



TR-4955: Disaster Recovery with FSx ONTAP and VMC (AWS VMware Cloud)

NetApp public and hybrid cloud solutions

NetApp

July 30, 2025

Table of Contents

TR-4955: Disaster Recovery with FSx ONTAP and VMC (AWS VMware Cloud)	1
Overview	1
Getting started	1
Deploy and configure VMware Cloud on AWS	1
Provision and configure FSx ONTAP	2
Deploy and configure SnapMirror to FSx ONTAP	2
DRO installation	2
Prerequisites	3
OS requirements	3
Install the package	3
DRO configuration	4
Resource groupings	6
Replication plans	7
Ransomware recovery	15
Benefits	15

TR-4955: Disaster Recovery with FSx ONTAP and VMC (AWS VMware Cloud)

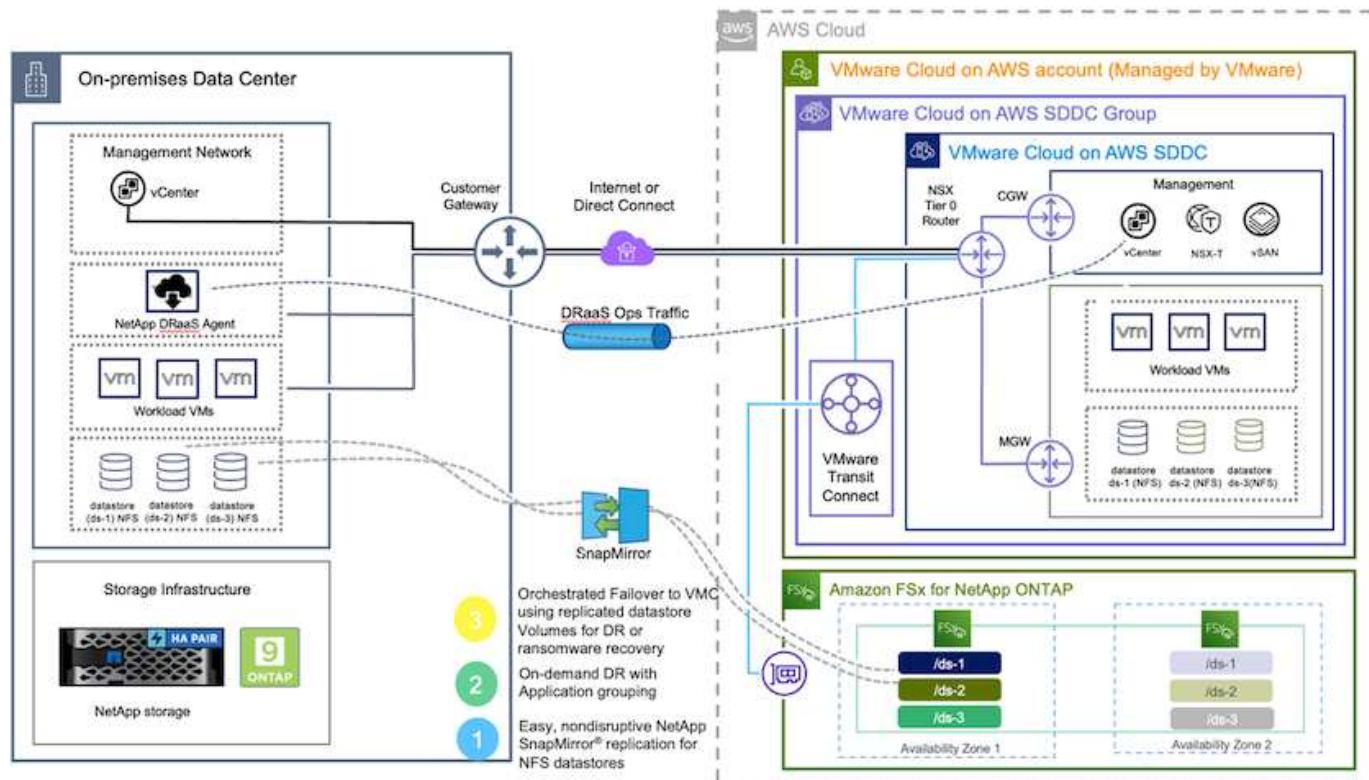
Disaster Recovery Orchestrator (DRO; a scripted solution with UI) can be used to seamlessly recover workloads replicated from on-premises to FSx ONTAP. DRO automates the recovery from the SnapMirror level, through VM registration to VMC, to network mappings directly on NSX-T. This feature is included with all VMC environments.

Niyaz Mohamed, NetApp

Overview

Disaster recovery to cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware). With NetApp SnapMirror technology, on-premises VMware workloads can be replicated to FSx ONTAP running in AWS.

