



# **TR-4955: Disaster Recovery with Azure NetApp Files (ANF) and Azure VMware Solution (AVS)**

NetApp public and hybrid cloud solutions

NetApp  
July 30, 2025

# Table of Contents

- TR-4955: Disaster Recovery with Azure NetApp Files (ANF) and Azure VMware Solution (AVS) . . . . . 1
  - Overview . . . . . 1
    - Prerequisites and general recommendations . . . . . 1
  - Getting started . . . . . 2
    - Deploy Azure VMware Solution . . . . . 2
    - Provision and configure Azure NetApp Files . . . . . 2
  - DRO installation . . . . . 3
  - DRO configuration . . . . . 4
    - Resource groupings . . . . . 7
    - Replication plans . . . . . 8
    - Ransomware recovery . . . . . 14
  - Conclusion . . . . . 14
    - Where to find additional information . . . . . 14

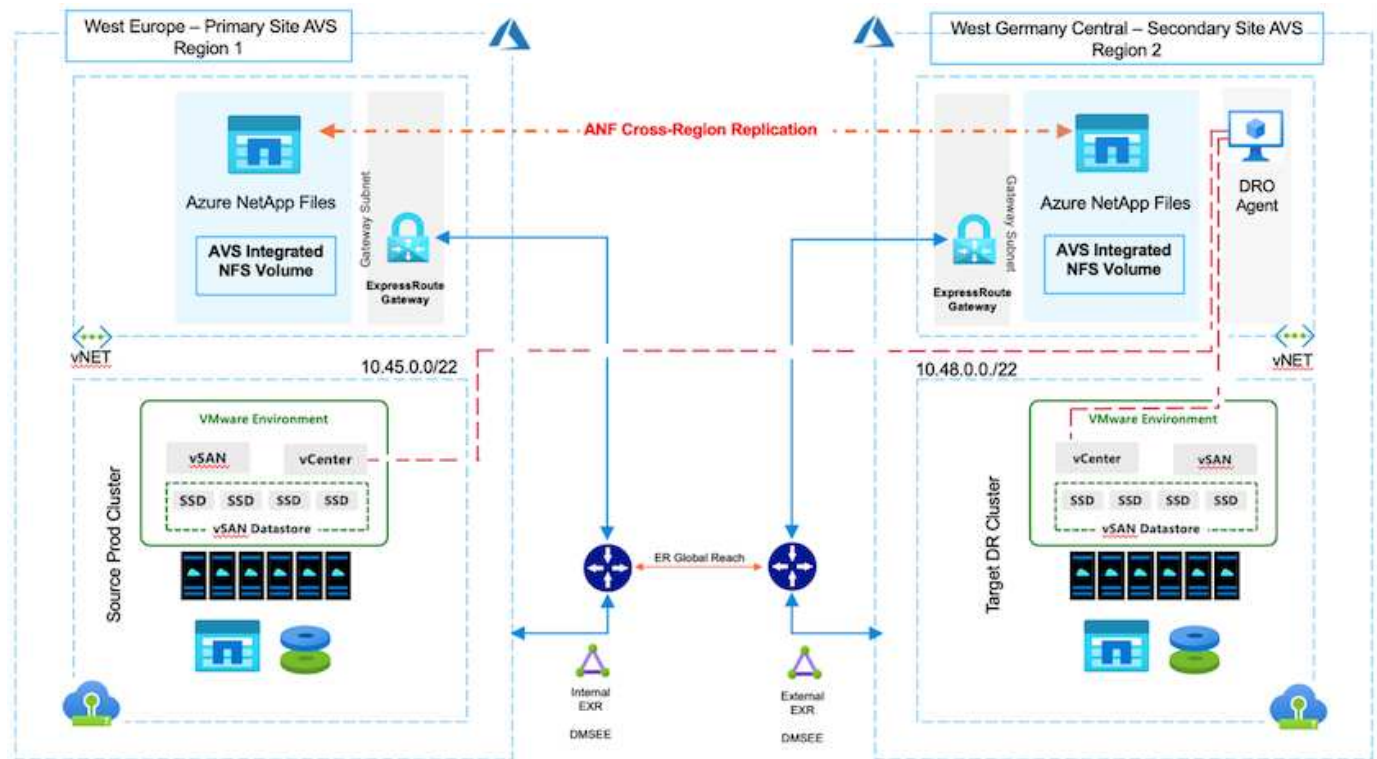
# TR-4955: Disaster Recovery with Azure NetApp Files (ANF) and Azure VMware Solution (AVS)

Disaster recovery using block-level replication between regions within the cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events (for example, ransomware).

