mounted to the ESXi hosts and that contain the Windows and Linux VMs.

Review the [Data protection workflow](#) section of the SCV documentation for more information on the steps involved in configuring backups.

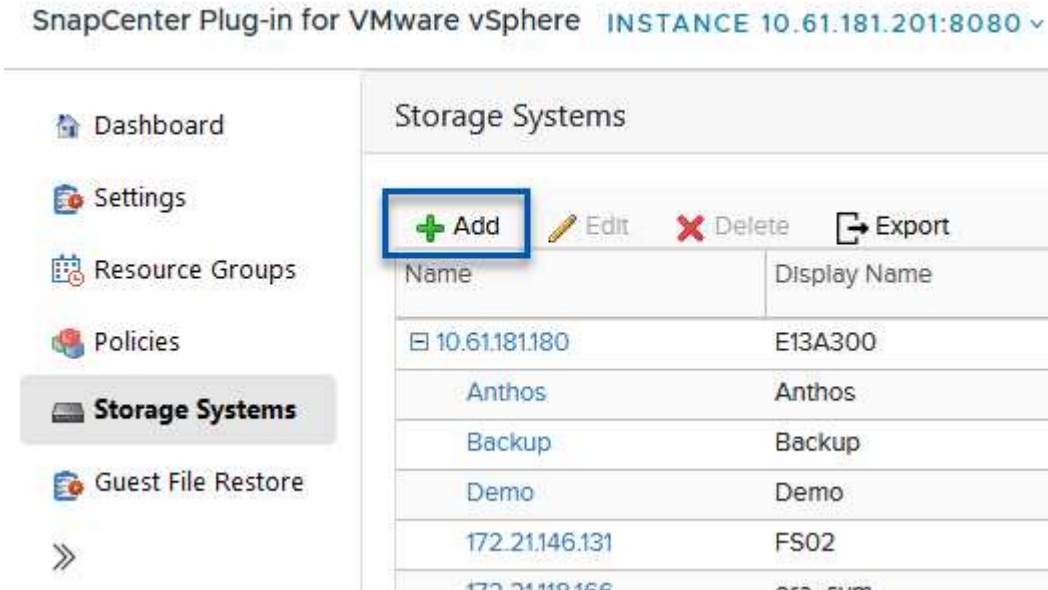
To configure backups of your virtual machines and datastores the following steps will need to be completed from the plug-in interface.

Discovery ONTAP storage systems

Discover the ONTAP storage clusters to be used for both primary and secondary backups.

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Storage Systems** in the left-hand menu and click on the **Add** button.

SnapCenter Plug-in for VMware vSphere **INSTANCE 10.61.181.201:8080** ▾



The screenshot shows the SnapCenter interface for the SnapCenter Plug-in for VMware vSphere. The left-hand menu is visible, with 'Storage Systems' selected. The main content area displays the 'Storage Systems' page, which includes a table of existing storage systems and a '+ Add' button highlighted with a blue box. The table has two columns: 'Name' and 'Display Name'. The existing storage systems listed are:

Name	Display Name
10.61.181.180	E13A300
Anthos	Anthos
Backup	Backup
Demo	Demo
172.21.146.131	FS02
172.21.146.155	FS03

2. Fill out the credentials and platform type for the primary ONTAP storage system and click on **Add**.

Add Storage System

Storage System	<input type="text" value="10.61.185.145"/>
Platform	<input type="text" value="All Flash FAS"/>
Authentication Method	<input checked="" type="radio"/> Credentials <input type="radio"/> Certificate
Username	<input type="text" value="admin"/>
Password	<input type="password" value="••••••••"/>
Protocol	<input type="text" value="HTTPS"/>
Port	<input type="text" value="443"/>
Timeout	<input type="text" value="60"/> <input type="text" value="Seconds"/>
<input type="checkbox"/> Preferred IP	<input type="text" value="Preferred IP"/>

Event Management System(EMS) & AutoSupport Setting

- Log Snapcenter server events to syslog
- Send AutoSupport Notification for failed operation to storage system

3. Repeat this procedure for the secondary ONTAP storage system.


Create SCV backup policies

Policies specify the retention period, frequency and replication options for the backups managed by SCV.

Review the [Create backup policies for VMs and datastores](#) section of the documentation for more information.

To create backup policies complete the following steps:

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Policies** in the left-hand menu and click on the **Create** button.



2. Specify a name for the policy, retention period, frequency and replication options, and snapshot label.

New Backup Policy

Name

Description

Retention ⓘ

Frequency

Replication

- Update SnapMirror after backup ⓘ
- Update SnapVault after backup ⓘ

Snapshot label

Advanced ▾

- VM consistency ⓘ
- Include datastores with independent disks

Scripts ⓘ



When creating a policy in the SnapCenter Plug-in you will see options for SnapMirror and SnapVault. If you choose SnapMirror, the retention schedule specified in the policy will be the same for both the primary and secondary snapshots. If you choose SnapVault, the retention schedule for the secondary snapshot will be based on a separate schedule implemented with the SnapMirror relationship. This is useful when you wish longer retention periods for secondary backups.



Snapshot labels are useful in that they can be used to enact policies with a specific retention period for the SnapVault copies replicated to the secondary ONTAP cluster. When SCV is used with BlueXP Backup and Restore, the Snapshot label field must either be blank or match the label specified in the BlueXP backup policy.

3. Repeat the procedure for each policy required. For example, separate policies for daily, weekly, and monthly backups.

Create resource groups

Resource groups contain the datastores and virtual machines to be included in a backup job, along with the associated policy and backup schedule.

Review the [Create resource groups](#) section of the documentation for more information.

To create resource groups complete the following steps.

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Resource Groups** in the left-hand menu and click on the **Create** button.



2. In the Create Resource Group wizard, enter a name and description for the group, as well as information required to receive notifications. Click on **Next**
3. On the next page select the datastores and virtual machines that wish to be included in the backup job and then click on **Next**.

Create Resource Group

1. General info & notification

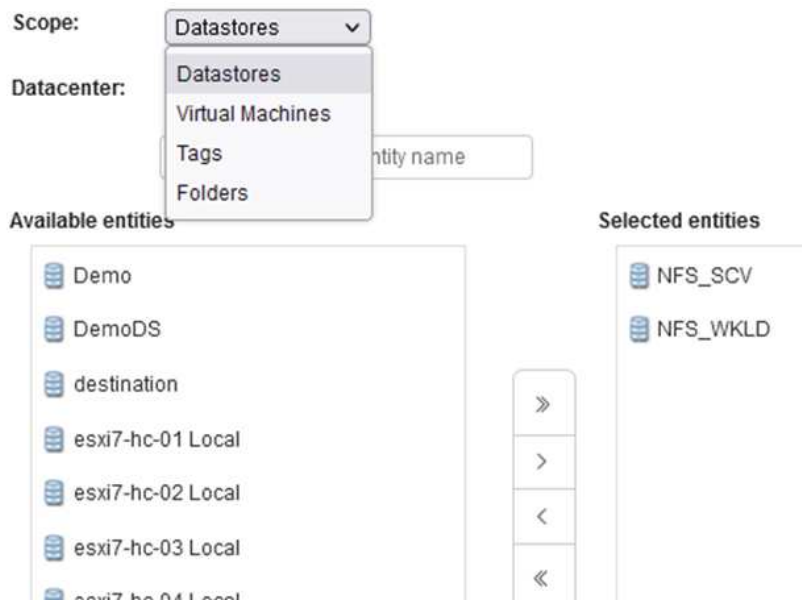
2. Resource

3. Spanning disks

4. Policies

5. Schedules

6. Summary





You have the option to select specific VMs or entire datastores. Regardless of which you choose, the entire volume (and datastore) is backed up since the backup is the result of taking a snapshot of the underlying volume. In most cases, it is easiest to choose the entire datastore. However, if you wish to limit the list of available VMs when restoring, you can choose only a subset of VMs for backup.

4. Choose options for spanning datastores for VMs with VMDKs that reside on multiple datastores and then click on **Next**.

Create Resource Group

1. General info & notification
 2. Resource
 3. Spanning disks
 4. Policies
 5. Schedules
 6. Summary

Always exclude all spanning datastores
 This means that only the datastores directly added to the resource group and the primary datastore of VMs directly added to the resource group will be backed up

Always include all spanning datastores
 All datastores spanned by all included VMs are included in this backup

Manually select the spanning datastores to be included
 You will need to modify the list every time new VMs are added

There are no spanned entities in the selected virtual entities list.



BlueXP backup and recovery does not currently support backing up VMs with VMDKs that span multiple datastores.

5. On the next page select the policies that will be associated with the resource group and click on **Next**.

Create Resource Group

1. General info & notification
 2. Resource
 3. Spanning disks
 4. Policies
 5. Schedules
 6. Summary

+ Create

<input type="checkbox"/>	Name	VM Consistent	Include independent di...	Schedule
<input checked="" type="checkbox"/>	Daily	No	No	Daily
<input type="checkbox"/>	FCD	No	Yes	On Demand Only
<input type="checkbox"/>	Monthly	No	No	Monthly
<input type="checkbox"/>	On Demand	No	No	On Demand Only
<input type="checkbox"/>	Weekly	No	No	Weekly



When backing up SCV managed snapshots to object storage using BlueXP backup and recovery, each resource group can only be associated with a single policy.

6. Select a schedule that will determine at what times the backups will run. Click on **Next**.

Create Resource Group

✓ 1. General info & notification

✓ 2. Resource

✓ 3. Spanning disks

✓ 4. Policies

✓ 5. Schedules

✓ 6. Summary

Daily



Type

Daily

Every

1

Day(s)

Starting

06/23/2023



At

07



00



PM



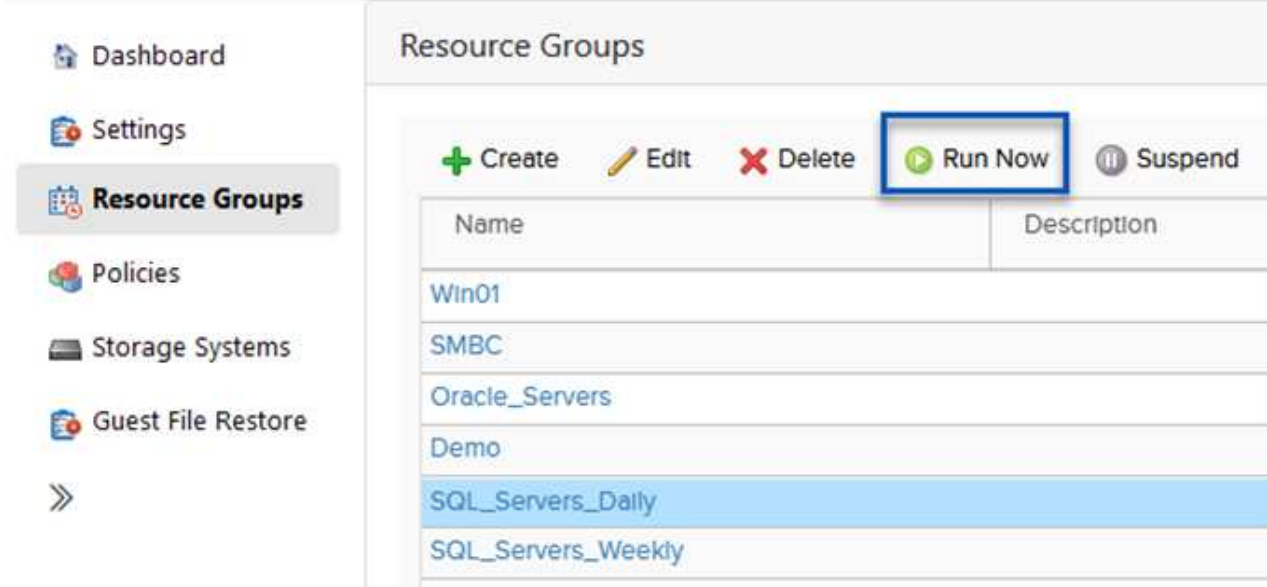
7. Finally, review the summary page and then on **Finish** to complete the resource group creation.

Run a backup job

In this final step, run a backup job and monitor its progress. At least one backup job must be successfully completed in SCV before resources can be discovered from BlueXP backup and recovery.

1. In the SnapCenter Plug-in for VMware vSphere navigate to **Resource Groups** in the left-hand menu.
2. To initiate a backup job, select the desired resource group and click the **Run Now** button.

SnapCenter Plug-in for VMware vSphere **INSTANCE 10.61.181.201:8080** ▾



The screenshot shows the SnapCenter interface for the SnapCenter Plug-in for VMware vSphere. The left-hand menu is visible, with 'Resource Groups' selected. The main content area displays a table of Resource Groups. The 'Run Now' button is highlighted with a blue box.

Name	Description
Win01	
SMBC	
Oracle_Servers	
Demo	
SQL_Servers_Daily	
SQL_Servers_Weekly	

3. To monitor the backup job, navigate to **Dashboard** on the left hand menu. Under **Recent Job Activities** click on the Job ID number to monitor the job progress.

Job Details : 2614 ↻ ✕

- ✔ Validate Retention Settings
- ✔ Quiescing Applications
- ✔ Retrieving Metadata
- ✔ Creating Snapshot copy
- ✔ Unquiescing Applications
- ✔ Registering Backup
- ✔ Backup Retention
- ✔ Clean Backup Cache
- ✔ Send EMS Messages
- ▶ (Job 2616)SnapVault Update

▶ Running, Start Time: 07/31/2023 07:24:40 PM.

CLOSE DOWNLOAD JOB LOGS

Configure Backups to Object Storage in BlueXP backup and recovery

For BlueXP to manage the data infrastructure effectively, it requires the prior installation of a Connector. The Connector executes the actions involved in discovering resources and managing data operations.

For more information on the BlueXP Connector refer to [Learn about Connectors](#) in the BlueXP documentation.


Once the connector is installed for the cloud provider being utilized, a graphic representation of the object storage will be viewable from the Canvas.

To configure BlueXP backup and recovery to backup data managed by SCV on-premises, complete the following steps:


Add working environments to the Canvas

The first step is to add the on-premises ONTAP storage systems to BlueXP

1. From the Canvas select **Add Working Environment** to begin.



2. Select **On-Premises** from the choice of locations and then click on the **Discover** button.



3. Fill out the credentials for the ONTAP storage system and click the **Discover** button to add the working environment.

ONTAP Cluster IP

10.61.181.180

User Name

admin

Password

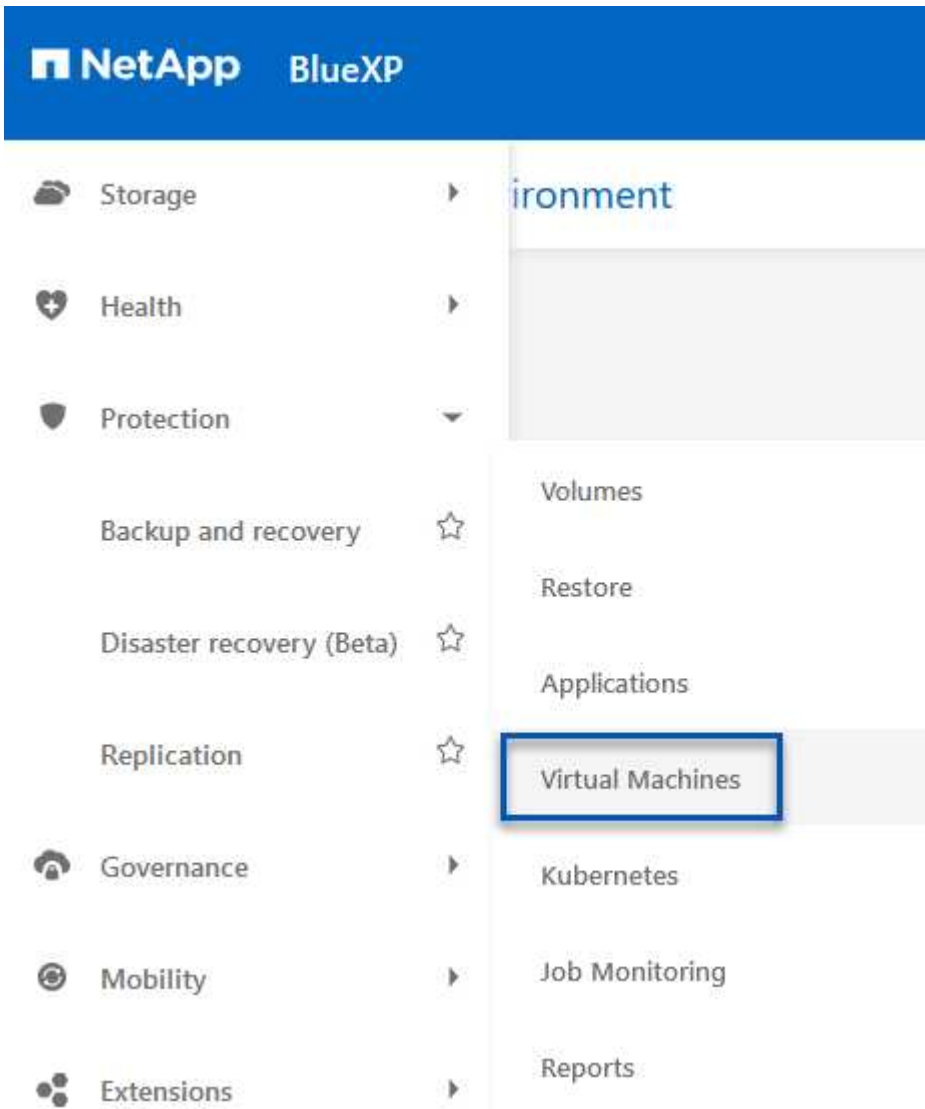
••••••••



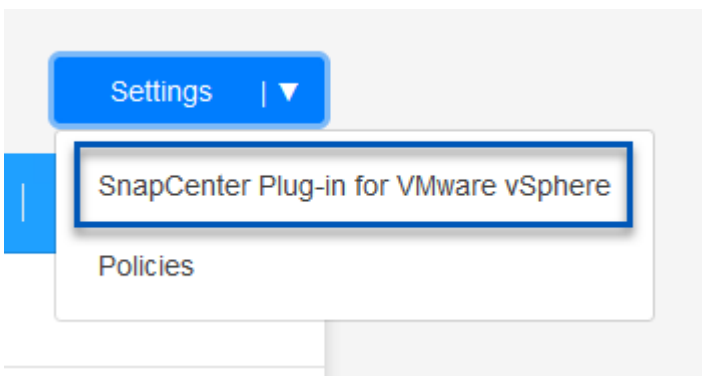
Discover on-premises SCV appliance and vCenter

To discover the on-premises datastore and virtual machine resources, add info for the SCV data broker and credentials for the vCenter management appliance.

1. From the BlueXP left-hand menu selection **Protection > Backup and recovery > Virtual Machines**



2. From the Virtual Machines main screen access the **Settings** drop down menu and select **SnapCenter Plug-in for VMware vSphere**.




3. Click on the **Register** button and then enter the IP address and port number for the SnapCenter Plug-in appliance and the username and password for the vCenter management appliance. Click on the **Register** button to begin the discovery process.