Disaster Recovery Orchestrator (DRO; a scripted solution with UI) can be used to seamlessly recover workloads replicated from on-premises to FSx ONTAP. DRO automates the recovery from the SnapMirror level, through VM registration to VMC, to network mappings directly on NSX-T. This feature is included with all VMC environments.



Getting started

Deploy and configure VMware Cloud on AWS

VMware Cloud on AWS provides a cloud-native experience for VMware-based workloads in the AWS

ecosystem. Each VMware Software-Defined Data Center (SDDC) runs in an Amazon Virtual Private Cloud (VPC) and provides a full VMware stack (including vCenter Server), NSX-T software-defined networking, vSAN software-defined storage, and one or more ESXi hosts that provide compute and storage resources to the workloads. To configure a VMC environment on AWS, follow the steps at this [link](#). A pilot-light cluster can also be used for DR purposes.



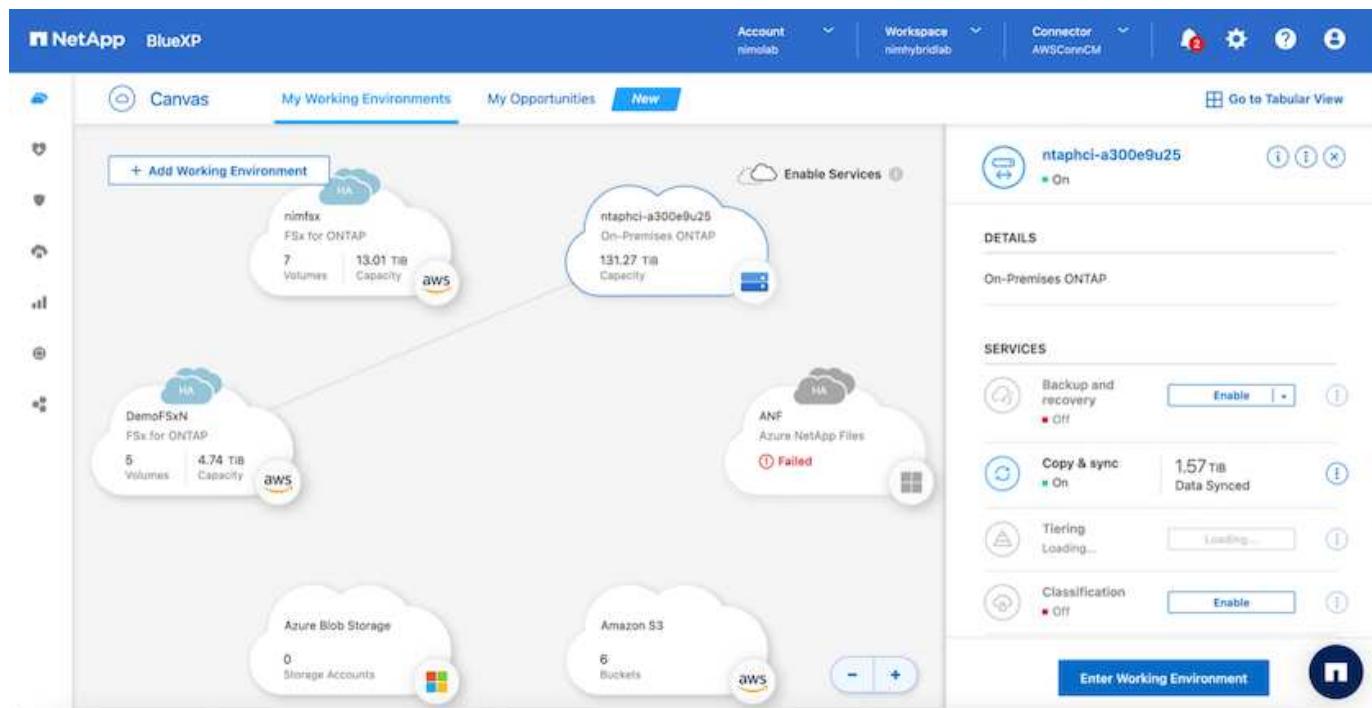
In the initial release, DRO supports an existing pilot-light cluster. On-demand SDDC creation will be available in an upcoming release.

Provision and configure FSx ONTAP

Amazon FSx ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on the popular NetApp ONTAP file system. Follow the steps at this [link](#) to provision and configure FSx ONTAP.

Deploy and configure SnapMirror to FSx ONTAP

The next step is to use NetApp BlueXP and discover the provisioned FSx ONTAP on AWS instance and replicate the desired datastore volumes from an on-premises environment to FSx ONTAP with the appropriate frequency and NetApp Snapshot copy retention:



Follow the steps in this [link](#) to configure BlueXP. You can also use the NetApp ONTAP CLI to schedule replication following this [link](#).