## Overview

With Azure NetApp files (ANF) cross-region volume replication, VMware workloads running on an Azure VMware Solution (AVS) SDDC site using Azure NetApp files volumes as an NFS datastore on the primary AVS site can be replicated to a designated secondary AVS site in the target recovery region.

Disaster Recovery Orchestrator (DRO) (a scripted solution with a UI) can be used to seamlessly recover workloads replicated from one AVS SDDC to another. DRO automates recovery by breaking replication peering and then mounting the destination volume as a datastore, through VM registration to AVS, to network mappings directly on NSX-T (included with all AVS private clouds).



## Prerequisites and general recommendations

- Verify that you have enabled cross-region replication by creating replication peering. See [Create volume replication for Azure NetApp Files](#).
- You must configure ExpressRoute Global Reach between the source and target Azure VMware Solution private clouds.
- You must have a service principal that can access resources.
- The following topology is supported: primary AVS site to secondary AVS site.

- Configure the [replication](#) schedule for each volume appropriately based on business needs and the data-change rate.



Cascading and fan- in and fan- out topologies are not supported.

## Getting started

### Deploy Azure VMware Solution

The [Azure VMware Solution](#) (AVS) is a hybrid cloud service that provides fully functional VMware SDDCs within a Microsoft Azure public cloud. AVS is a first-party solution fully managed and supported by Microsoft and verified by VMware that uses Azure infrastructure. Therefore, customers get VMware ESXi for compute virtualization, vSAN for hyper-converged storage, and NSX for networking and security, all while taking advantage of Microsoft Azure's global presence, class-leading data- center facilities, and proximity to the rich ecosystem of native Azure services and solutions. A combination of Azure VMware Solution SDDC and Azure NetApp Files provides the best performance with minimal network latency.

To configure an AVS private cloud on Azure, follow the steps in this [link](#) for NetApp documentation and in this [link](#) for Microsoft documentation. A pilot- light environment set up with a minimal configuration can be used for DR purposes. This setup only contains core components to support critical applications, and it can scale out and spawn more hosts to take the bulk of the load if a failover occurs.




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

### Provision and configure Azure NetApp Files

[Azure NetApp Files](#) is a high-performance, enterprise-class, metered file- storage service. Follow the steps in this [link](#) to provision and configure Azure NetApp Files as a NFS datastore to optimize AVS private cloud deployments.

#### Create volume replication for Azure NetApp Files-powered datastore volumes

The first step is to set up cross- region replication for the desired datastore volumes from the AVS primary site to the AVS secondary site with the appropriate frequencies and retentions.



Follow the steps in this [link](#) to set up cross-region replication by creating replication peering. The service level for the destination capacity pool can match that of the source capacity pool. However, for this specific use case, you can select the standard service level and then [modify the service level](#) in the event of a real disaster or DR simulations.



A cross- region replication relationship is a prerequisite and must be created beforehand.

# DRO installation

To get started with DRO, use the Ubuntu operating system on the designated Azure virtual machine and make sure you meet the prerequisites. Then install the package.

## Prerequisites:

- Service principal that can access resources.
- Make sure that appropriate connectivity exists to the source and destination SDDC and Azure NetApp Files instances.
- DNS resolution should be in place if you are using DNS names. Otherwise, use IP addresses for vCenter.

## OS requirements:

- Ubuntu Focal 20.04 (LTS)The following packages must be installed on the designated agent virtual machine:
- Docker
- Docker- compose
- JqChange `docker.sock` to this new permission: `sudo chmod 666 /var/run/docker.sock`.



The `deploy.sh` script executes all required prerequisites.

The steps are as follows:

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

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




The agent must be installed in the secondary AVS site region or in the primary AVS site region in a separate AZ than the SDDC.