Register SnapCenter Plug-in for VMware vSphere


SnapCenter Plug-in for VMware vSphere	Username
<input type="text" value="10.61.181.201"/>	<input type="text" value="administrator@vsphere.local"/>
Port	Password
<input type="text" value="8144"/>	<input type="password" value="••••••••"/>

4. The progress of jobs can be monitored from the Job Monitoring tab.


Job Name: Discover Virtual Resources from SnapCenter Plugin for VMWare vSphere
Job Id: 559167ba-8876-45db-9131-b918a165d0a1




Other
Job Type



Jul 31 2023, 9:18:22 pm
Start Time



Jul 31 2023, 9:18:26 pm
End Time



Success
Job Status

Sub-Jobs(2) Collapse All ^

Job Name	Job ID	Start Time	End Time	Duration
Discover Virtual Resources from SnapCenter Plu...	559167ba-8876-45db-...	Jul 31 2023, 9:18:22 pm	Jul 31 2023, 9:18:26 pm	4 Seconds
Discovering Virtual Resources	99446761-f997-4c80-8...	Jul 31 2023, 9:18:22 pm	Jul 31 2023, 9:18:24 pm	2 Seconds
Registering Datastores	b7ab4195-1ee5-40ff-9a...	Jul 31 2023, 9:18:24 pm	Jul 31 2023, 9:18:26 pm	2 Seconds

5. Once discovery is complete you will be able to view the datastores and virtual machines across all discovered SCV appliances.

4 Working Environments

6 Datastores

14 Virtual Machines

Datastore Protection

4 Protected

2 Unprotected

6 Datastores

Filter By +

VM View

Settings


Datastore	Datastore Type	vCenter	Policy Name	Protection Status
NFS_SCV	NFS	vcsa7-hc.sddc.netapp.com		Unprotected
OTS_DS01	NFS	172.21.254.160	1 Year Daily LTR	Protected
SCV_WKLD	NFS	vcsa7-hc.sddc.netapp.com	1 Year Daily LTR	Protected
NFS_SQL	NFS	vcsa7-hc.sddc.netapp.com	1 Year Daily LTR	Protected
NFS_SQL2	NFS	vcsa7-hc.sddc.netapp.com	1 Year Daily LTR	Protected
SCV_DEMO	NFS	vcsa7-hc.sddc.netapp.com		Unprotected

Create BlueXP backup policies

In BlueXP backup and recovery for virtual machines, create policies to specify the retention period, backup source and the archival policy.

For more information on creating policies refer to [Create a policy to back up datastores](#).

1. From the BlueXP backup and recovery for virtual machines main page, access the **Settings** drop down menu and select **Policies**.



2. Click on **Create Policy** to access the **Create Policy for Hybrid Backup** window.
 - a. Add a name for the policy
 - b. Select the desired retention period
 - c. Select if backups will be sourced from the primary or secondary on-premises ONTAP storage system
 - d. Optionally, specify after what period of time backups will be tiered to archival storage for additional cost savings.

Create Policy for Hybrid Backup

Policy Details

Policy Name
12 week - daily backups

Retention ⓘ

Daily ^

Backups to retain: 84 SnapMirror Label: Daily

Weekly Setup Retention Weekly ∨

Monthly Setup Retention Monthly ∨

Backup Source

Primary

Secondary

Archival Policy ⓘ

Backups reside in standard storage for frequently accessed data. Optionally, you can tier backups to archival storage for further cost optimization.

Tier Backups to Archival

Archival After (Days)



The SnapMirror Label entered here is used to identify which backups to apply the policy too. The label name must match the label name in the corresponding on-premises SCV policy.

3. Click on **Create** to complete the policy creation.

Backup datastores to Amazon Web Services

The final step is to activate data protection for the individual datastores and virtual machines. The following steps outline how to activate backups to AWS.

For more information refer to [Back up datastores to Amazon Web Services](#).

1. From the BlueXP backup and recovery for virtual machines main page, access the settings drop down for the datastore to be backed up and select **Activate Backup**.

Datastore	Datastore Type	vCenter	Policy Name	Protection Status
NFS_SCV	NFS	vcsa7-hc.sddc.netapp.com		Unprotected
OTS_DS01	NFS	172.21.254.160	1 Year Daily LTR	Protected
SCV_WKLD	NFS	vcsa7-hc.sddc.netapp.com	1 Year Daily LTR	Protected

2. Assign the policy to be used for the data protection operation and click on **Next**.

1 Assign Policy 2 Add Working Environments 3 Select Provider 4 Configure Provider 5 Review

Assign Policy

21 Policies

	Policy Name	SnapMirror Label	Retention Count	Backup Source	Archival Policy
<input type="radio"/>	5 Year Daily LTR	daily	daily : 1830	Primary	Not Active
<input checked="" type="radio"/>	5 Year Daily LTR	daily	daily : 1830	Primary	Not Active
<input type="radio"/>	7 Year Weekly LTR	weekly	weekly : 370	Primary	Not Active

3. At the **Add Working Environments** page the datastore and working environment with a check mark should appear if the working environment has been previously discovered. If the working environment has not been previously discovered you can add it here. Click on **Next** to continue.


1 Assign Policy 2 Add Working Environments 3 Select Provider 4 Configure Provider 5 Review

Add Working Environments


Provide ONTAP cluster (working environment) details that you want Cloud Manager to discover. Working environment details will appear for all volumes that reside on the same cluster. You will need to enter multiple working environments when volumes reside on different clusters.

SVM	Volume	Working Environment	
EHC_NFS	NFS_SCV	<input checked="" type="checkbox"/> OnPremWorkingEnvironment-6MzE27u1	Edit

4. At the **Select Provider** page click on AWS and then click on the **Next** button to continue.



5. Fill out the provider specific credential information for AWS including the AWS access key and secret key, region, and archival tier to be used. Also, select the ONTAP IP space for the on-premises ONTAP storage system. Click on **Next**.



6. Finally, review the backup job details and click on the **Activate Backup** button to initiate data protection of the datastore.

Review

Policy	5 Year Daily LTR
SVM	EHC_NFS
Volumes	NFS_SCV
Working Environment	OnPremWorkingEnvironment-6MzE27u1
Backup Source	Primary
Cloud Service Provider	AWS
AWS Account	[REDACTED]
AWS Access Key	[REDACTED]
Region	US East (N. Virginia)
IP space	Default
Tier Backups to Archival	No

Previous

Activate Backup



At this point data transfer may not immediately begin. BlueXP backup and recovery scans for any outstanding snapshots every hour and then transfers them to object storage.

Restoring Virtual Machines in the case of data loss

Ensuring the safeguarding of your data is only one aspect of comprehensive data protection. Equally crucial is the ability to promptly restore data from any location in the event of data loss or a ransomware attack. This capability is vital for maintaining seamless business operations and meeting recovery point objectives.

NetApp offers a highly adaptable 3-2-1 strategy, providing customized control over retention schedules at the

primary, secondary, and object storage locations. This strategy provides the flexibility to tailor data protection approaches to specific needs.

This section provides an overview of the data restoration process from both the SnapCenter Plug-in for VMware vSphere and BlueXP backup and recovery for virtual machines.


Restoring Virtual Machines from SnapCenter Plug-in for VMware vSphere

For this solution virtual machines were restored to original and alternate locations. Not all aspects of SCV's data restoration capabilities will be covered in this solution. For in depth information on all that SCV has to offer refer to the [Restore VMs from backups](#) in the product documentation.


Restore virtual machines from SCV

Complete the following steps to restore a virtual machine restore from primary or secondary storage.

1. From the vCenter client navigate to **Inventory > Storage** and click on the datastore that contains the virtual machines you wish to restore.
2. From the **Configure** tab click on **Backups** to access the list of available backups.



3. Click on a backup to access the list of VMs and then select a VM to restore. Click on **Restore**.



4. From the Restore wizard select to restore the entire virtual machine or a specific VMDK. Select to install to the original location or alternate location, provide VM name after restore, and destination datastore. Click **Next**.

Restore ✕

✓ 1. Select scope

2. Select location

3. Summary

Restore scope Entire virtual machine ▾

Restart VM

Restore Location

Original Location
(This will restore the entire VM to the original Hypervisor with the original settings. Existing VM will be unregistered and replaced with this VM.)

Alternate Location
(This will create a new VM on selected vCenter and Hypervisor with the customized settings.)

Destination vCenter Server 10.61.181.210 ▾

Destination ESXi host esxi7-hc-04.sddc.netapp.com ▾

Network Management 181 ▾

VM name after restore SQL_SRV_08_restored

Select Datastore: NFS_SCV ▾

BACK
NEXT
FINISH
CANCEL

5. Choose to backup from the primary or secondary storage location.

Restore ✕

✓ 1. Select scope

2. Select location

3. Summary

Destination datastore	Locations
SCV_DEMO	(Primary) SCV:SCV_DEMO ▾
	(Primary) SCV:SCV_DEMO
	(Secondary) EHC_NFS:SCV_DEMO_dest

6. Finally, review a summary of the backup job and click on Finish to begin the restore process.

Restoring Virtual Machines from BlueXP backup and recovery for virtual machines


BlueXP backup and recovery for virtual machines allows restores of virtual machines to their original location. Restore functions are accessed through the BlueXP web console.

For more information refer to [Restore virtual machines data from the cloud](#).

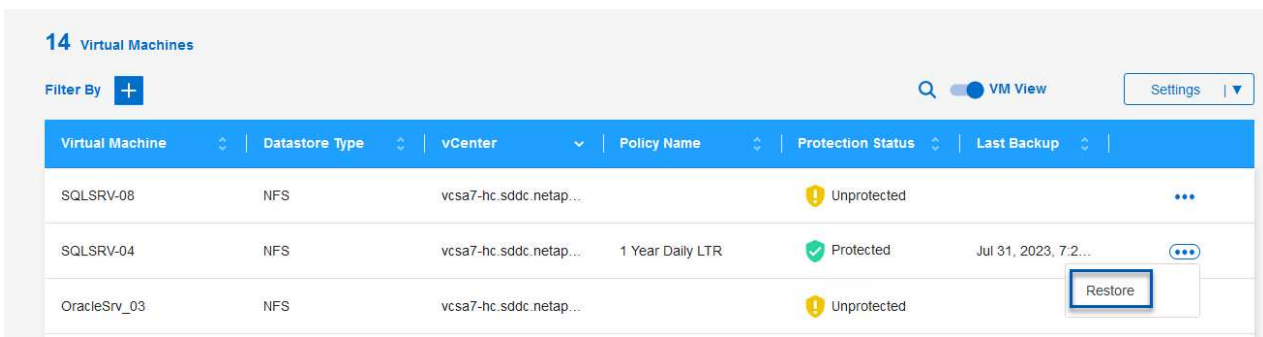
Restore virtual machines from BlueXP backup and recovery

To restore a virtual machine from BlueXP backup and recovery, complete the following steps.


1. Navigate to **Protection > Backup and recovery > Virtual Machines** and click on Virtual Machines to view the list of virtual machines available to be restored.



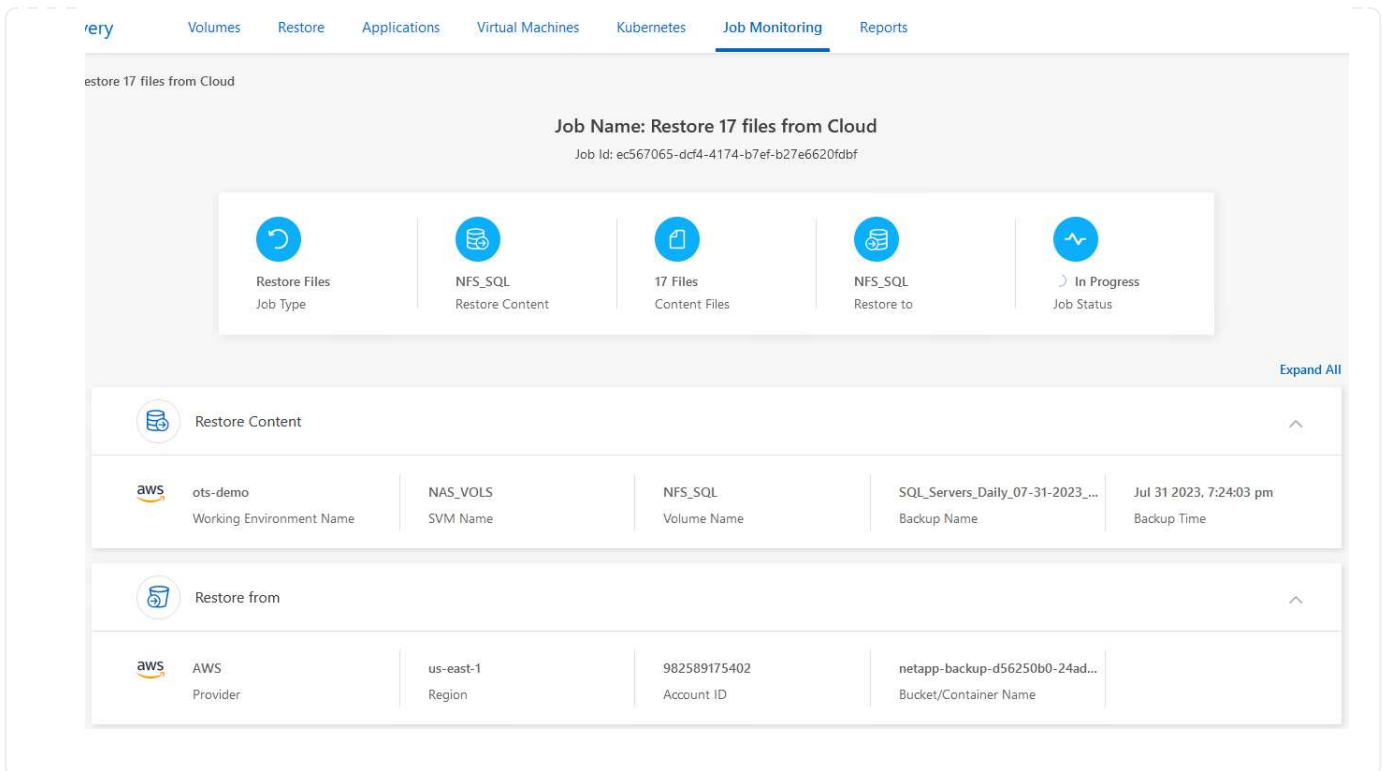
2. Access the settings drop down menu for the VM to be restored and select



3. Select the backup to restore from and click on **Next**.



4. Review a summary of the backup job and click on **Restore** to start the restore process.
5. Monitor the progress of the restore job from the **Job Monitoring** tab.



Conclusion

The 3-2-1 backup strategy, when implemented with SnapCenter Plug-in for VMware vSphere and BlueXP backup and recovery for virtual machines, offers a robust, reliable, and cost-effective solution for data protection. This strategy not only ensures data redundancy and accessibility but also provides the flexibility of restoring data from any location and from both on-premises ONTAP storage systems and cloud based object storage.

The use case presented in this documentation focuses on proven data protection technologies that highlight the integration between NetApp, VMware, and the leading cloud providers. The SnapCenter Plug-in for VMware vSphere provides seamless integration with VMware vSphere, allowing for efficient and centralized management of data protection operations. This integration streamlines the backup and recovery processes for virtual machines, enabling easy scheduling, monitoring, and flexible restore operations within the VMware ecosystem. BlueXP backup and recovery for virtual machines provides the one (1) in 3-2-1 by providing secure, air-gapped backups of virtual machine data to cloud based object storage. The intuitive interface and logical workflow provide a secure platform for long-term archival of critical data.

Additional Information

To learn more about the technologies presented in this solution refer to the following additional information.

- [SnapCenter Plug-in for VMware vSphere documentation](#)
- [BlueXP documentation](#)

DR using BlueXP DRaaS

Overview

Disaster Recovery is foremost in the minds of every VMware administrator. Because

VMware encapsulates entire servers into a series of files that make up the virtual machine; administrators take advantage of block storage-based techniques such as clones, snapshots and replicas to protect these VMs. ONTAP arrays offer built-in replication to transfer volume data, and therefore the virtual machines residing on the designated datastore LUNs, from one site to another. BlueXP DRaaS integrates with vSphere and automates the entire workflow for seamless failover and failback in the event of disaster. By combining storage replication with intelligent automation, administrators now have a manageable way to not only configure, automate, and test disaster recovery plans, but the means to easily run them in the case of a disaster.

Most time-consuming parts of a DR failover in a VMware vSphere environment is the execution of the steps necessary to inventory, register, reconfigure, and power up VMs at the DR site. An ideal solution has both a low RPO (as measured in minutes) and a low RTO (measured in minutes to hours). One factor that is often overlooked in a DR solution is the ability to test the DR solution efficiently on a periodic interval.

To architect a DR solution, keep the following factors in mind:

- The recovery time objective (RTO). The RTO is how quickly a business can recover from a disaster, or, more specifically, how long it takes to execute the recovery process to make business services available again.
- The recovery point objective (RPO). The RPO is how old the recovered data is after it has been made available, relative to the time that the disaster occurred.
- Scalability and adaptability. This factor includes the ability to grow storage resources incrementally as demand increases.