A SnapMirror relationship is a prerequisite and must be created beforehand.

DRO installation

To get started with DRO, use the Ubuntu operating system on a designated EC2 instance or virtual machine to make sure you meet the prerequisites. Then install the package.

Prerequisites

- Make sure that connectivity to the source and destination vCenter and storage systems exists.
- DNS resolution should be in place if you are using DNS names. Otherwise, you should use IP addresses for the vCenter and storage systems.
- Create a user with root permissions. You can also use sudo with an EC2 instance.

OS requirements

- Ubuntu 20.04 (LTS) with minimum of 2GB and 4 vCPUs
- The following packages must be installed on the designated agent VM:
 - Docker
 - Docker-compose
 - Jq

Change permissions on docker.sock: sudo chmod 666 /var/run/docker.sock.



The deploy.sh script executes all the required prerequisites.

Install the package

1. Download the installation package on the designated virtual machine:

```
git clone https://github.com/NetApp/DRO-AWS.git
```



The agent can be installed on-premises or within an AWS VPC.

2. Unzip the package, run the deployment script, and enter the host IP (for example, 10.10.10.10).

```
tar xvf DRO-prereq.tar
```

3. Navigate to the directory and run the deploy script as follows:

```
sudo sh deploy.sh
```

4. Access the UI using:

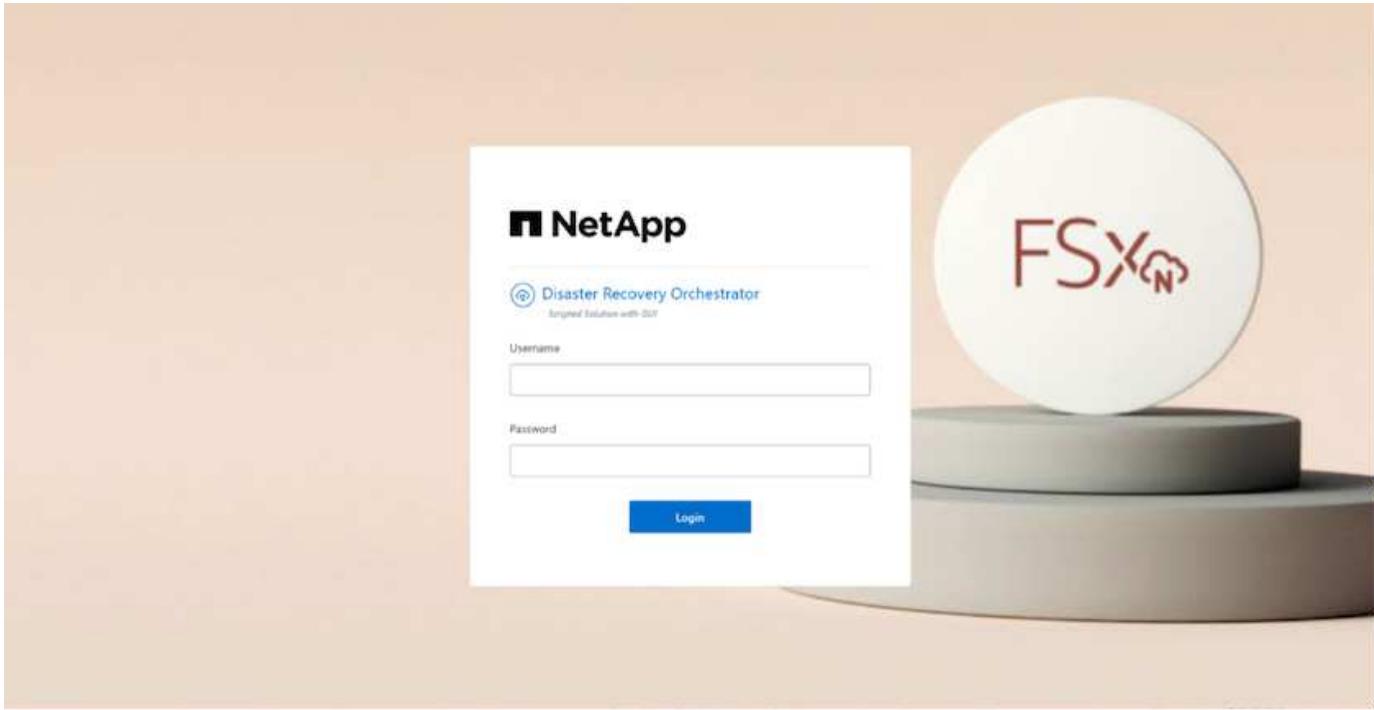
```
https://<host-ip-address>
```

with the following default credentials:

Username: admin
Password: admin



The password can be changed using the "Change Password" option.