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

```
tar xvf draas_package.tar
Navigate to the directory and run the deploy script as below:
sudo sh deploy.sh
```

3. Access the UI using the following credentials:

- Username: admin
- Password: admin



## DRO configuration

After Azure NetApp Files and AVS have been configured properly, you can begin configuring DRO to automate the recovery of workloads from the primary AVS site to the secondary AVS site. NetApp recommends deploying the DRO agent in the secondary AVS site and configuring the ExpressRoute gateway connection so that the DRO agent can communicate via the network with the appropriate AVS and Azure NetApp Files components.

The first step is to Add credentials. DRO requires permission to discover Azure NetApp Files and the Azure VMware Solution. You can grant the required permissions to an Azure account by creating and setting up an Azure Active Directory (AD) application and by obtaining the Azure credentials that DRO needs. You must bind the service principal to your Azure subscription and assign it a custom role that has the relevant required permissions. When you add source and destination environments, you are prompted to select the credentials associated with the service principal. You need to add these credentials to DRO before you can click Add New Site.

To perform this operation, complete the following steps:

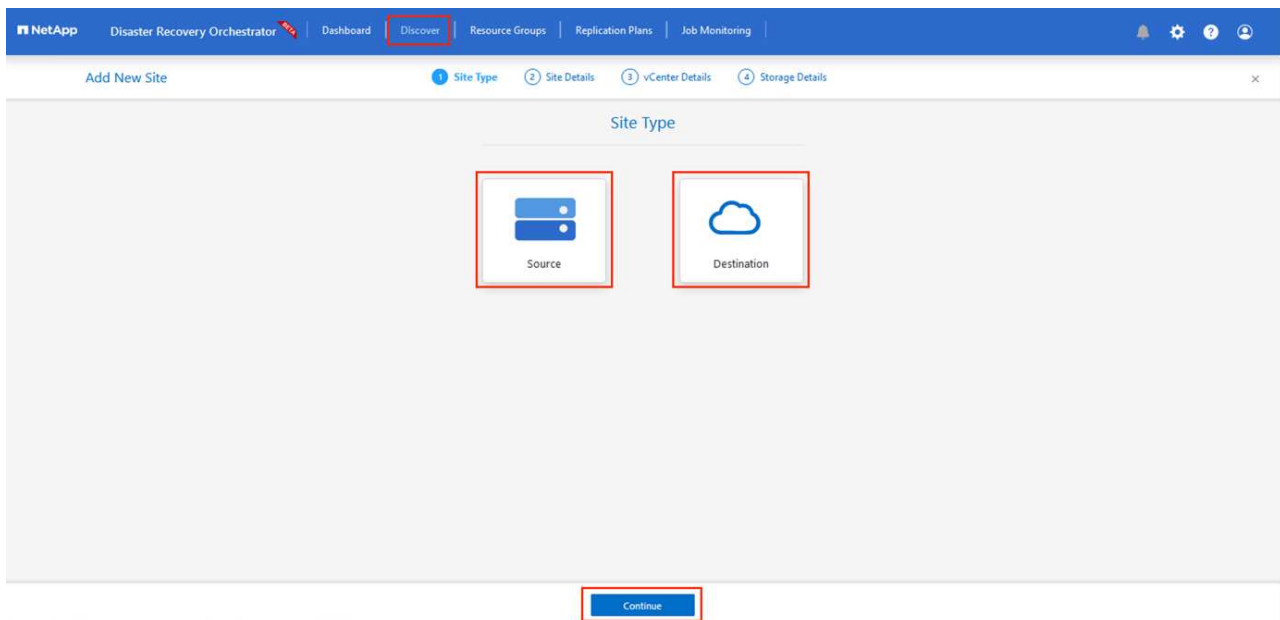
1. Open DRO in a supported browser and use the default username and password (admin/admin). The password can be reset after the first login using the Change Password option.
2. In the upper right of the DRO console, click the **Settings** icon, and select **Credentials**.
3. Click Add New Credential and follow the steps in the wizard.
4. To define the credentials, enter information about the Azure Active Directory service principal that grants the required permissions:
  - Credential name
  - Tenant ID
  - Client ID
  - Client secret
  - Subscription ID

You should have captured this information when you created the AD application.