For more technical information on the available solutions, please see:

- [DR using BlueXP DRaaS for NFS Datastores](#)
- [DR using BlueXP DRaaS for VMFS Datastores](#)

DR using BlueXP DRaaS for NFS Datastores

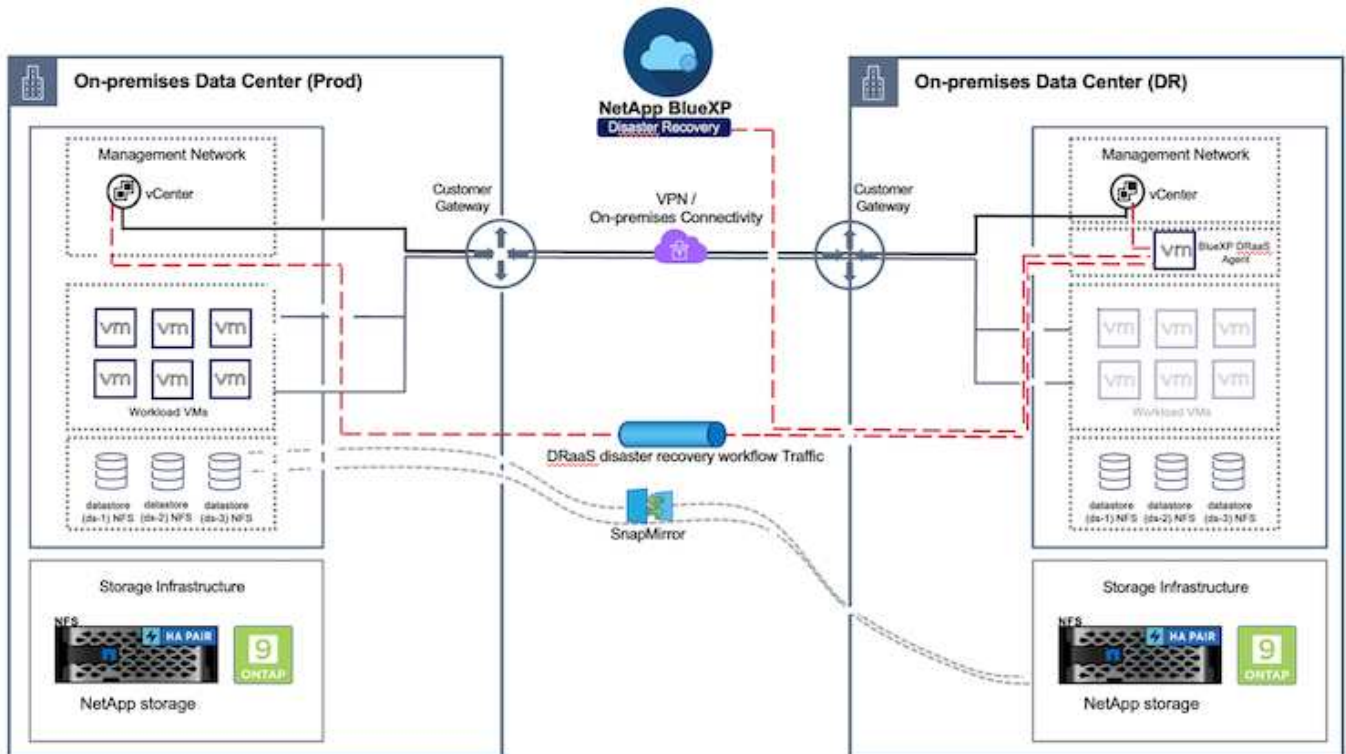
Implementing disaster recovery through block-level replication from the production site to the disaster recovery site is a resilient and cost-effective method for safeguarding workloads against site outages and data corruption events, such as ransomware attacks. Using NetApp SnapMirror replication, VMware workloads running on on-premises ONTAP systems with NFS datastore can be replicated to another ONTAP storage system located in a designated recovery datacenter where VMware is also deployed.

This section of the document describes the configuration of BlueXP DRaaS to set up disaster recovery for on-premises VMware VMs to another designated site. As part of this setup, the BlueXP account, BlueXP connector, the ONTAP arrays added within BlueXP workspace which is needed to enable communication from VMware vCenter to the ONTAP storage. In addition, this document details how to configure replication between sites and how to setup and test a recovery plan. The last section has instructions for performing a full site failover and how to failback when the primary site is recovered and brought online.

Utilizing the BlueXP disaster recovery service, integrated into the NetApp BlueXP console, companies can easily discover their on-premises VMware vCenters and ONTAP storage. Organizations can then create resource groupings, create a disaster recovery plan, associate it with resource groups, and test or execute failover and failback. SnapMirror provides storage-level block replication to keep the two sites up to date with

incremental changes, resulting in a Recovery Point Objective (RPO) of up to 5 minutes. Additionally, it is possible to simulate disaster recovery procedures without affecting production or incurring additional storage costs.


BlueXP disaster recovery leverages ONTAP's FlexClone technology to create a space-efficient copy of the NFS datastore from the last replicated Snapshot on the disaster recovery site. After completing the disaster recovery test, customers can easily delete the test environment without impacting actual replicated production resources. In case of an actual failover, the BlueXP disaster recovery service orchestrates all the necessary steps to automatically bring up the protected virtual machines on the designated disaster recovery site with just a few clicks. The service will also reverse the SnapMirror relationship to the primary site and replicate any changes from the secondary to the primary for a failback operation, when needed. All these capabilities come at a fraction of the cost compared to other well-known alternatives.



Getting started

To get started with BlueXP disaster recovery, use BlueXP console and then access the service.

1. Log in to BlueXP.
2. From the BlueXP left navigation, select Protection > Disaster recovery.
3. The BlueXP disaster recovery Dashboard appears.




Before configuring disaster recovery plan, ensure the following pre-requisites are met:

- BlueXP Connector is set up in NetApp BlueXP.
- BlueXP connector instance have connectivity to the source and destination vCenter and storage systems.
- NetApp Data ONTAP cluster to provide storage NFS datastores.
- On-premises NetApp storage systems hosting NFS datastores for VMware are added in BlueXP.
- DNS resolution should be in place when using DNS names. Otherwise, use IP addresses for the vCenter.
- SnapMirror replication is configured for the designated NFS based datastore volumes.
- Make sure that the environment has supported versions of vCenter Server and ESXi servers.

Once the connectivity is established between the source and destination sites, proceed with configuration steps, which should take couple of clicks and about 3 to 5 minutes.



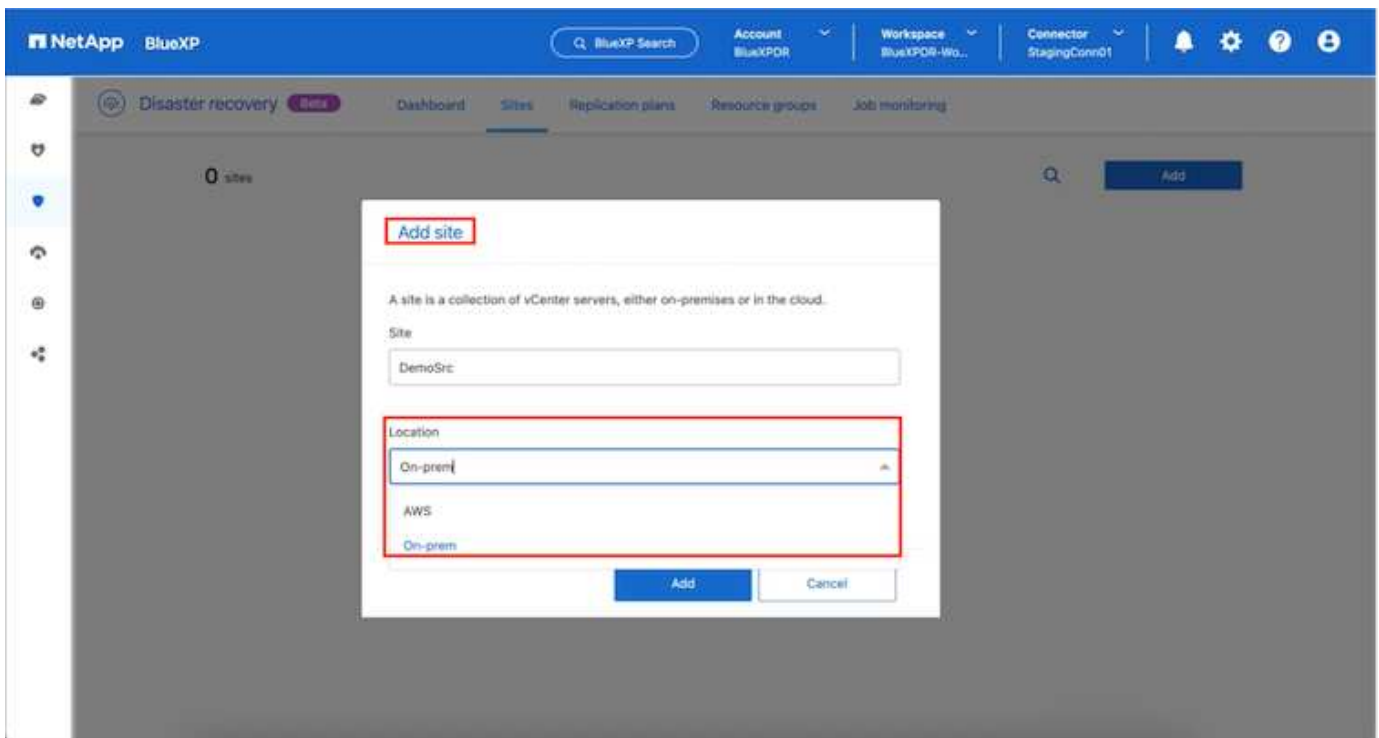
NetApp recommends deploying the BlueXP connector in the destination site or in a third site, so that the BlueXP connector can communicate through the network with source and destination resources.



BlueXP disaster recovery configuration

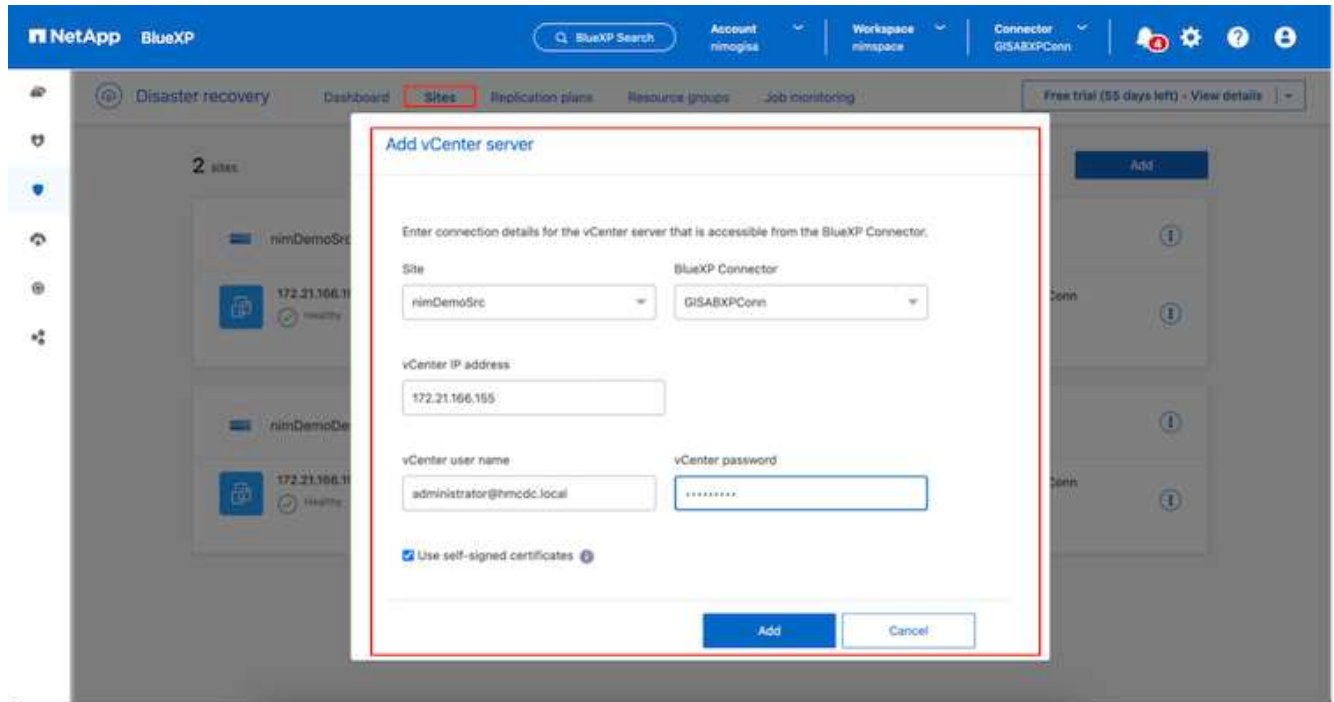
The first step in preparing for disaster recovery is to discover and add the on-premises vCenter and storage resources to BlueXP disaster recovery.

Open BlueXP console and select **Protection > Disaster Recovery** from left navigation. Select **Discover vCenter servers** or use top menu, Select **Sites > Add > Add vCenter**.




Add the following platforms:

- **Source.** On-premises vCenter.



- **Destination.** VMC SDDC vCenter.



Once the vCenters are added, automated discovery is triggered.

Configuring Storage replication between source site array and destination site array

SnapMirror provides data replication in a NetApp environment. Built on NetApp Snapshot® technology, SnapMirror replication is extremely efficient because it replicates only the blocks that have been changed or added since the previous update. SnapMirror is easily configured by using either NetApp OnCommand®

System Manager or the ONTAP CLI. BlueXP DRaaS also creates the SnapMirror relationship provided cluster and SVM peering is configured beforehand.

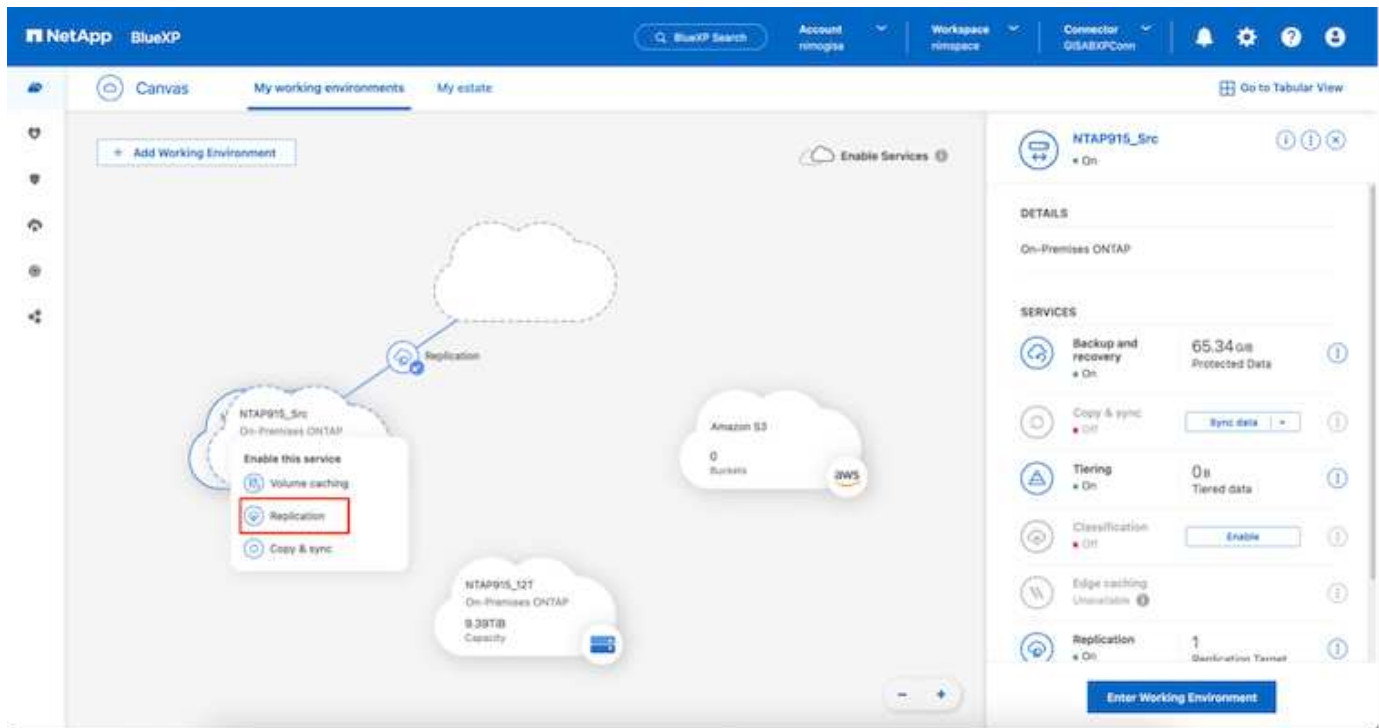
For cases in which the primary storage is not completely lost, SnapMirror provides an efficient means of resynchronizing the primary and DR sites. SnapMirror can resynchronize the two sites, transferring only changed or new data back to the primary site from the DR site by simply reversing the SnapMirror relationships. This means replication plans in BlueXP DRaaS can be resynchronized in either direction after a failover without recopying the entire volume. If a relationship is resynchronized in the reverse direction, only new data that was written since the last successful synchronization of the Snapshot copy is sent back to the destination.



If SnapMirror relationship is already configured for the volume via CLI or System Manager, BlueXP DRaaS picks up the relationship and continues with the rest of the workflow operations.


How to set it up for VMware Disaster Recovery

The process to create SnapMirror replication remains the same for any given application. The process can be manual or automated. The easiest way is to leverage BlueXP to configure SnapMirror replication by using simple drag & drop of the source ONTAP system in the environment onto the destination to trigger the wizard that guides through the rest of the process.



BlueXP DRaaS can also automate the same provided the following two criteria's are met:

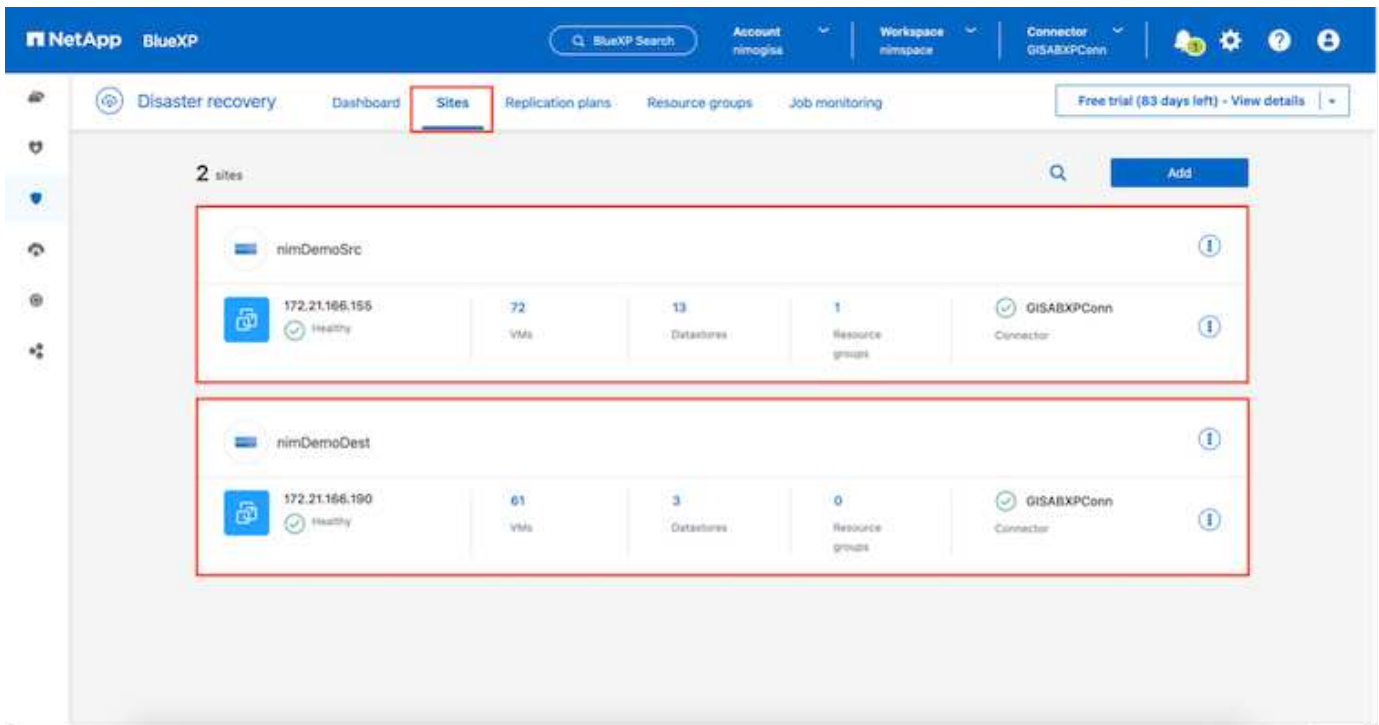
- Source and destination clusters have a peer relationship.
- Source SVM and destination SVM have a peer relationship.