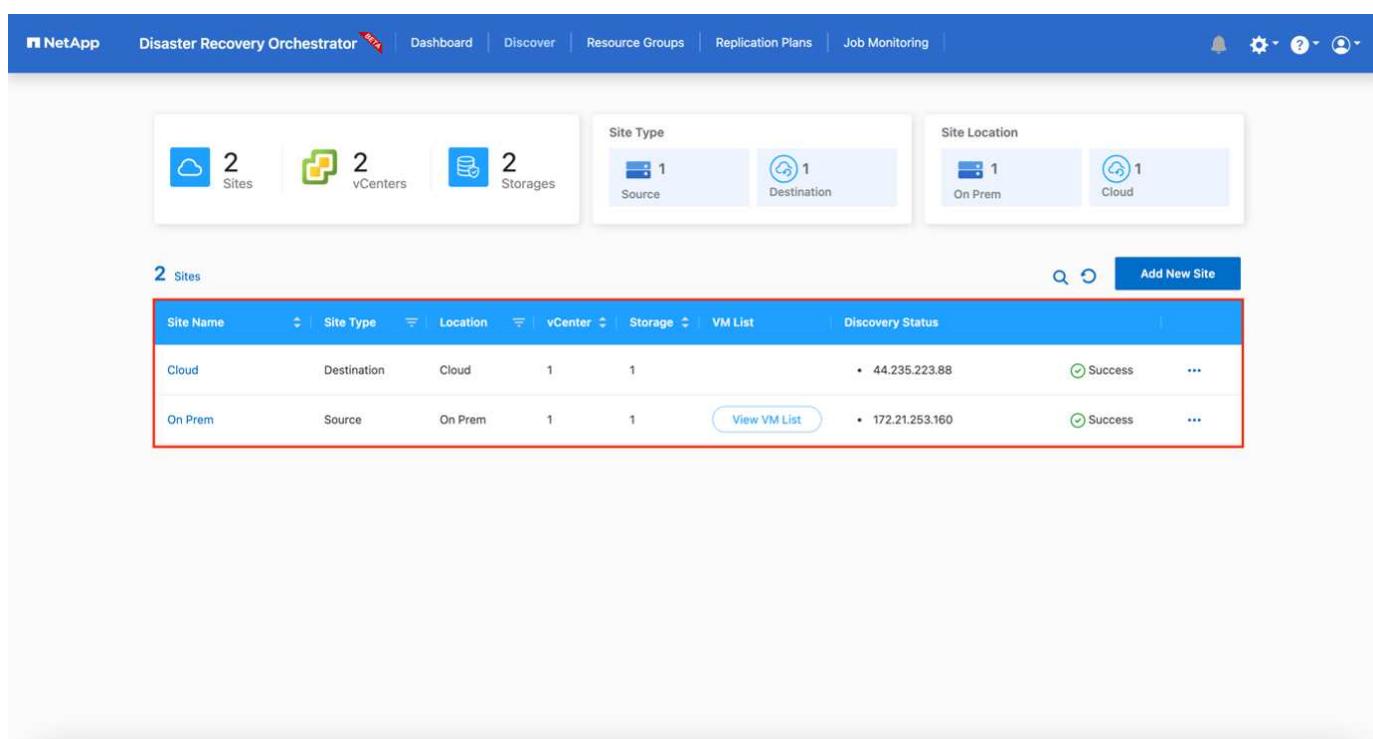
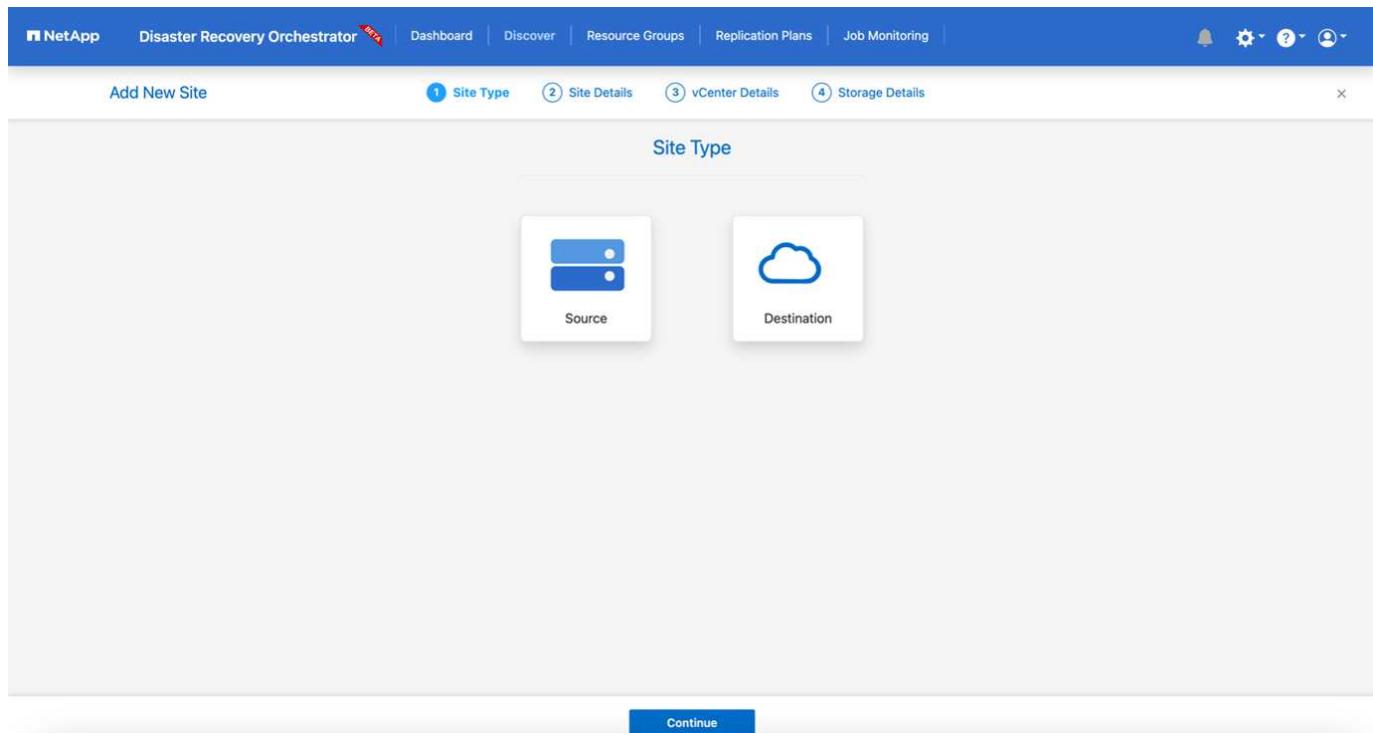
DRO configuration