5. Confirm the details about the new credentials and click Add Credential.

After you add the credentials, it's time to discover and add the primary and secondary AVS sites (both vCenter and the Azure NetApp files storage account) to DRO. To add the source and destination site, complete the following steps:

6. Go to the **Discover** tab.
7. Click **Add New Site**.
8. Add the following primary AVS site (designated as **Source** in the console).
  - SDDC vCenter
  - Azure NetApp Files storage account
9. Add the following secondary AVS site (designated as **Destination** in the console).
  - SDDC vCenter
  - Azure NetApp Files storage account



10. Add site details by clicking **Source**, entering a friendly site name, and select the connector. Then click **Continue**.



For demonstration purposes, adding a source site is covered in this document.

11. Update the vCenter details. To do this, select the credentials, Azure region, and resource group from the dropdown for the primary AVS SDDC.
12. DRO lists all the available SDDCs within the region. Select the designated private cloud URL from the dropdown.
13. Enter the `cloudadmin@vsphere.local` user credentials. This can be accessed from Azure Portal. Follow the steps mentioned in this [link](#). Once done, click **Continue**.

14. Select the Source Storage details (ANF) by selecting the Azure Resource group and NetApp account.
15. Click **Create Site**.



The screenshot shows the NetApp Disaster Recovery Orchestrator (DRO) Dashboard. The top navigation bar includes links for Dashboard, Discover, Resource Groups, Replication Plans, and Job Monitoring. The main content area displays site configuration metrics: 2 Sites, 2 vCenters, and 2 Storages. Below these, a table lists the configured sites:

Site Name	Site Type	Location	vCenter	Storage	VM List	Discovery Status
DemoDest	Destination	Cloud	1	1	https://10.75.0.2/	Success
DemoSRC	Source	Cloud	1	1	View VM List   https://172.30.156.2/	Success

Once added, DRO performs automatic discovery and displays the VMs that have corresponding cross- region replicas from the source site to the destination site. DRO automatically detects the networks and segments used by the VMs and populates them.

The screenshot shows the NetApp Disaster Recovery Orchestrator (DRO) VM List page for the DemoSRC site. The page displays 7 Datastores and 128 Virtual Machines. A table lists the VMs with their status and protection details:

VM Name	VM Status	VM State	DataStore	CPU	Memory (MB)
HCIbench_2.5.1	Not Protected	Powered On	vsanDatastore	8	8192
hci-fio-datastore-13984-0-1	Not Protected	Powered Off	HCIstDS	32	65536
ICCAz005-WD-R1	Not Protected	Powered On	vsanDatastore	8	14336
ICCAz005-FIE-R1	Not Protected	Powered On	vsanDatastore	8	3072
ICCAz005-IX-R1	Not Protected	Powered On	vsanDatastore	8	3072
HCI_Demo_OS	Not Protected	Powered Off	Demo002	1	2048
hci-nim-datastore-13984-0-1	Not Protected	Powered Off	HCIstDS	24	49152


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

## Resource groupings

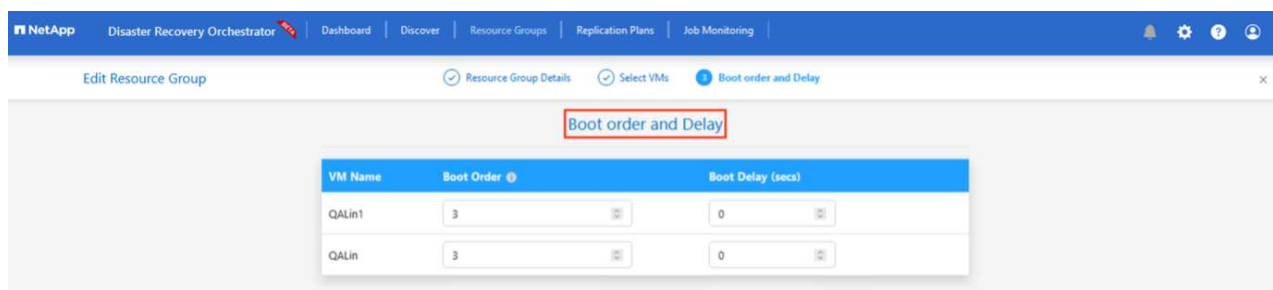
After the platforms have been added, group the VMs you want to recover into resource groups. DRO resource groups allow you to group a set of dependent VMs into logical groups that contain their boot orders, boot delays, and optional application validations that can be executed upon recovery.