If SnapMirror relationship is already configured for the volume via CLI, BlueXP DRaaS picks up the relationship and continues with the rest of the workflow operations.



What can BlueXP disaster recovery do for you?

After the source and destination sites are added, BlueXP disaster recovery performs automatic deep discovery and displays the VMs along with associated metadata. BlueXP disaster recovery also automatically detects the networks and port groups used by the VMs and populates them.




After the sites have been added, VMs can be grouped into resource groups. BlueXP disaster recovery

resource groups allow you to group a set of dependent VMs into logical groups that contain their boot orders and boot delays that can be executed upon recovery. To start creating resource groups, navigate to **Resource Groups** and click **Create New Resource Group**.

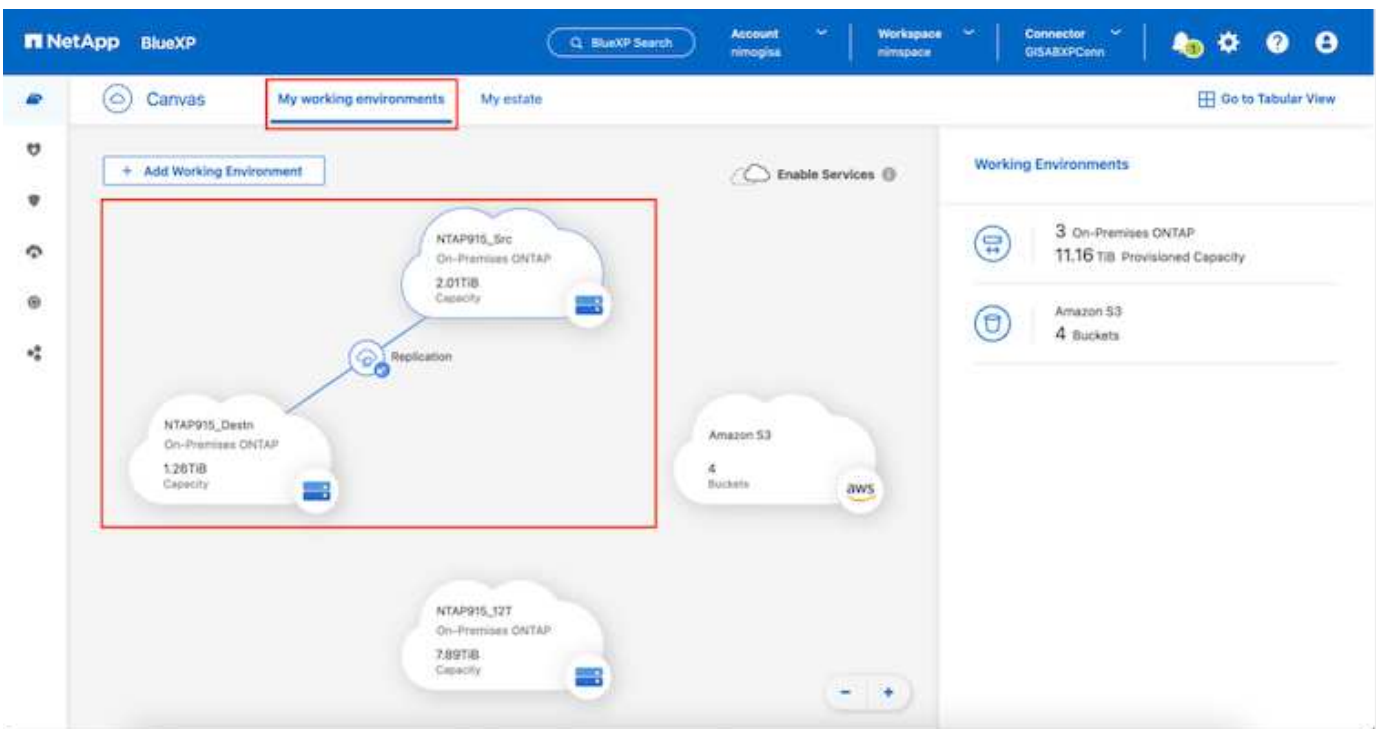


The resource group can also be created while creating a replication plan.

The boot order of the VMs can be defined or modified during the creation of resource groups by using simple drag and drop mechanism.



Once the resource groups are created, the next step is to create the execution blueprint or a plan to recover virtual machines and applications in the event of a disaster. As mentioned in the prerequisites, SnapMirror replication can be configured beforehand or DRaaS can configure it using the RPO and retention count specified during creation of the replication plan.



NetApp BlueXP

Account nimogisa | Workspace nimspace | Connector GISABXPCann

Replication

Volume Relationships (8)

Health Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Successful Transfer
	NTAP915_Src	NTAP915_Destn				30.3 MB
✓	Demo_TPS_DS01 NTAP915_Src	Demo_TPS_DS01_Copy NTAP915_Destn	13 seconds	idle	snapmirrored	Aug 5, 2024, 6:15 388.63 MiB
✓	Src_250_Vol01 NTAP915_Src	Src_250_Vol01_Copy NTAP915_Destn	4 seconds	idle	snapmirrored	Aug 16, 2024, 12: 79.23 MiB
✓	Src_NFS_DS03 NTAP915_Src	Src_NFS_DS03_CP NTAP915_Destn	12 seconds	idle	snapmirrored	Aug 16, 2024, 12: 24.64 MiB
✓	Src_NFS_DS04 NTAP915_Src	Src_NFS_DS04_CP NTAP915_Destn	3 seconds	idle	snapmirrored	Aug 16, 2024, 12: 47.38 MiB
✓	Src_JSCSI_DS04 NTAP915_Src	Src_JSCSI_DS04_copy NTAP915_Destn	4 seconds	idle	snapmirrored	Aug 16, 2024, 12: 108.87 MiB
✓	nimpra NTAP915_Src	nimpra_dest NTAP915_Destn	2 seconds	idle	snapmirrored	Aug 16, 2024, 12: 3.48 KiB

Configure the replication plan by selecting the source and destination vCenter platforms from the drop down and pick the resource groups to be included in the plan, along with the grouping of how applications should be restored and powered on and mapping of clusters and networks. To define the recovery plan, navigate to the **Replication Plan** tab and click **Add Plan**.

First, select the source vCenter and then select the destination vCenter.

NetApp BlueXP

Account nimogisa | Workspace nimspace | Connector GISABXPCann

Add replication plan

1 vCenter servers | 2 Applications | 3 Resource mapping | 4 Recurrence | 5 Review

Replication plan name
DemoNFSDS03RP

Select a source vCenter where your data exists, to replicate to the selected target vCenter.

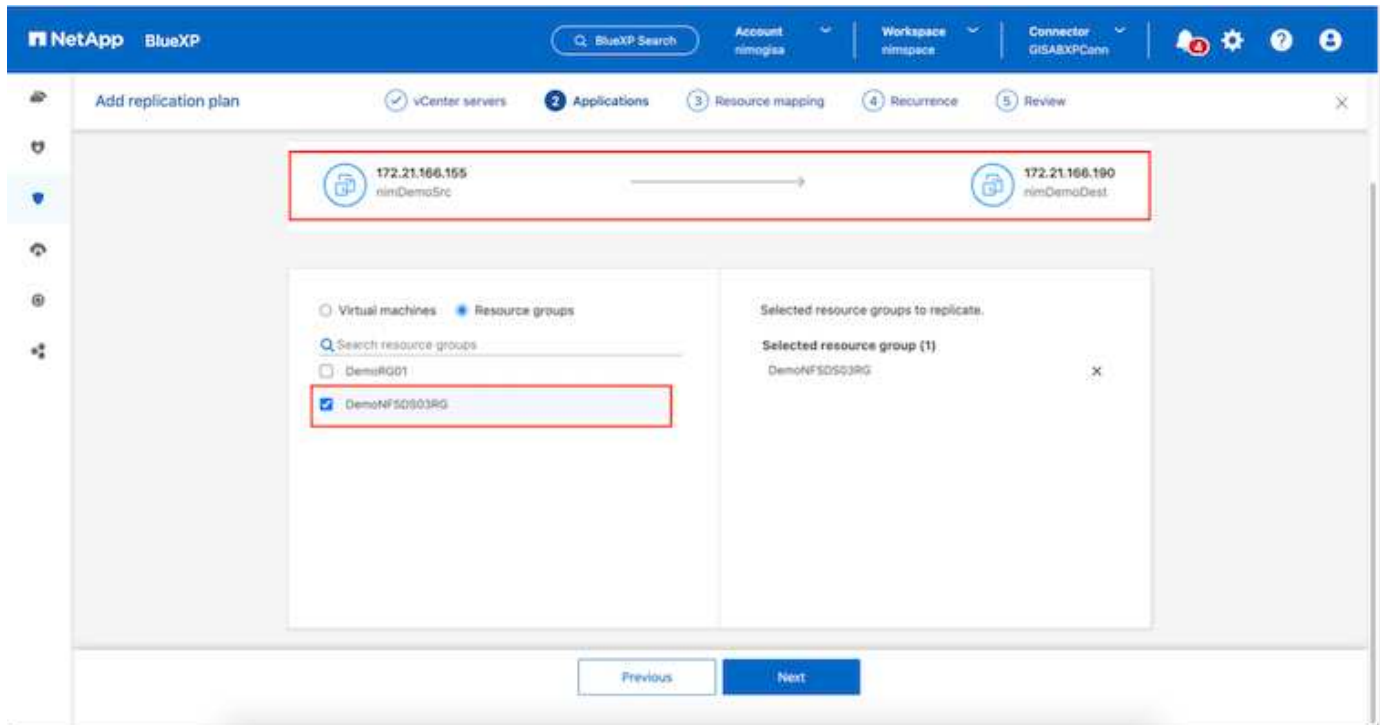
Source vCenter: 172.21.166.155

Target vCenter: 172.21.166.190

Cancel | Next

The next step is to select existing resource groups. If no resource groups created, then the wizard helps to group the required virtual machines (basically create functional resource groups) based on the recovery objectives. This also helps define the operation sequence of how application virtual machines should be

restored.




Resource group allows to set boot order using the drag and drop functionality. It can be used to easily modify the order in which the VMs would be powered on during the recovery process.

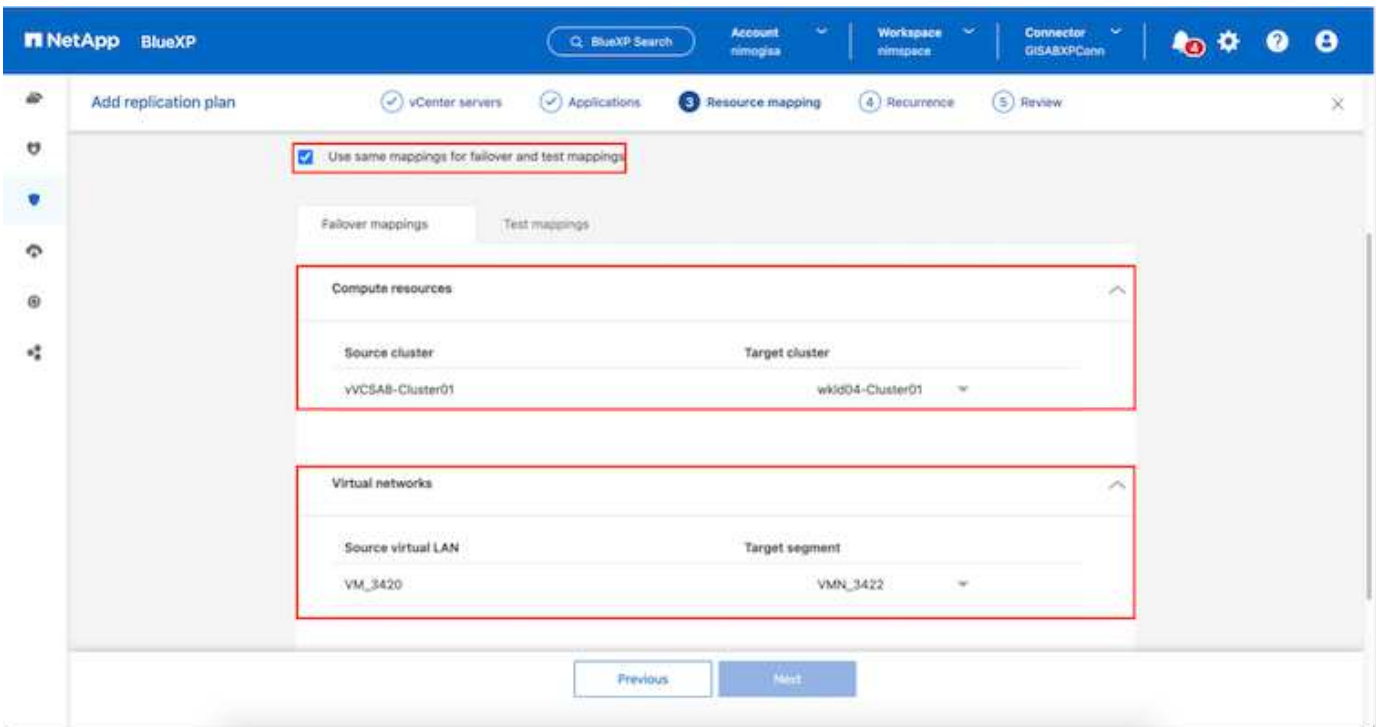


Each virtual machine within a resource group is started in sequence based on the order. Two resource groups are started in parallel.


The below screenshot shows the option to filter virtual machines or specific datastores based on organizational requirements if resource groups are not created beforehand.



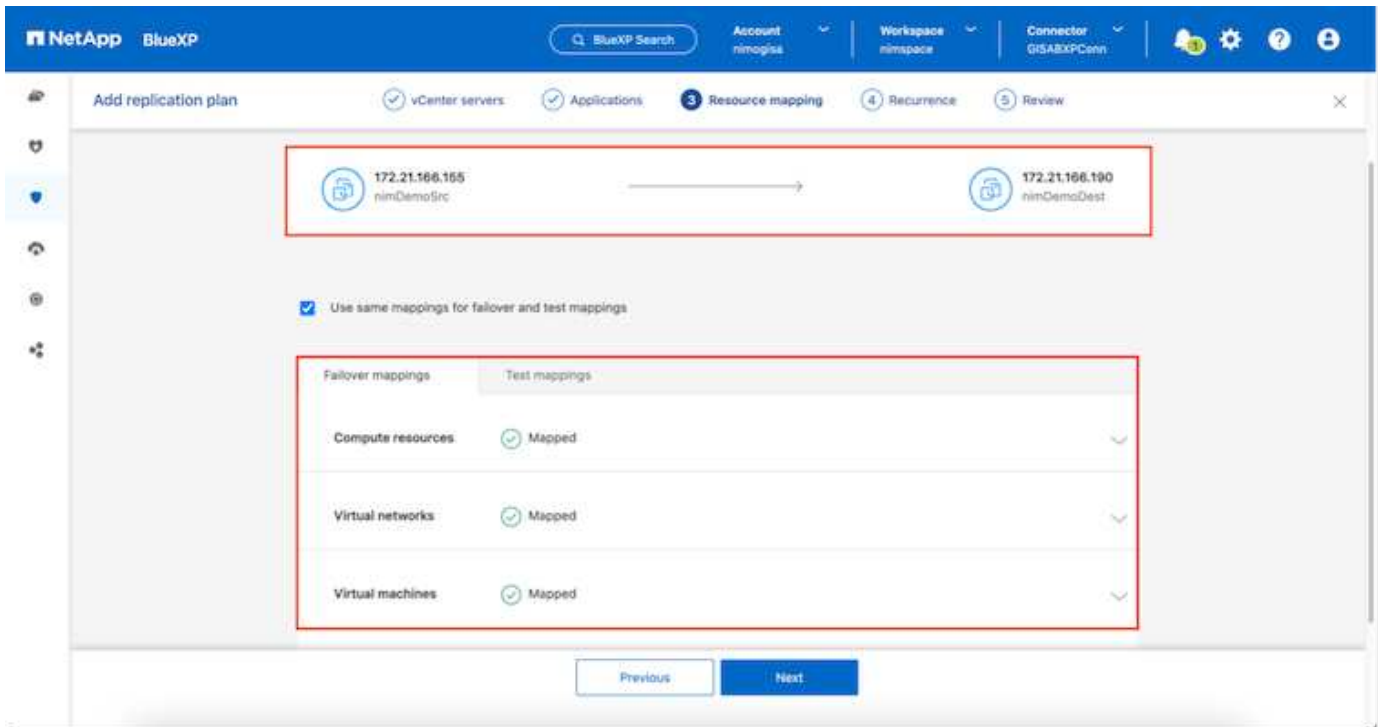
Once the resource groups are selected, create the failover mappings. In this step, specify how the resources from the source environment maps to the destination. This includes compute resources, virtual networks, IP customization, pre- and post-scripts, boot delays, application consistency and so on. For detailed information, refer to [Create a replication plan](#).



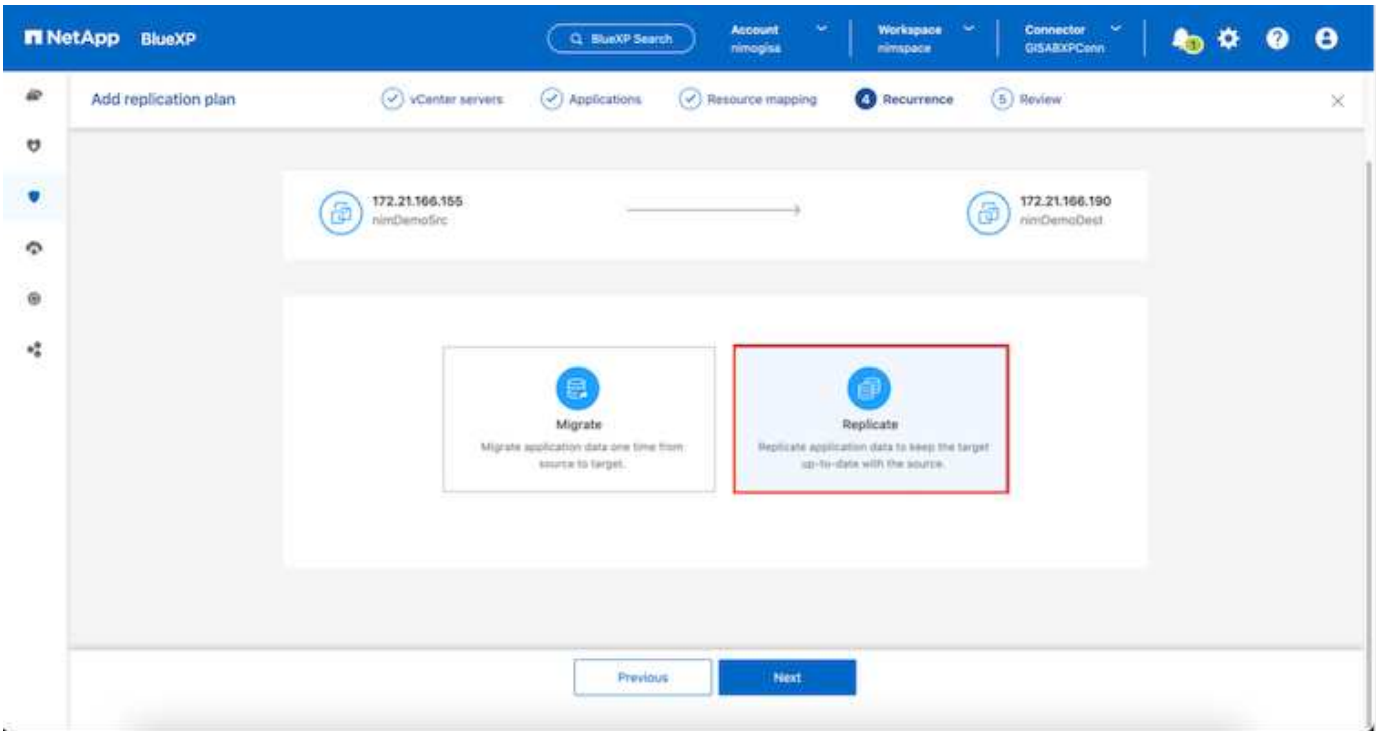
By default, same mapping parameters are used for both test and failover operations. To set different mappings for test environment, select the Test mapping option after unchecking the checkbox as shown below:



Once the resource mapping is complete, click Next.