After FSx ONTAP and VMC have been configured properly, you can begin configuring DRO to automate the recovery of on-premises workloads to VMC by using the read-only SnapMirror copies on FSx ONTAP.

NetApp recommends deploying the DRO agent in AWS and also to the same VPC where FSx ONTAP is deployed (it can be peer connected too), so that the DRO agent can communicate through the network with your on-premises components as well as with the FSx ONTAP and VMC resources.

The first step is to discover and add the on-premises and cloud resources (both vCenter and storage) to DRO. Open DRO in a supported browser and use the default username and password (admin/admin) and Add Sites. Sites can also be added using the Discover option. Add the following platforms:

- On-premises
 - On-premises vCenter
 - ONTAP storage system
- Cloud
 - VMC vCenter
 - FSx ONTAP



Once added, DRO performs automatic discovery and displays the VMs that have corresponding SnapMirror replicas from the source storage to FSx ONTAP. DRO automatically detects the networks and portgroups used by the VMs and populates them.

VM List

Site: On Prem | vCenter: 172.21.253.160

VM Protection: 3 Protected, 216 Unprotected

VM Name	VM Status	VM State (1)	DataStore	CPU	Memory (MB)
a300-vcsa02	Not Protected	Powered On	A300_NFS_DS04	16	65536
PFsense	Not Protected	Powered On	A300_NFS_DS04	4	8192
PFsense260	Not Protected	Powered On	A300_NFS_DS04	4	16384
NimDC02	Not Protected	Powered On	A300_NFS_DS04	4	8192
jh8Bhja-187	Not Protected	Powered On	A300_NFS_DS04	4	16384
jhNmo-187	Not Protected	Powered On	A300_NFS_DS04	4	16384
NimMSdesktop	Not Protected	Powered On	A300_NFS_DS04	8	12288

The next step is to group the required VMs into functional groups to serve as resource groups.

Resource groupings

After the platforms have been added, you can group the VMs you want to recover into resource groups. DRO resource groups allow you to group a set of dependent VMs into logical groups that contain their boot orders, boot delays, and optional application validations that can be executed upon recovery.