To start creating resource groups, click the **Create New Resource Group** menu item.

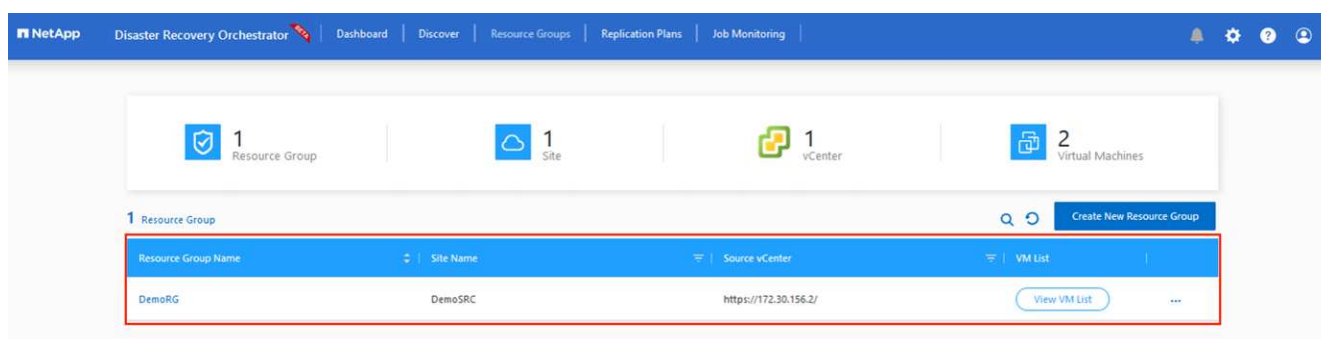
1. Access **Resource Groups** and click **Create New Resource Group**.



2. Under New Resource Group, select the source site from the dropdown and click **Create**.
3. Provide the resource group details and click **Continue**.
4. Select appropriate VMs using the search option.
5. Select the **Boot Order** and **Boot Delay** (secs) for all the selected VMs. Set the order of the power-on sequence by selecting each virtual machine and setting up the priority for it. The default value for all virtual machines is 3. The options are as follows:
  - The first virtual machine to power on
  - Default
  - The last virtual machine to power on