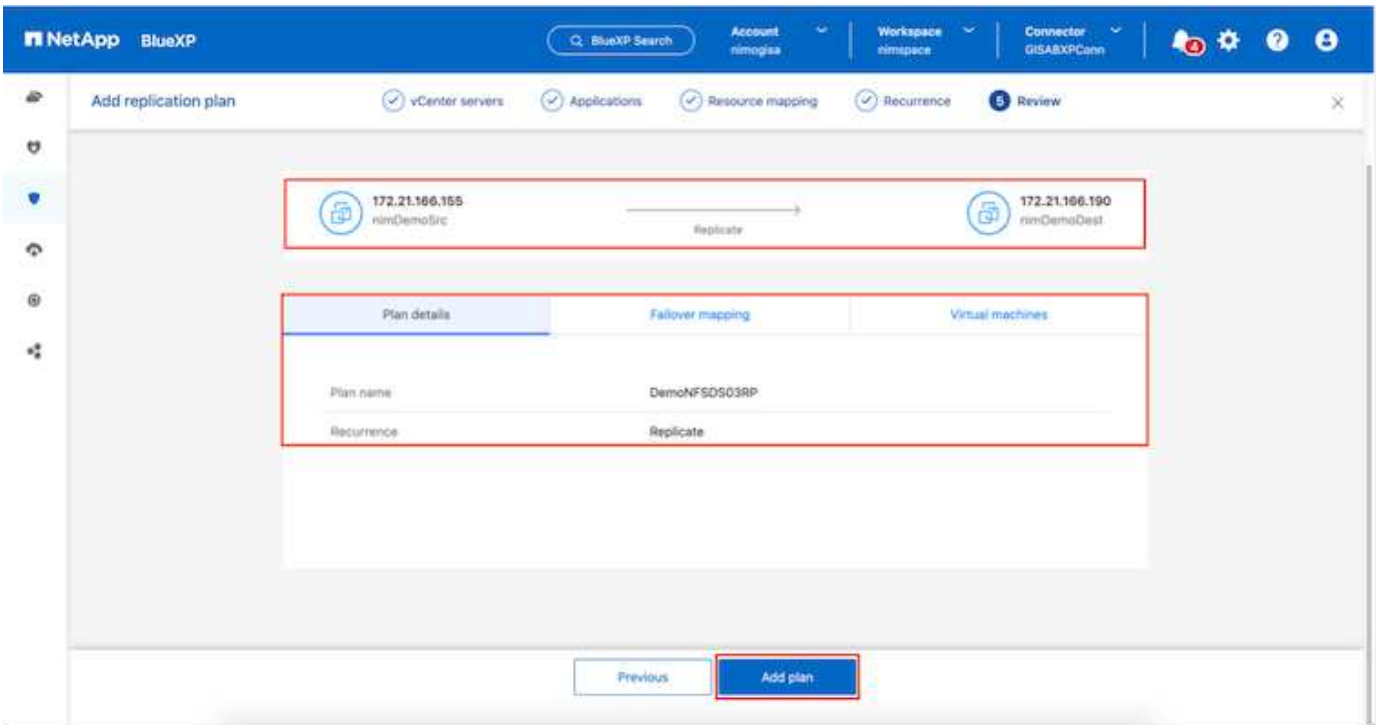
Select the recurrence type. In simple words, select Migrate (one time migration using failover) or recurring continuous replication option. In this walkthrough, Replicate option is selected.




Once done, review the created mappings and then click on **Add plan**.



VMs from different volumes and SVMs can be included in a replication plan. Depending on the VM placement (be it on same volume or separate volume within the same SVM, separate volumes on different SVMs), the BlueXP disaster recovery creates a Consistency Group Snapshot.






BlueXP DRaaS consists of the following workflows:

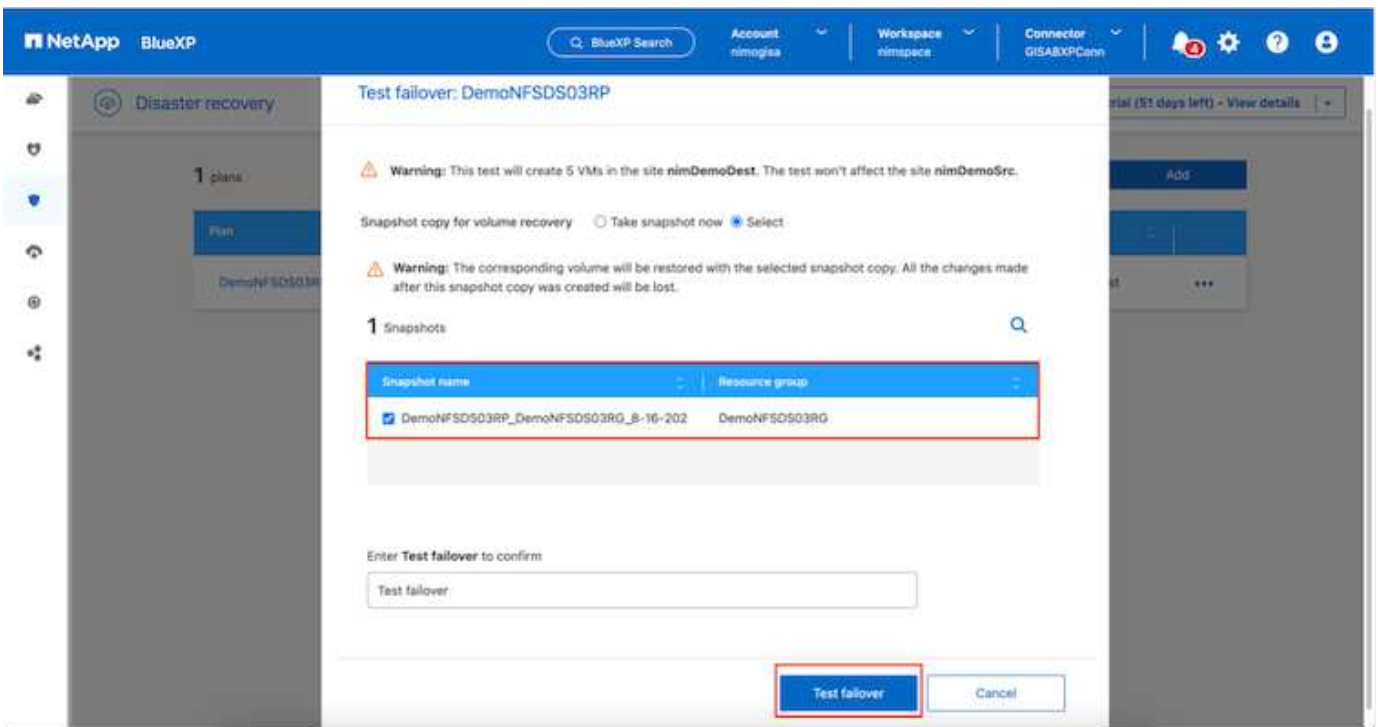
- Test failover (including periodic automated simulations)
- Cleanup failover test
- Failover
- Failback

Test failover

Test failover in BlueXP DRaaS is an operational procedure that allows VMware administrators to fully validate their recovery plans without disrupting their production environments.



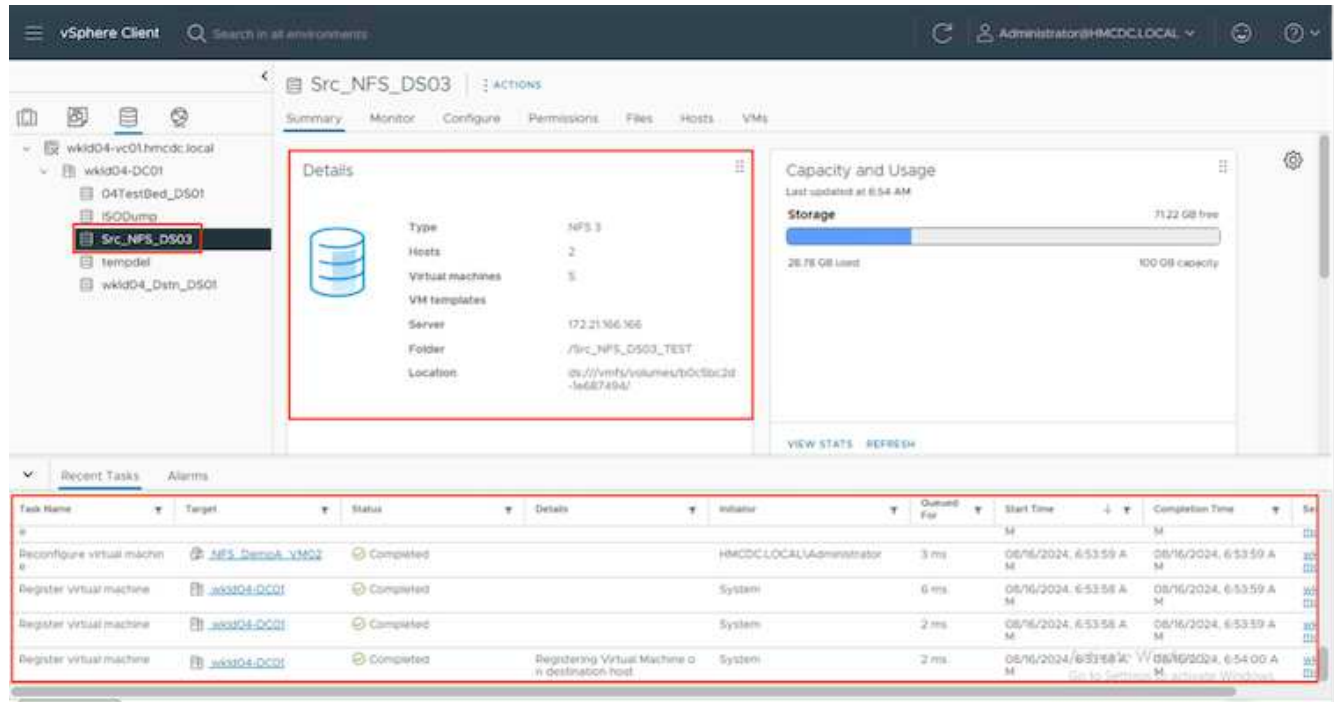
BlueXP DRaaS incorporates the ability to select the snapshot as an optional capability in the test failover operation. This capability allows the VMware administrator to verify that any changes that were recently made in the environment are replicated to the destination site and thus are present during the test. Such changes include patches to the VM guest operating system



When the VMware administrator runs a test failover operation, BlueXP DRaaS automates the following tasks:


- Triggering SnapMirror relationships to update storage at the destination site with any recent changes that were made at the production site.
- Creating NetApp FlexClone volumes of the FlexVol volumes on the DR storage array.

- Connecting the NFS datastores in the FlexClone volumes to the ESXi hosts at the DR site.
- Connecting the VM network adapters to the test network specified during the mapping.
- Reconfiguring the VM guest operating system network settings as defined for the network at the DR site.
- Executing any custom commands that have been stored in the replication plan.
- Powering on the VMs in the order that is defined in the replication plan.



Cleanup failover test Operation

The cleanup failover test operation occurs after the replication plan test has been completed and the VMware administrator responds to the cleanup prompt.




This action will reset the virtual machines (VMs) and the status of the replication plan to the ready state.

When the VMware administrator performs a recovery operation, BlueXP DRaaS completes the following process:

1. It powers off each recovered VM in the FlexClone copy that was used for testing.
2. It deletes the FlexClone volume that was used to present the recovered VMs during the test.


Planned Migration and Fail over

BlueXP DRaaS has two methods for performing a real failover: planned migration and fail over. The first method, planned migration, incorporates VM shutdown and storage replication synchronization into the process to recover or effectively move the VMs to the destination site. Planned migration requires access to the source site. The second method, failover, is an planned/unplanned failover in which the VMs are recovered at the destination site from the last storage replication interval that was able to complete. Depending on the RPO that was designed into the solution, some amount of data loss can be expected in the DR scenario.



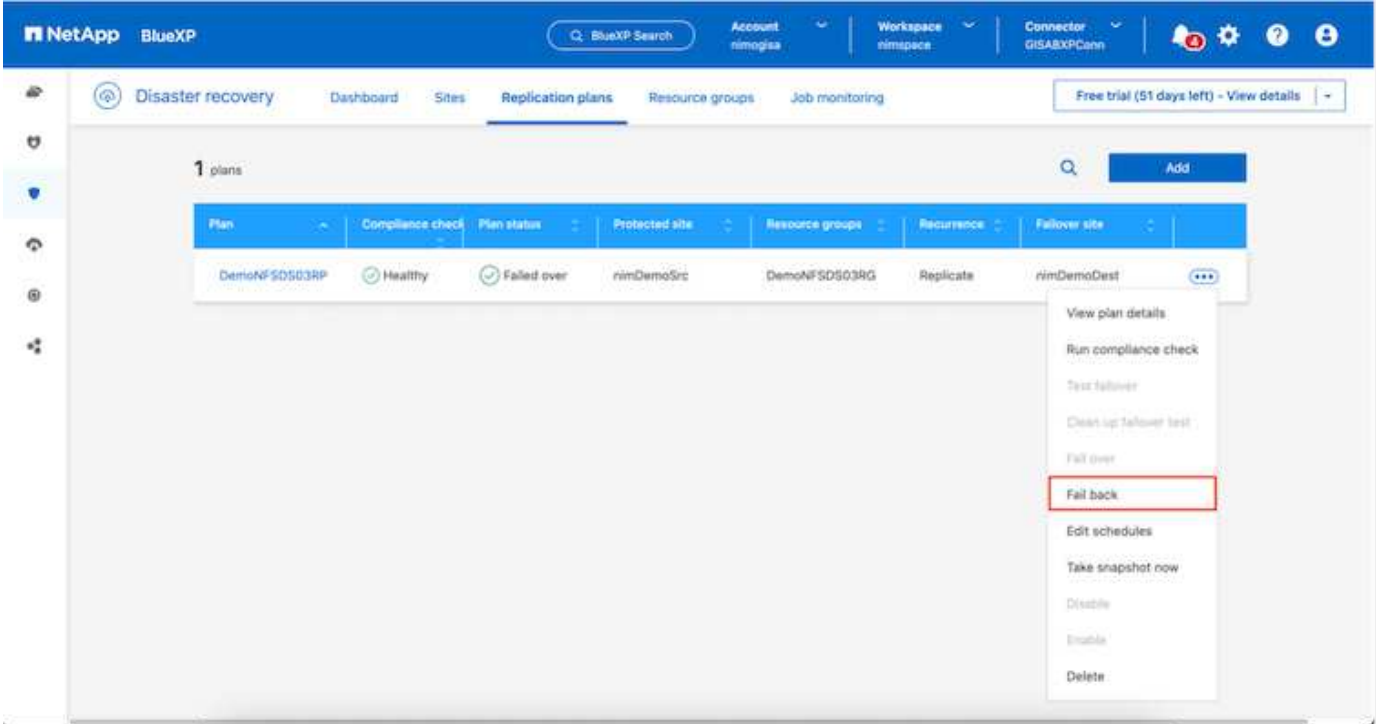
When the VMware administrator performs a failover operation, BlueXP DRaaS automates the following tasks:

- Break and fail over the NetApp SnapMirror relationships.
- Connect the replicated NFS datastores to the ESXi hosts at the DR site.
- Connect the VM network adapters to the appropriate destination site network.
- Reconfigure the VM guest operating system network settings as defined for the network at the destination site.
- Execute any custom commands (if any) that have been stored in the replication plan.
- Power on the VMs in the order that was defined in the replication plan.



Failback

A failback is an optional procedure that restores the original configuration of the source and destination sites after a recovery.



VMware administrators can configure and run a failback procedure when they are ready to restore services to the original source site.

NOTE: BlueXP DRaaS replicates (resyncs) any changes back to the original source virtual machine before reversing the replication direction. This process starts from a relationship that has completed failing over to a

target and involves the following steps:

- Power off and unregister the virtual machines and volumes on the destination site are unmounted.
- Break the SnapMirror relationship on the original source is broken to make it read/write.
- Resynchronize the SnapMirror relationship to reverse the replication.
- Mount the volume on the source, power on and register the source virtual machines.

For more details about accessing and configuring BlueXP DRaaS, see the [Learn about BlueXP Disaster Recovery for VMware](#).

Monitoring and Dashboard


From BlueXP or the ONTAP CLI, you can monitor the replication health status for the appropriate datastore volumes, and the status of a failover or test failover can be tracked via Job Monitoring.

ID	Status	Workload	Name	Start time	End time	
d923e507-b2c2-401	In pro...	Backup	Backup job for Replication Plan:DemoNF...	08/16/2024, 04:5...	-	Cancel job?
3549cc9c-aa4e-45e	Succe...	Backup	Initialize Backup of DemoNFSDS03RP for...	08/16/2024, 04:5...	08/16/2024, 04:5...	
5cb01bcc-9ea8-4af1	Succe...	Backup	Backup job for Replication Plan:DemoNF...	08/16/2024, 04:4...	08/16/2024, 04:5...	
a2f225d9-b7be-4c2f	Succe...	Backup	Initialize Backup of DemoNFSDS03RP for...	08/16/2024, 04:4...	08/16/2024, 04:4...	
2f8b44d4-4be2-46t	Succe...	Compliance	Compliance check for Replication Plan: D...	08/16/2024, 04:4...	08/16/2024, 04:4...	
398bc6a3-ata8-45d	Succe...	Compliance	Initialize Compliance of DemoNFSDS03R...	08/16/2024, 04:4...	08/16/2024, 04:4...	
97fdbed8-6f77-459f	Succe...	Backup	Backup job for Replication Plan:DemoNF...	08/16/2024, 04:4...	08/16/2024, 04:4...	
bffc016e-ca3a-4D9d	Succe...	Backup	Initialize Backup of DemoNFSDS03RP for...	08/16/2024, 04:4...	08/16/2024, 04:4...	
cde759a8-ebef-498b	Succe...	Backup	Backup job for Replication Plan:DemoNF...	08/16/2024, 04:3...	08/16/2024, 04:4...	
a414daba-9830-4c5	Succe...	Backup	Initialize Backup of DemoNFSDS03RP for...	08/16/2024, 04:3...	08/16/2024, 04:3...	