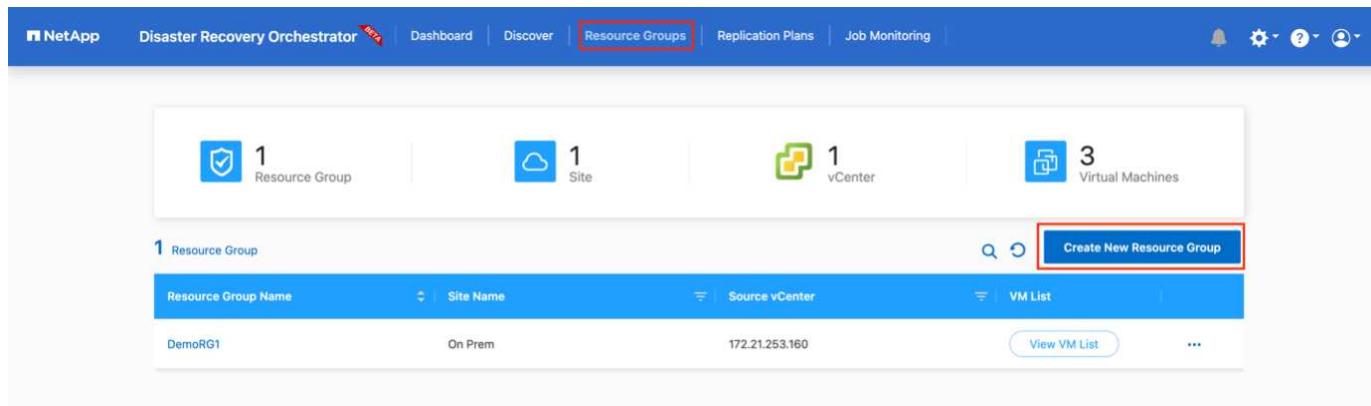
To start creating resource groups, complete the following steps:

1. Access **Resource Groups**, and click **Create New Resource Group**.
2. Under **New resource group**, select the source site from the dropdown and click **Create**.
3. Provide **Resource Group Details** and click **Continue**.
4. Select the appropriate VMs using the search option.
5. Select the boot order and boot delay (secs) for the selected VMs. Set the order of the power-on sequence by selecting each VM and setting up the priority for it. Three is the default value for all VMs.

Options are as follows:

- 1 – The first virtual machine to power on
- 3 – Default
- 5 – The last virtual machine to power on

6. Click **Create Resource Group**.



NetApp Disaster Recovery Orchestrator Dashboard Discover Resource Groups Replication Plans Job Monitoring

1 Resource Group 1 Site 1 vCenter 3 Virtual Machines

1 Resource Group

Resource Group Name Site Name Source vCenter VM List

DemoRG1 On Prem 172.21.253.160 View VM List

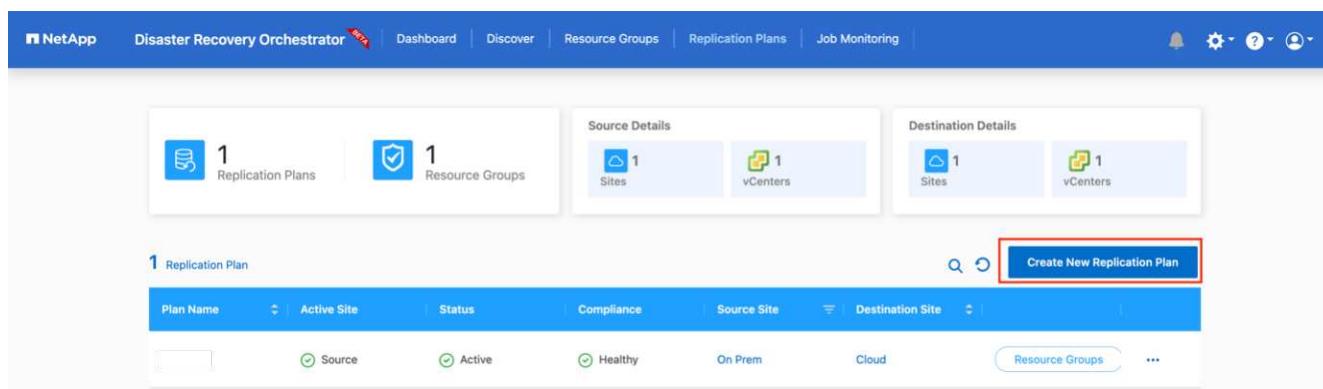
Create New Resource Group

Replication plans

You need a plan to recover applications in the event of a disaster. Select the source and destination vCenter platforms from the drop down and pick the resource groups to be included in this plan, along with the grouping of how applications should be restored and powered on (for example, domain controllers, then tier-1, then tier-2, and so on). Such plans are sometimes also called blueprints. To define the recovery plan, navigate to the **Replication Plan** tab and click **New Replication Plan**.

To start creating a replication plan, complete the following steps:

1. Access **Replication Plans**, and click **Create New Replication Plan**.



NetApp Disaster Recovery Orchestrator Dashboard Discover Resource Groups Replication Plans Job Monitoring

1 Replication Plans 1 Resource Groups

Source Details Destination Details

1 Sites 1 vCenters 1 Sites 1 vCenters

1 Replication Plan

Plan Name Active Site Status Compliance Source Site Destination Site

Source Active Healthy On Prem Cloud Resource Groups

Create New Replication Plan

2. Under **New Replication Plan**, provide a name for the plan and add recovery mappings by selecting the source site, associated vCenter, destination site, and associated vCenter.