If a job is currently in progress or queued, and you wish to stop it, there is an option to cancel it.

With the BlueXP disaster recovery dashboard, confidently evaluate the status of disaster recovery sites and replication plans. This enables administrators to swiftly identify healthy, disconnected, or degraded sites and plans.



This provides a powerful solution to handle a tailored and customized disaster recovery plan. Failover can be done as planned failover or failover with a click of a button when disaster occurs and decision is made to activate the DR site.

To learn more about this process, feel free to follow the detailed walkthrough video or use the [solution simulator](#).

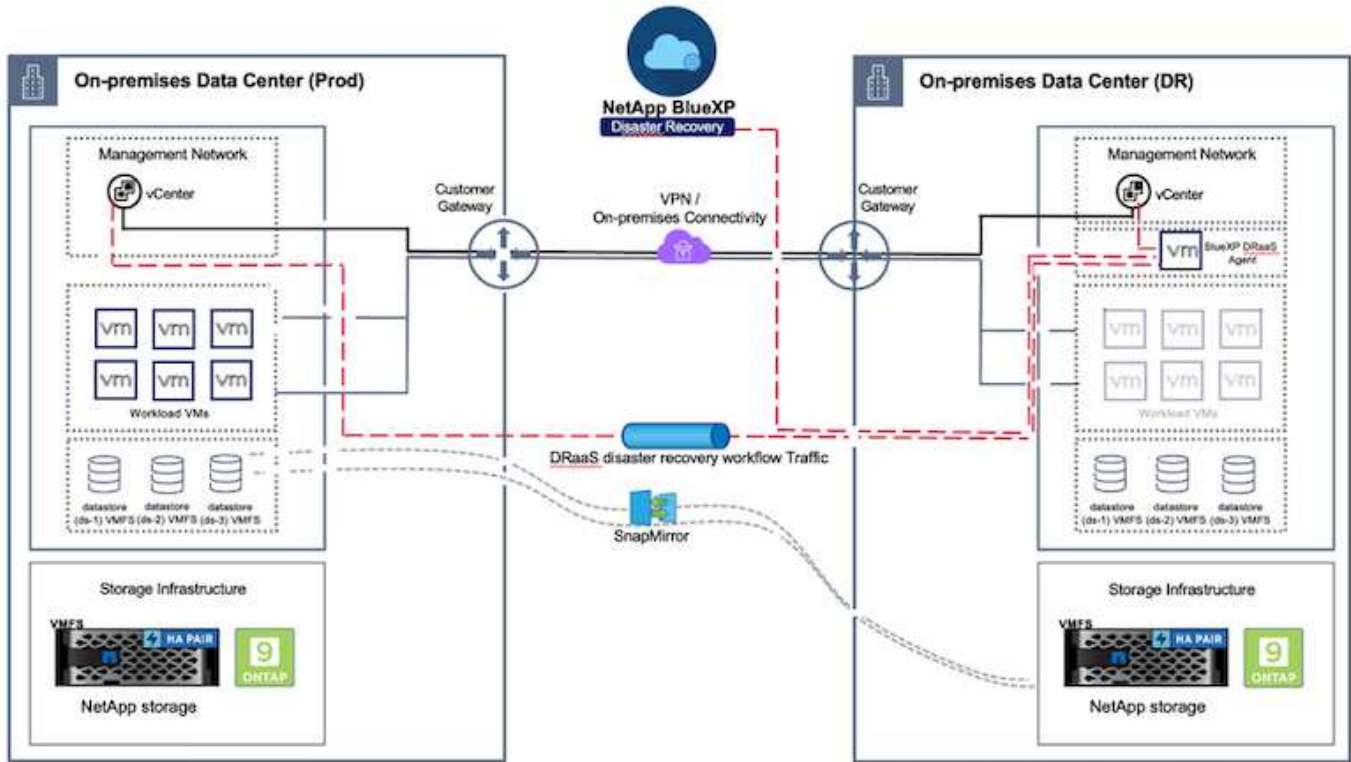
DR using BlueXP DRaaS for VMFS Datastores

Disaster recovery using block-level replication from production site to disaster recovery site is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events, like ransomware attacks. With NetApp SnapMirror replication, VMware workloads running on-premises ONTAP systems using VMFS datastore can be replicated to another ONTAP storage system in a designated recovery datacenter where VMware resides

This section of the document describes the configuration of BlueXP DRaaS to set up disaster recovery for on-premises VMware VMs to another designated site. As part of this setup, the BlueXP account, BlueXP connector, the ONTAP arrays added within BlueXP workspace which is needed to enable communication from VMware vCenter to the ONTAP storage. In addition, this document details how to configure replication between sites and how to setup and test a recovery plan. The last section has instructions for performing a full site failover and how to failback when the primary site is recovered and brought online.

Using the BlueXP disaster recovery service, which is integrated into the NetApp BlueXP console, customers can discover their on-premises VMware vCenters along with ONTAP storage, create resource groupings, create a disaster recovery plan, associate it with resource groups, and test or execute failover and failback. SnapMirror provides storage-level block replication to keep the two sites up to date with incremental changes, resulting in a RPO of up to 5 minutes. It is also possible to simulate DR procedures as a regular drill without impacting the production and replicated datastores or incurring additional storage costs. BlueXP disaster recovery takes advantage of ONTAP's FlexClone technology to create a space-efficient copy of the VMFS datastore from the last replicated Snapshot on the DR site. Once the DR test is complete, customers can

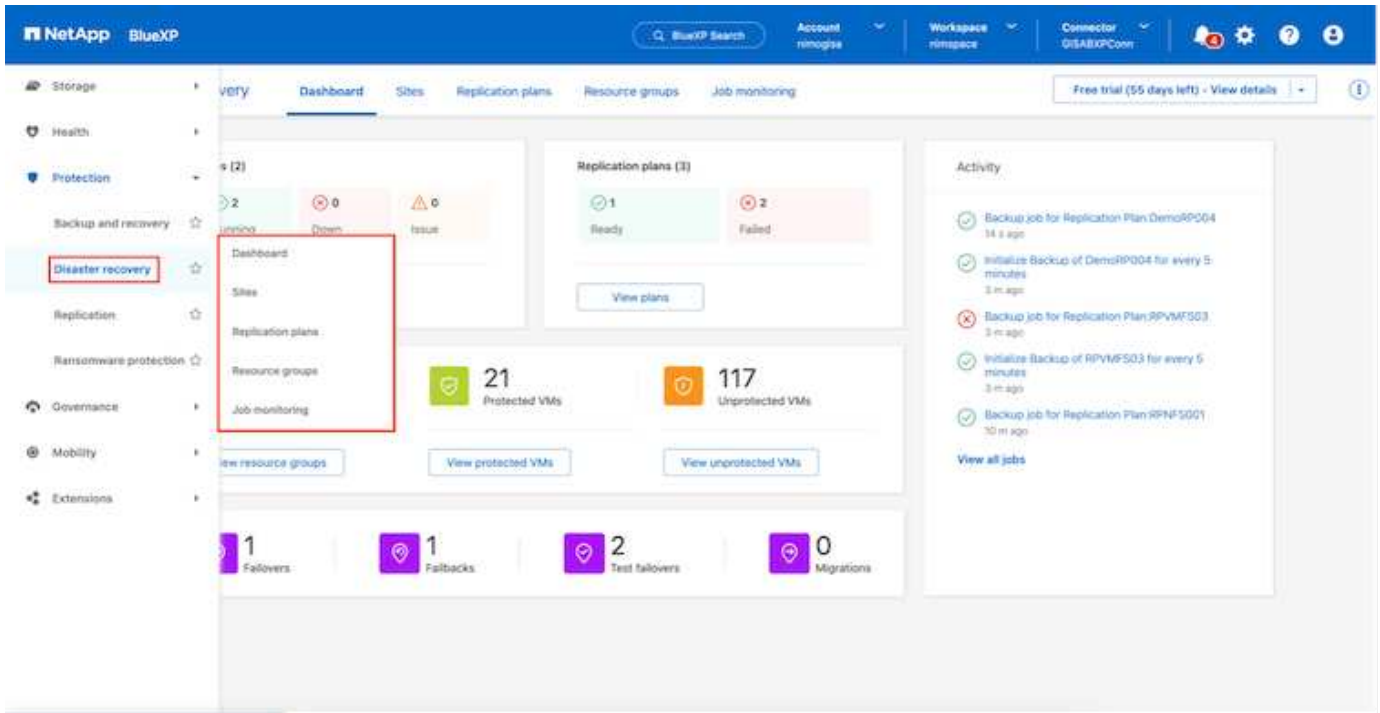
simply delete the test environment, again without any impact to actual replicated production resources. When there is a need (planned or unplanned) for actual failover, with a few clicks, the BlueXP disaster recovery service will orchestrate all the steps needed to automatically bring up the protected virtual machines on designated disaster recovery site. The service will also reverse the SnapMirror relationship to the primary site and replicate any changes from secondary to primary for a failback operation, when needed. All of these can be achieved with a fraction of cost compared to other well-known alternatives.



Getting started

To get started with BlueXP disaster recovery, use BlueXP console and then access the service.

1. Log in to BlueXP.
2. From the BlueXP left navigation, select Protection > Disaster recovery.
3. The BlueXP disaster recovery Dashboard appears.




Before configuring disaster recovery plan, ensure the following pre-requisites are met:

- BlueXP Connector is set up in NetApp BlueXP. The connector should be deployed in AWS VPC.
- BlueXP connector instance have connectivity to the source and destination vCenter and storage systems.
- On-premises NetApp storage systems hosting VMFS datastores for VMware are added in BlueXP.
- DNS resolution should be in place when using DNS names. Otherwise, use IP addresses for the vCenter.
- SnapMirror replication is configured for the designated VMFS based datastore volumes.

Once the connectivity is established between the source and destination sites, proceed with configuration steps, which should take about 3 to 5 minutes.



NetApp recommends deploying the BlueXP connector in the disaster recovery site or in a third site, so that the BlueXP connector can communicate through the network with source and destination resources during real outages or natural disasters.




Support for on-premises to on-premises VMFS datastores is in technology preview while writing this document. The capability is supported with both FC and iSCSI protocol based VMFS datastores.

BlueXP disaster recovery configuration

The first step in preparing for disaster recovery is to discover and add the on-premises vCenter and storage resources to BlueXP disaster recovery.




Ensure the ONTAP storage systems are added to the working environment within the canvas. Open BlueXP console and select **Protection > Disaster Recovery** from left navigation. Select **Discover vCenter servers** or use top menu, Select **Sites > Add > Add vCenter**.




Add the following platforms:

- **Source.** On-premises vCenter.



- **Destination.** VMC SDDC vCenter.



Once the vCenters are added, automated discovery is triggered.

Configuring Storage replication between source and destination site

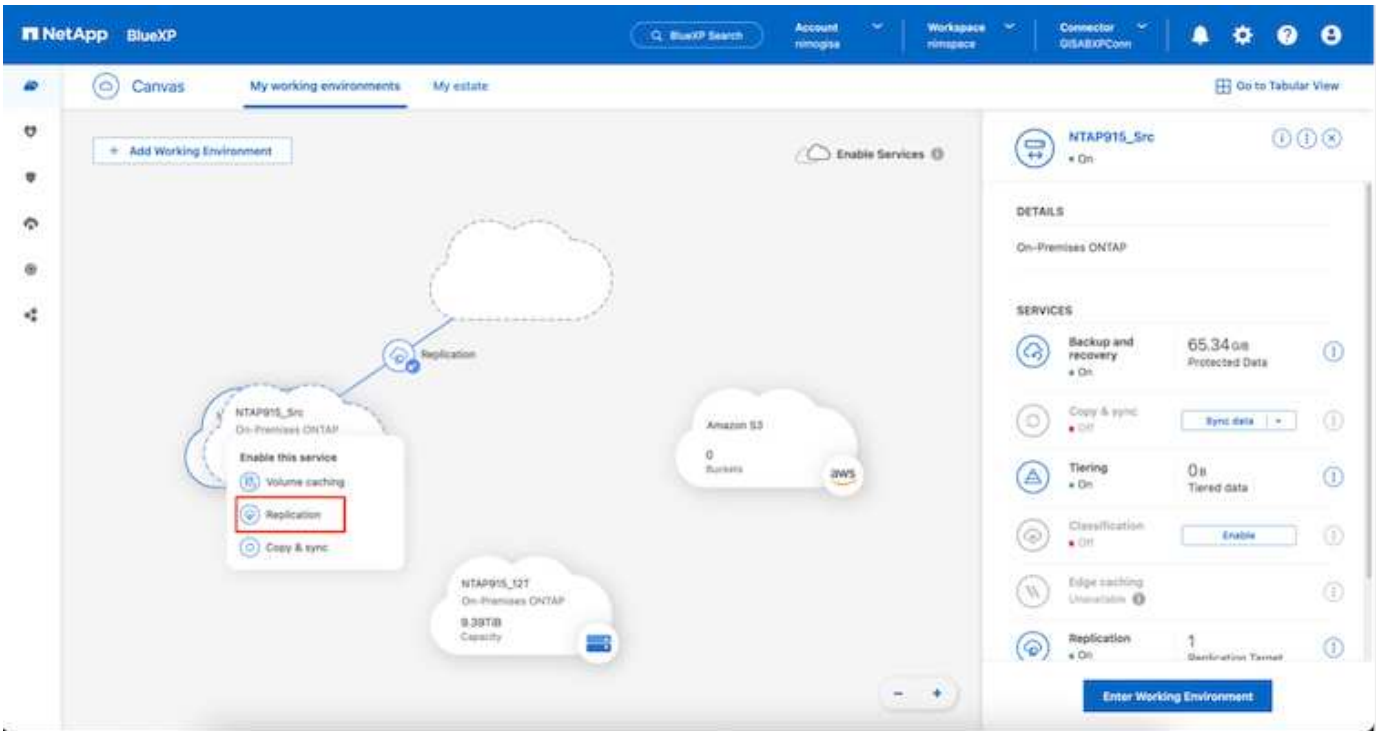
SnapMirror makes use of ONTAP snapshots to manage the transfer of data from one location to another. Initially, a full copy based on a snapshot of the source volume is copied over to the destination to perform a baseline synchronization. As data changes occur at the source, a new snapshot is created and compared to the baseline snapshot. The blocks found to have changed are then replicated to the destination, with the newer snapshot becoming the current baseline, or newest common snapshot. This enables the process to be repeated and incremental updates to be sent to the destination.

When a SnapMirror relationship has been established, the destination volume is in an online read-only state, and so is still accessible. SnapMirror works with physical blocks of storage, rather than at a file or other logical level. This means that the destination volume is an identical replica of the source, including snapshots, volume settings, etc. If ONTAP space efficiency features, such as data compression and data deduplication, are being used by the source volume, the replicated volume will retain these optimizations.

Breaking the SnapMirror relationship makes the destination volume writable and would typically be used to perform a failover when SnapMirror is being used to synchronize data to a DR environment. SnapMirror is sophisticated enough to allow the data changed at the failover site to be efficiently resynchronized back to the primary system, should it later come back online, and then allow for the original SnapMirror relationship to be re-established.


How to set it up for VMware Disaster Recovery

The process to create SnapMirror replication remains the same for any given application. The process can be manual or automated. The easiest way is to leverage BlueXP to configure SnapMirror replication by using simple drag & drop of the source ONTAP system in the environment onto the destination to trigger the wizard that guides through the rest of the process.



BlueXP DRaaS can also automate the same provided the following two criteria's are met:

- Source and destination clusters have a peer relationship.
- Source SVM and destination SVM have a peer relationship.



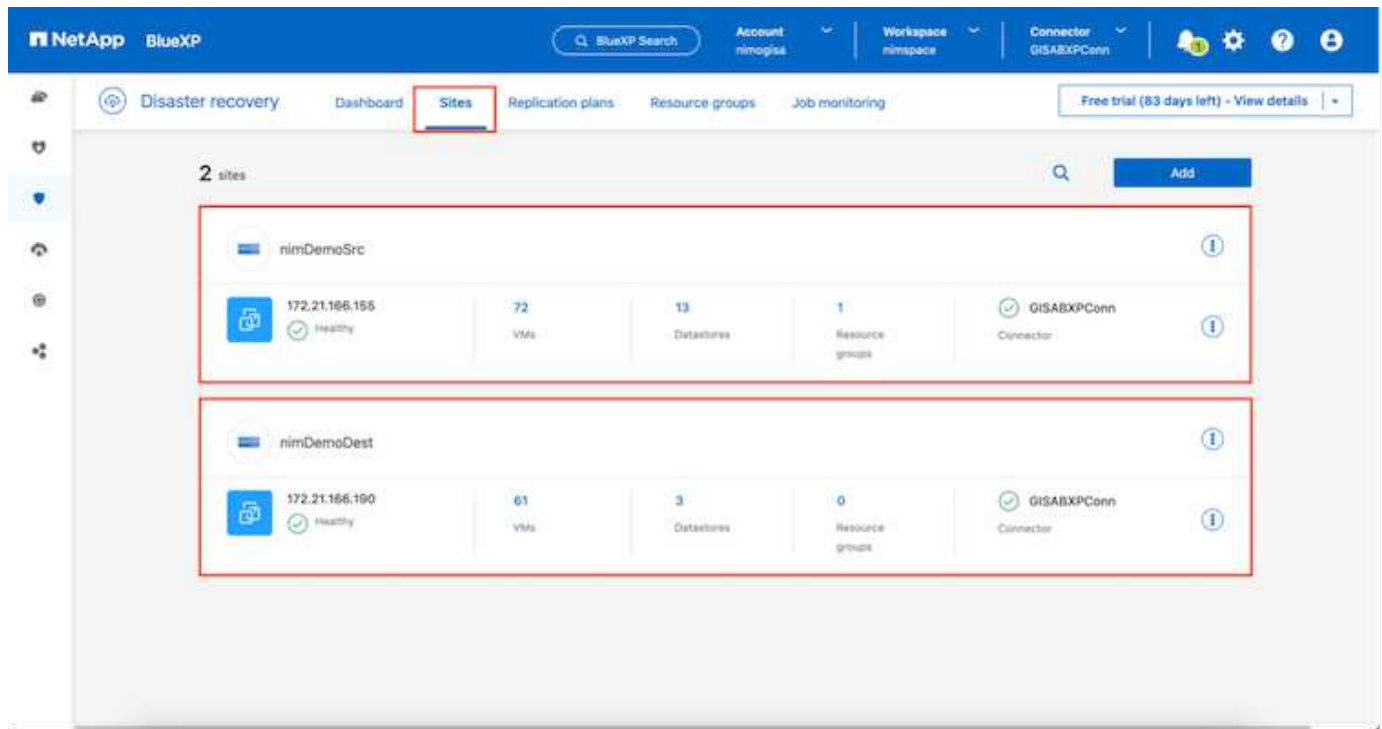
If SnapMirror relationship is already configured for the volume via CLI, BlueXP DRaaS picks up the relationship and continues with the rest of the workflow operations.




Apart from the above approaches, SnapMirror replication can also be created via ONTAP CLI or System Manager. Irrespective of the approach used to synchronize the data using SnapMirror, BlueXP DRaaS orchestrates the workflow for seamless and efficient disaster recovery operations.

What can BlueXP disaster recovery do for you?

After the source and destination sites are added, BlueXP disaster recovery performs automatic deep discovery and displays the VMs along with associated metadata. BlueXP disaster recovery also automatically detects the networks and port groups used by the VMs and populates them.

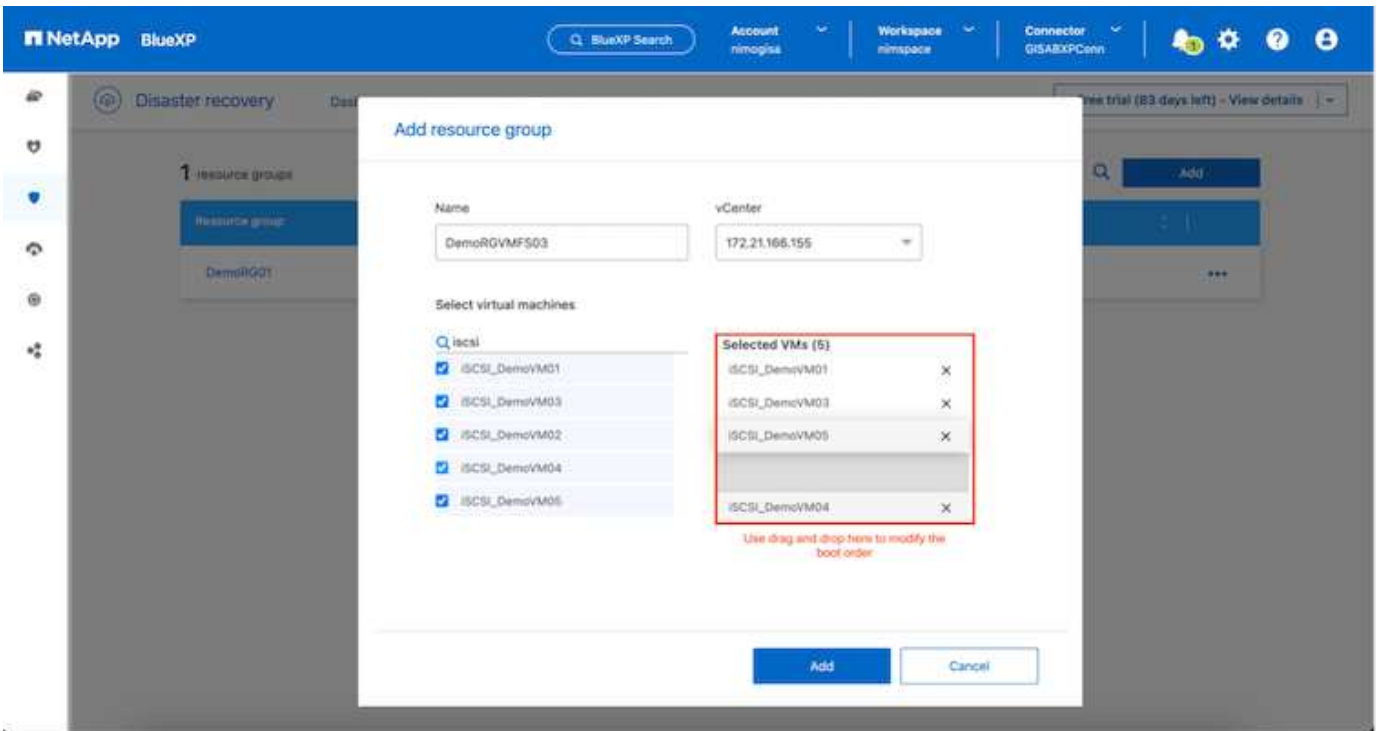


After the sites have been added, VMs can be grouped into resource groups. BlueXP disaster recovery resource groups allow you to group a set of dependent VMs into logical groups that contain their boot orders and boot delays that can be executed upon recovery. To start creating resource groups, navigate to **Resource Groups** and click **Create New Resource Group**.





The resource group can also be created while creating a replication plan.

The boot order of the VMs can be defined or modified during the creation of resource groups by using simple drag and drop mechanism.




Once the resource groups are created, the next step is to create the execution blueprint or a plan to recover virtual machines and applications in the event of a disaster. As mentioned in the prerequisites, SnapMirror replication can be configured beforehand or DRaaS can configure it using the RPO and retention count specified during creation of the replication plan.

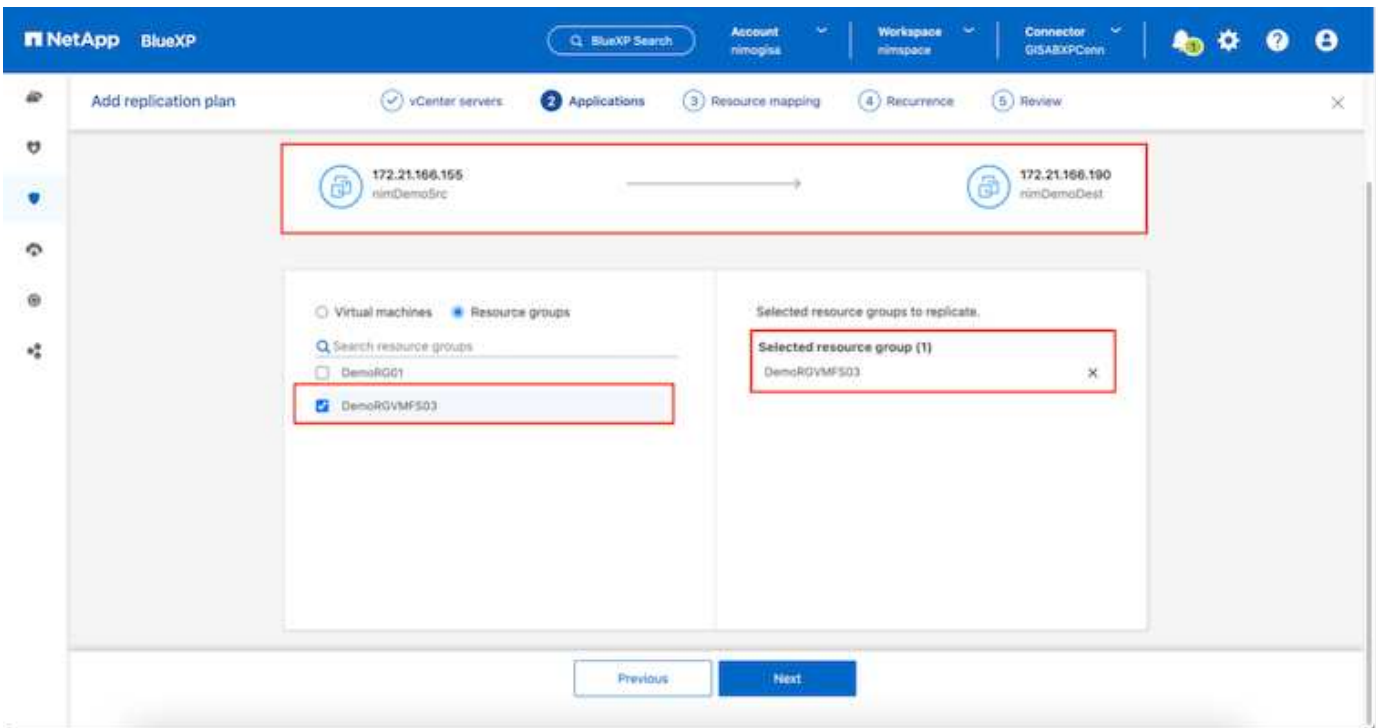


Configure the replication plan by selecting the source and destination vCenter platforms from the drop down and pick the resource groups to be included in the plan, along with the grouping of how applications should be restored and powered on and mapping of clusters and networks. To define the recovery plan, navigate to the **Replication Plan** tab and click **Add Plan**.

First, select the source vCenter and then select the destination vCenter.



The next step is to select existing resource groups. If no resource groups created, then the wizard helps to group the required virtual machines (basically create functional resource groups) based on the recovery objectives. This also helps define the operation sequence of how application virtual machines should be restored.

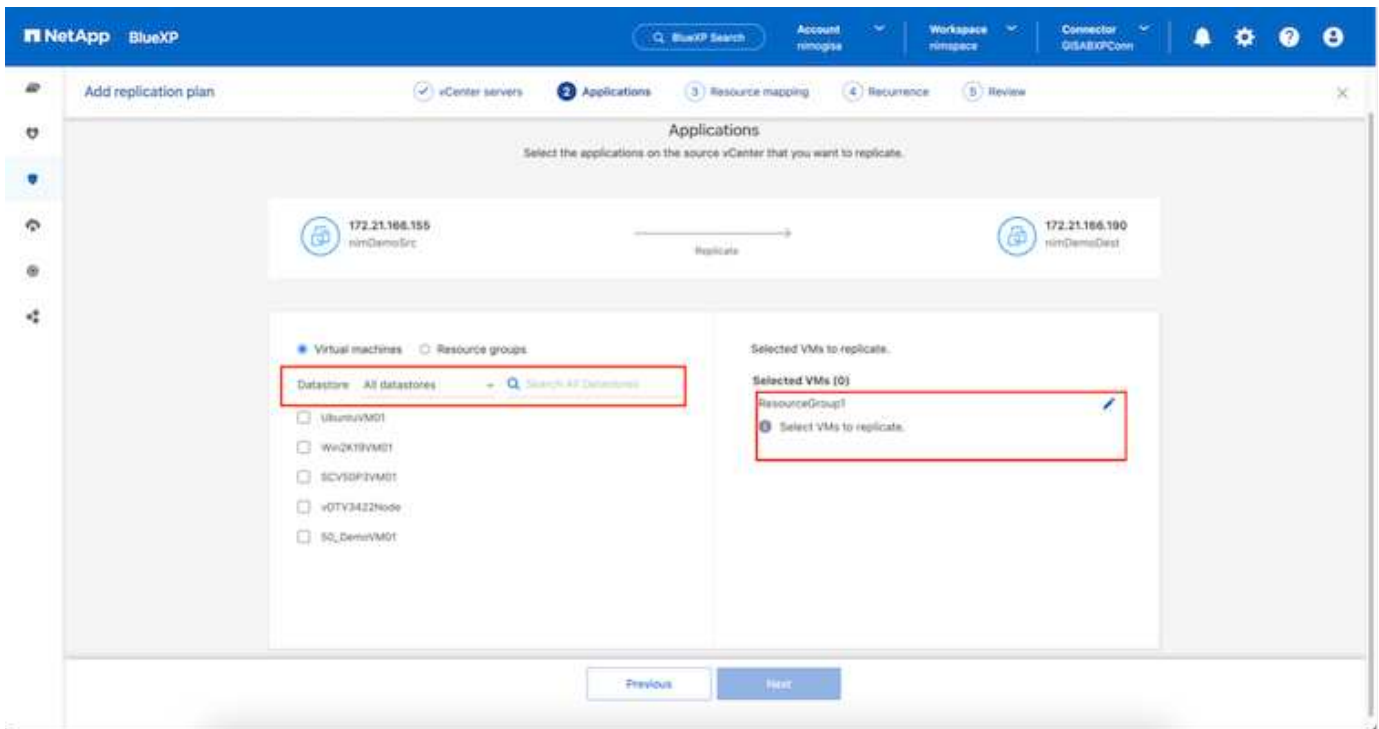


Resource group allows to set boot order using the drag and drop functionality. It can be used to easily modify the order in which the VMs would be powered on during the recovery process.

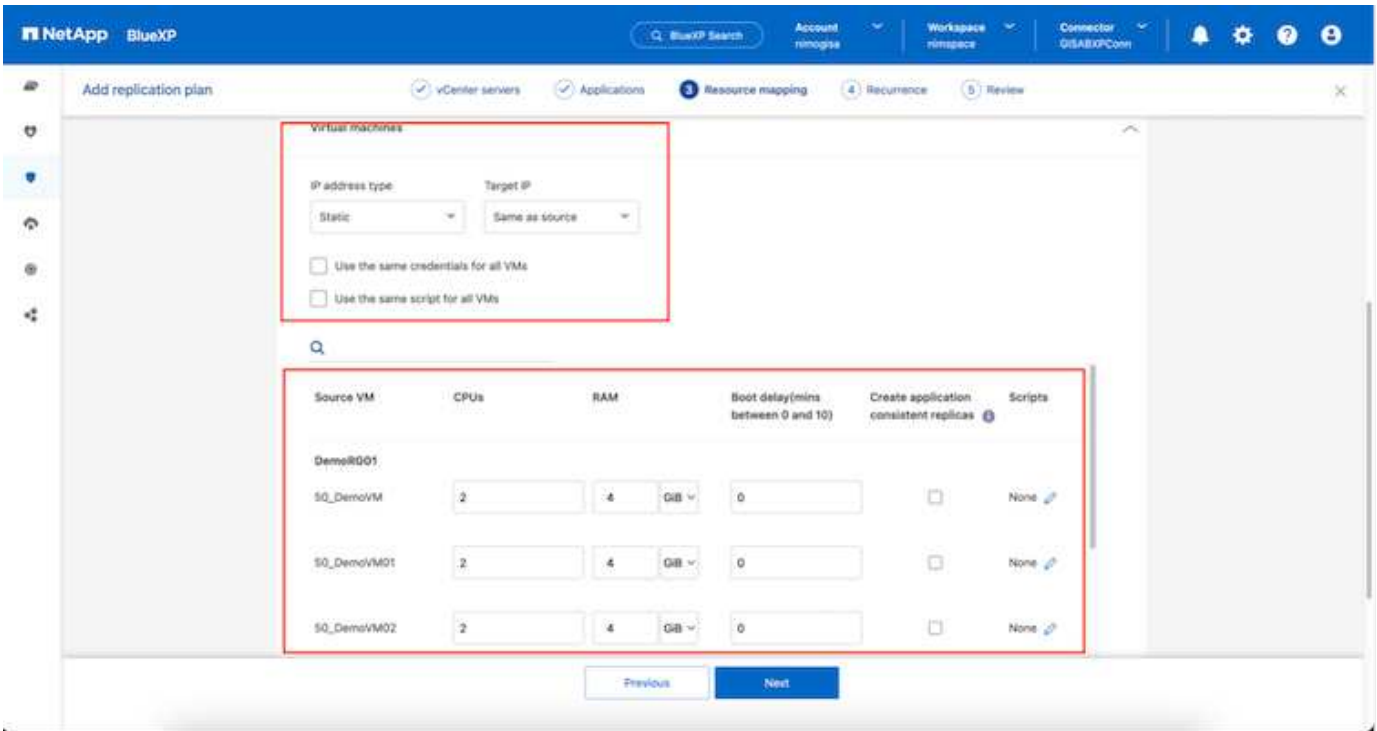


Each virtual machine within a resource group is started in sequence based on the order. Two resource groups are started in parallel.


The below screenshot shows the option to filter virtual machines or specific datastores based on organizational requirements if resource groups are not created beforehand.



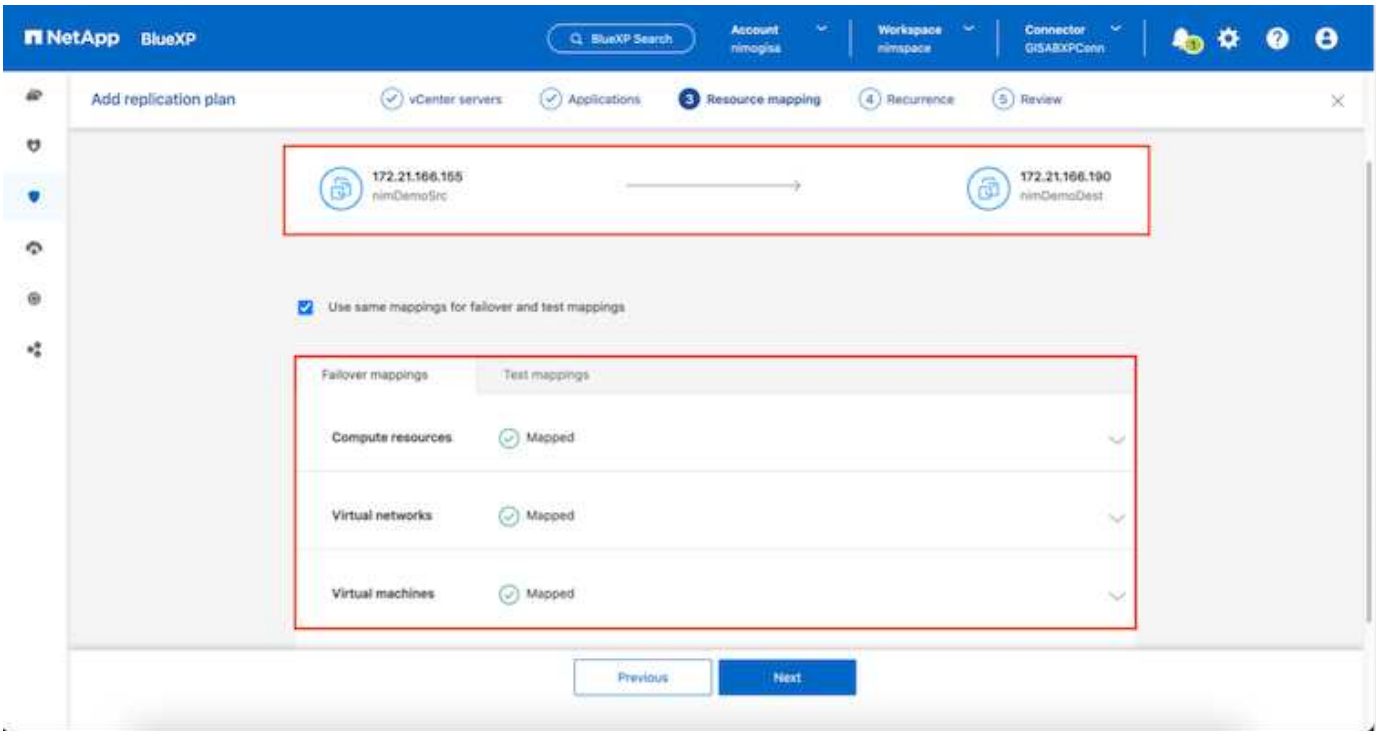
Once the resource groups are selected, create the failover mappings. In this step, specify how the resources from the source environment maps to the destination. This includes compute resources, virtual networks. IP customization, pre- and post-scripts, boot delays, application consistency and so on. For detailed information, refer to [Create a replication plan](#).