NetApp Disaster Recovery Orchestrator

Replication Plans

Create New Replication Plan

1 Replication Plan and Site Details 2 Select Resource Groups 3 Set Execution Order 4 Set VM Details

Replication Plan Details

Plan Name:

Recovery Mapping

Source Site: Select Source Site

Destination Site: Select Destination Site

Source vCenter: Select Source vCenter

Destination vCenter: Select Destination vCenter

Pre-requisite - You must configure SnapMirror relationships between the source site and target site to create successful replication plan

Continue

3. After Recovery mapping is completed, select the cluster mapping.

NetApp Disaster Recovery Orchestrator

Replication Plans

Create New Replication Plan

1 Replication Plan and Site Details 2 Select Resource Groups 3 Set Execution Order 4 Set VM Details

Replication Plan Details

Plan Name: DemoRP

Recovery Mapping

Source Site: On Prem

Destination Site: Cloud

Source vCenter: 172.21.253.160

Destination vCenter: 44.235.223.88

Cluster Mapping

Source Site Resource: TempCluster

Destination Site Resource: Cluster-1

Add

Source Resource	Destination Resource	
A300-Cluster01	Cluster-1	Delete

Continue

4. Select **Resource Group Details** and click **Continue**.

5. Set the execution order for the resource group. This option enables you to select the sequence of operations when multiple resource groups exist.

6. After you are done, select the network mapping to the appropriate segment. The segments should already be provisioned within VMC, so select the appropriate segment to map the VM.

7. Based on the selection of VMs, datastore mappings are automatically selected.



SnapMirror is at the volume level. Therefore, all VMs are replicated to the replication destination. Make sure to select all VMs that are part of the datastore. If they are not selected, only the VMs that are part of the replication plan are processed.

Replication Plan Details

Select Execution Order

Resource Group Name	Execution Order
DemoRG1	3

Network Mapping

No more Source/Destination network resources available for mapping

Source Resource	Destination Resource	
VLAN 3375	sddc-cgw-network-1	Delete

DataStore Mapping

Source DataStore	Destination Volume
DRO_Mini	DRO_Mini_copy

Previous Continue

- Under the VM details, you can optionally resize the VM's CPU and RAM parameters; this can be very helpful when recovering large environments to smaller target clusters or for conducting DR tests without having to provision a one-to-one physical VMware infrastructure. Also, you can modify the boot order and boot delay (seconds) for all the selected VMs across the resource groups. There is an additional option to modify the boot order if there are any changes required from those selected during the resource-group boot-order selection. By default, the boot order selected during resource-group selection is used; however, any modifications can be performed at this stage.

9. Click **Create Replication Plan**.

After the replication plan is created, the failover option, the test-failover option, or the migrate option can be exercised depending on the requirements. During the failover and test-failover options, the most recent SnapMirror Snapshot copy is used, or a specific Snapshot copy can be selected from a point-in-time Snapshot copy (per the retention policy of SnapMirror). The point-in-time option can be very helpful if you are facing a corruption event like ransomware, where the most recent replicas are already compromised or encrypted. DRO shows all available points in time. To trigger failover or test failover with the configuration specified in the replication plan, you can click **Failover** or **Test failover**.

NetApp Disaster Recovery Orchestrator

Dashboard | Discover | Resource Groups | Replication Plans | Job Monitoring | ? ! ? !

Source Details: 1 Site, 1 vCenter

Destination Details: 1 Site, 1 vCenter

2 Replication Plans

Plan Name	Active Site	Status	Compliance	Source Site	Destination Site	Actions
DemoRP	Source	Active	Healthy	On Prem	Cloud	Resource Groups ...
DemoRP	Source	Active	Healthy	On Prem	Cloud	Resource

Plan Details

Edit Plan

Failover

Test Failover

Migrate

Run Compliance

Delete Plan

Failover Details

Volume Snapshot Details

- Use latest snapshot ⓘ
- Select specific snapshot ⓘ

Start Failover

The replication plan can be monitored in the task menu:

Job Monitoring

Back

Failover Steps

Replication Plan: DemoRP

- Breaking SnapMirror relationships (in parallel) Success 11.3 Seconds
- Mounting volumes and creating datastores (in parallel) Success 34.7 Seconds
- Registering VMs (in parallel) Success 13.2 Seconds
- Powering on VMs in protection group - DemoRG1 - in target Success 95.8 Seconds
- Updating replication status Success 0.5 Seconds

After failover is triggered, the recovered items can be seen in the VMC vCenter (VMs, networks, datastores). By default, the VMs are recovered to the Workload folder.