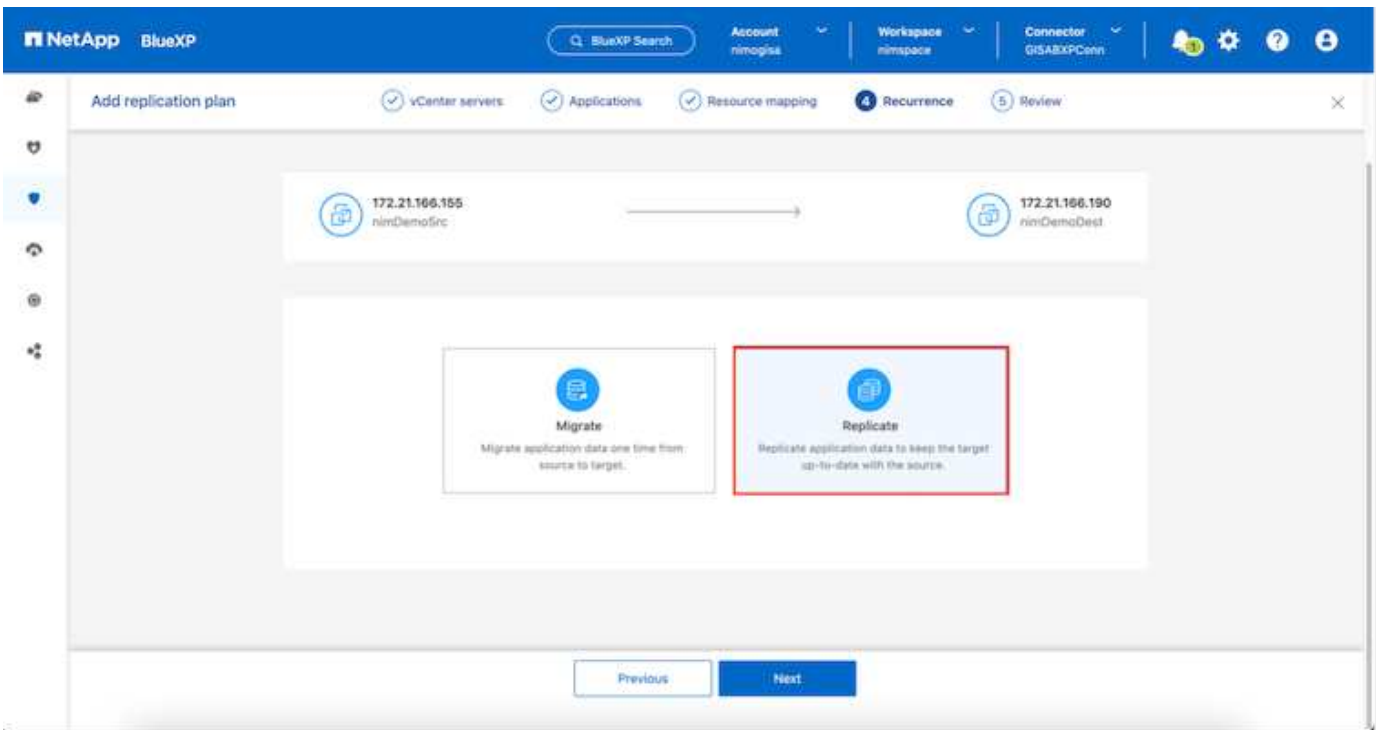
By default, same mapping parameters are used for both test and failover operations. To apply different mappings for test environment, select the Test mapping option after unchecking the checkbox as shown below:




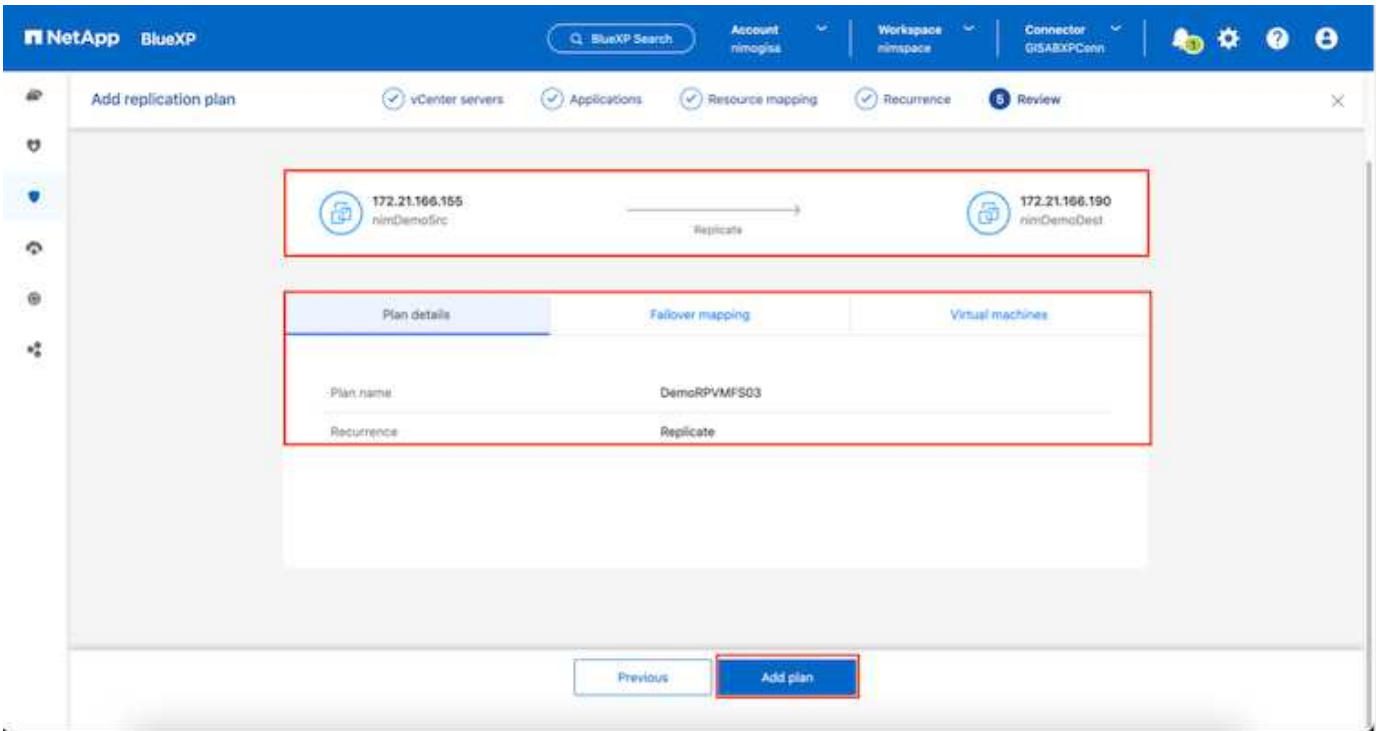
Once the resource mapping is complete, click Next.



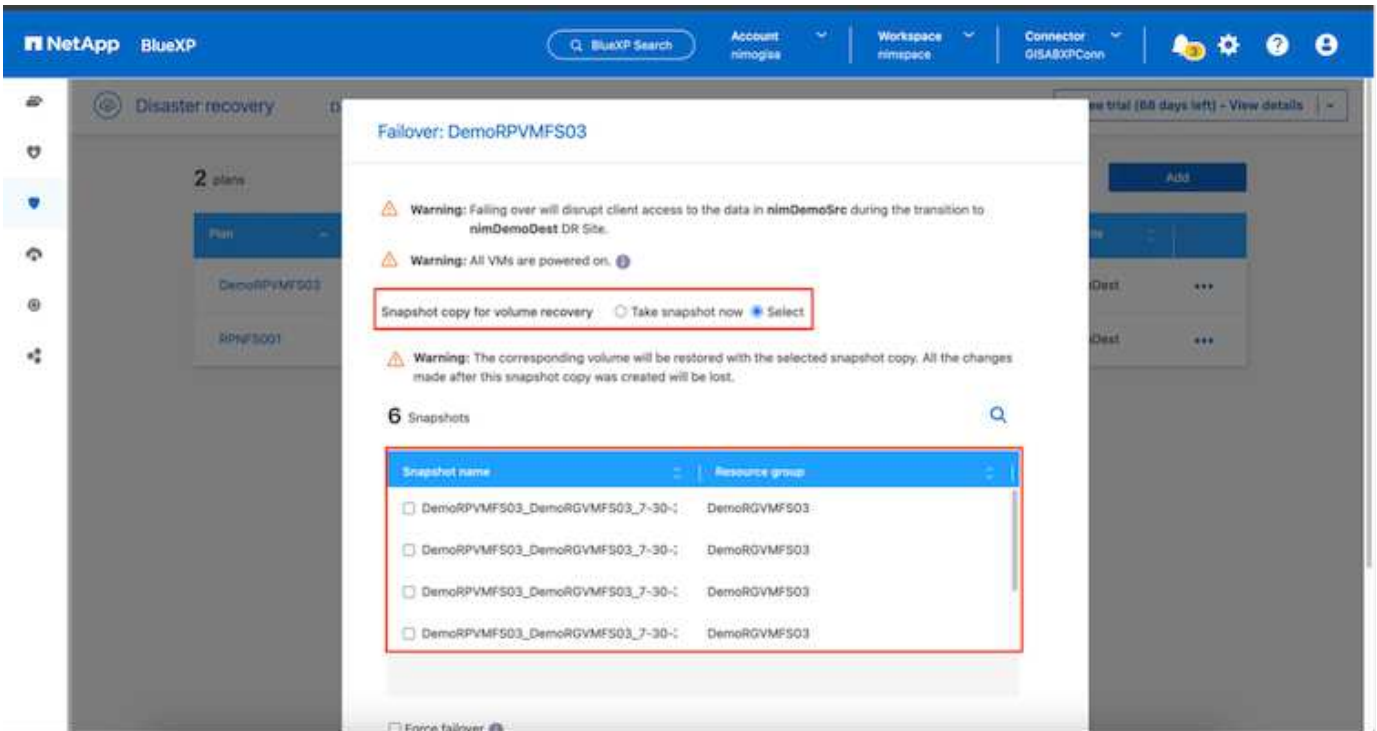
Select the recurrence type. In simple words, select Migrate (one time migration using failover) or recurring continuous replication option. In this walkthrough, Replicate option is selected.




Once done, review the created mappings and then click on Add plan.



Once the replication plan is created, failover can be performed depending on the requirements by selecting the failover option, test-failover option, or the migrate option. BlueXP disaster recovery ensures that the replication process is being executed according to the plan every 30 minutes. During the failover and test-failover options, you can use the most recent SnapMirror Snapshot copy, or you can select a specific Snapshot copy from a point-in-time Snapshot copy (per the retention policy of SnapMirror). The point-in-time option can be very helpful if there is a corruption event like ransomware, where the most recent replicas are already compromised or encrypted. BlueXP disaster recovery shows all available recovery points.



To trigger failover or test failover with the configuration specified in the replication plan, click on **Failover** or **Test failover**.



What happens during a failover or test failover operation?


During a test failover operation, BlueXP disaster recovery creates a FlexClone volume on the destination ONTAP storage system using the latest Snapshot copy or a selected snapshot of the destination volume.




A test failover operation creates a cloned volume on the destination ONTAP storage system.




Running a test recovery operation does not affect the SnapMirror replication.



During the process, BlueXP disaster recovery does not map the original target volume. Instead, it makes a new FlexClone volume from the selected Snapshot and a temporary datastore backing the FlexClone volume is mapped to the ESXi hosts.






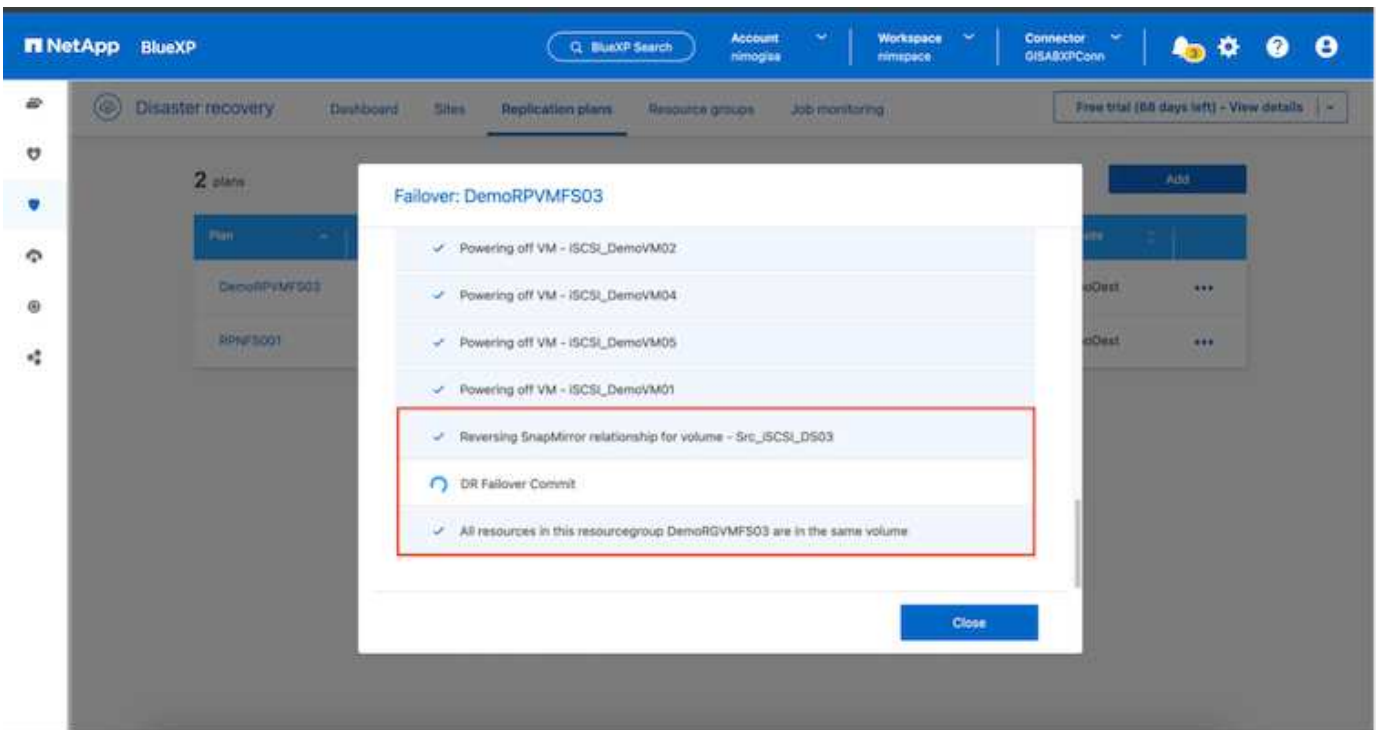
When the test failover operation completes, the cleanup operation can be triggered using “**Clean Up failover test**”. During this operation, BlueXP disaster recovery destroys the FlexClone volume that was used in the operation.

In the event of real disaster event occurs, BlueXP disaster recovery performs the following steps:

1. Breaks the SnapMirror relationship between the sites.
2. Mounts the VMFS datastore volume after resignature for immediate use.
3. Register the VMs
4. Power on VMs

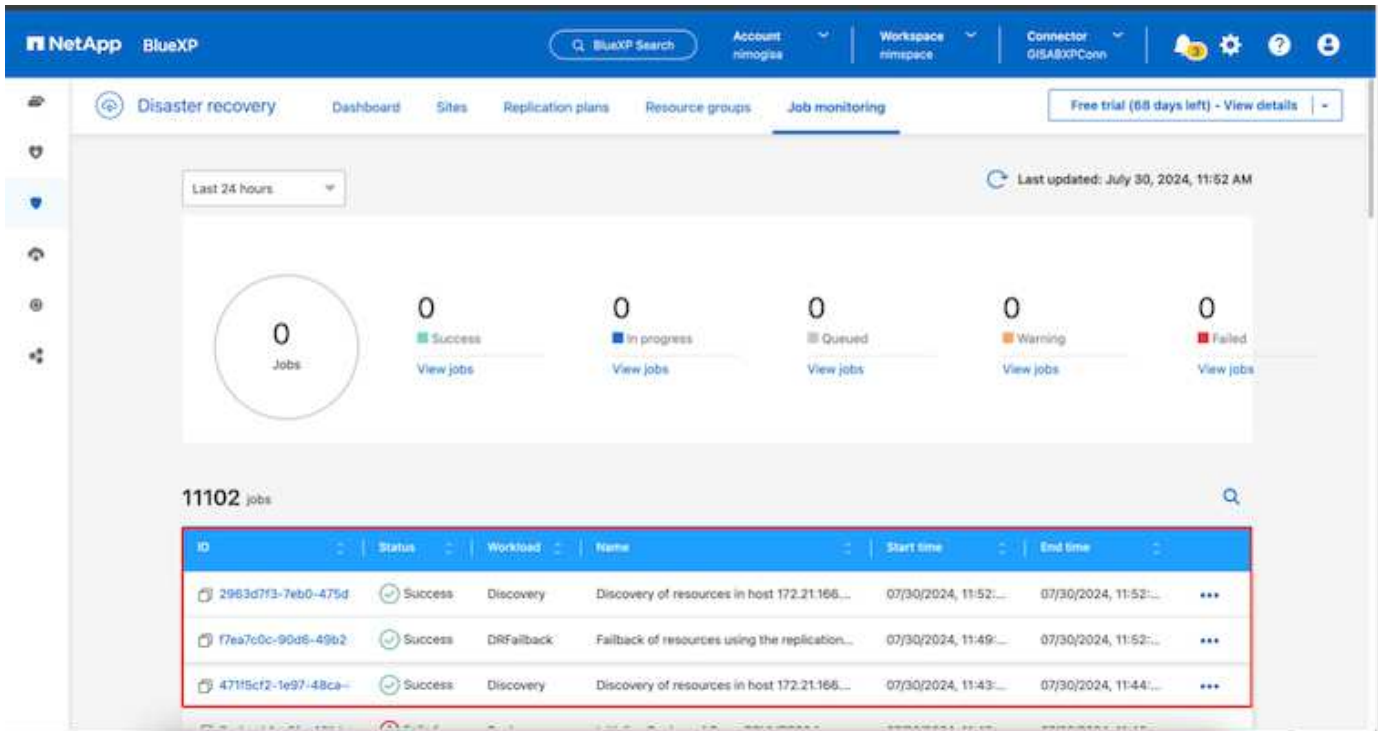


Once the primary site is up and running, BlueXP disaster recovery enables reverse resync for SnapMirror and enables failback, which again can be performed with the click of a button.



And if migrate option is chosen, it is considered as a planned failover event. In this case, an additional step is triggered which is to shut down the virtual machines at the source site. The rest of the steps remains the same as failover event.

From BlueXP or the ONTAP CLI, you can monitor the replication health status for the appropriate datastore volumes, and the status of a failover or test failover can be tracked via Job Monitoring.



This provides a powerful solution to handle a tailored and customized disaster recovery plan. Failover can be done as planned failover or failover with a click of a button when disaster occurs and decision is made to activate the DR site.

To learn more about this process, feel free to follow the detailed walkthrough video or use the [solution simulator](#).

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.