Disaster Recovery Orchestrator

Dashboard | Discover | Resource Groups | Replication Plans | Job Monitoring

Protected VMs

vCenter Summary

Replication Camera

Application Plans

Failback can be triggered at the replication-plan level. For a test failover, the tear-down option can be used to roll back the changes and remove the FlexClone relationship. Failback related to failover is a two-step process. Select the replication plan and select **Reverse data sync**.

NetApp Disaster Recovery Orchestrator

Source Details: 1 Site, 1 vCenter

Destination Details: 1 Site, 1 vCenter

2 Replication Plans

Plan Name	Active Site	Status	Compliance	Source Site	Destination Site
DemoRP	Destination	Running In Failover	Healthy	On Prem	Cloud
DemoRP	Source	Active	Healthy	On Prem	Cloud

Resource Groups: Plan Details, Reverse Data Sync, Fallback

Back

Reverse Data Sync Steps

Replication Plan: DemoRP

- Powering off VMs in protection group - DemoRG1 - in source (In progress)
- Reversing SnapMirror relationships (in parallel) (Initialized)

Once completed, you can trigger failback to move back to original production site.

NetApp Disaster Recovery Orchestrator

Source Details: 1 Site, 1 vCenter

Destination Details: 1 Site, 1 vCenter

2 Replication Plans

Plan Name	Active Site	Status	Compliance	Source Site	Destination Site
DemoRP	Destination	Active	Healthy	On Prem	Cloud
DemoRP	Source	Active	Healthy	On Prem	Cloud

Resource Groups: Plan Details, Failback

NetApp Disaster Recovery Orchestrator

Back

Fallback Steps
Replication Plan: DemoRP

Powering off VMs in protection group - DemoRG1 - in target	In progress	- ⓘ
Unregistering VMs in target (in parallel)	Initialized	- ⓘ
Unmounting volumes in target (in parallel)	Initialized	- ⓘ
Breaking reverse SnapMirror relationships (in parallel)	Initialized	- ⓘ
Updating VM networks (in parallel)	Initialized	- ⓘ
Powering on VMs in protection group - DemoRG1 - in source	Initialized	- ⓘ
Deleting reverse SnapMirror relationships (in parallel)	Initialized	- ⓘ
Resuming SnapMirror relationships to target (in parallel)	Initialized	- ⓘ

From NetApp BlueXP, we can see that replication health has broken off for the appropriate volumes (those that were mapped to VMC as read-write volumes). During test failover, DRO does not map the destination or replica volume. Instead, it makes a FlexClone copy of the required SnapMirror (or Snapshot) instance and exposes the FlexClone instance, which does not consume additional physical capacity for FSx ONTAP. This process makes sure that the volume is not modified and replica jobs can continue even during DR tests or triage workflows. Additionally, this process makes sure that, if errors occur or corrupted data is recovered, the recovery can be cleaned up without the risk of the replica being destroyed.

NetApp Disaster Recovery Orchestrator

Protected VMs
3 Protected, 216 Unprotected

Environments
2 Virtual Environments, 2 Storage Environments

vCenter Summary
3 Clusters, 22 Folders, 23 Datastores, 45 Networks

Topology Canvas
On Prem: 172.21.253.160, 172.21.254.210 → Cloud: 44.235.223.88, 10.49.0.191

Execution Jobs
3 Total Jobs, 0 In Progress

Replication Plans

Replication Plan	Active Site	Status
DemoRP	Source	Active

Ransomware recovery

Recovering from ransomware can be a daunting task. Specifically, it can be hard for IT organizations to pinpoint where the safe point of return is and, once that is determined, to protect recovered workloads from reoccurring attacks from, for example, sleeping malware or vulnerable applications.

DRO addresses these concerns by enabling you to recover your system from any available point in time. You can also recover workloads to functional and yet isolated networks so that applications can function and communicate with each other in a location where they are not exposed to north-south traffic. This gives your security team a safe place to conduct forensics and make sure there is no hidden or sleeping malware.

Benefits

- Use of the efficient and resilient SnapMirror replication.
- Recovery to any available point in time with Snapshot copy retention.
- Full automation of all required steps to recover hundreds to thousands of VMs from the storage, compute, network, and application validation steps.
- Workload recovery with ONTAP FlexClone technology using a method that doesn't change the replicated volume.
 - Avoids risk of data corruption for volumes or Snapshot copies.
 - Avoids replication interruptions during DR test workflows.
 - Potential use of DR data with cloud computing resources for workflows beyond DR such as DevTest, security testing, patch or upgrade testing, and remediation testing.
- CPU and RAM optimization to help lower cloud costs by allowing recovery to smaller compute clusters.

Copyright information

Copyright © 2026 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.