6. Click **Create Resource Group**.

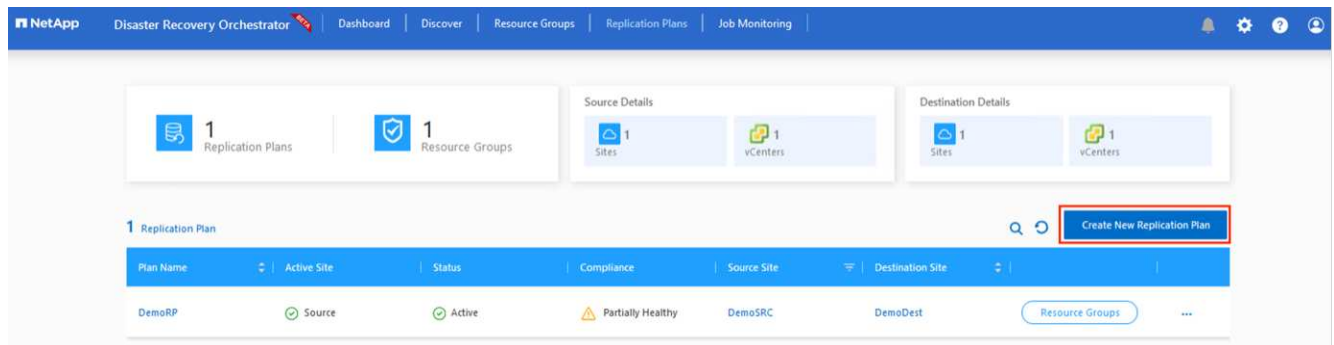


## Replication plans

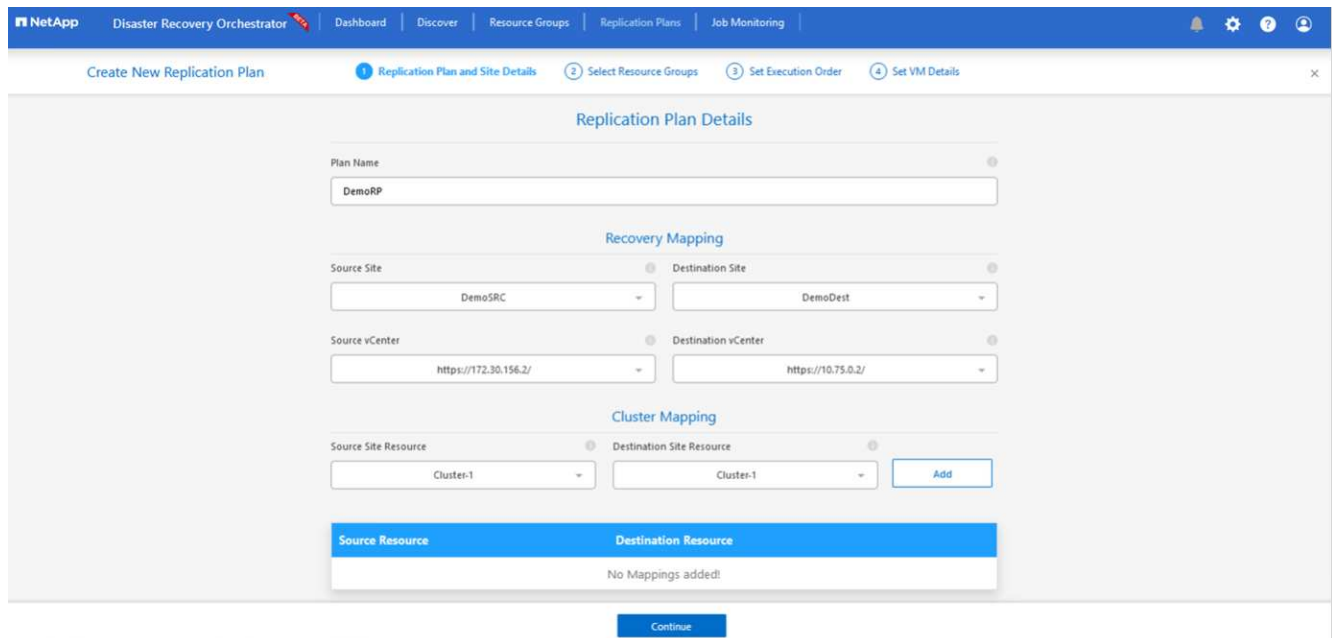
You must have a plan to recover applications in the event of a disaster. Select the source and destination vCenter platforms from the drop down, pick the resource groups to be included in this plan, and also include the grouping of how applications should be restored and powered on (for example, domain controllers, tier-1, tier-2, and so on). Plans are often called blueprints as well. 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. Navigate to **Replication Plans** and click **Create New Replication Plan**.



2. On the **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.



3. After recovery mapping is complete, select the **Cluster Mapping**.

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

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: DemoSRC | Destination Site: DemoDest

Source vCenter: https://172.30.156.2/ | Destination vCenter: https://10.75.0.2/

#### Cluster Mapping

No more Source/Destination cluster resources available for mapping

Source Resource	Destination Resource	
Cluster-1	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. Once done, set network mapping to the appropriate segment. The segments should already be provisioned on the secondary AVS cluster, and, to map the VMs to those, select the appropriate segment.
7. Datastore mappings are automatically selected based on the selection of VMs.



Cross- region replication (CRR) is at the volume level. Therefore, all VMs residing on the respective volume are replicated to the CRR destination. Make sure to select all VMs that are part of the datastore, because only virtual machines that are part of the replication plan are processed.

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

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

#### Select Execution Order

Resource Group Name	Execution Order
DemoRG	3

#### Network Mapping

No more Source/Destination network resources available for mapping

Source Resource	Destination Resource	
SepSeg	SegDR	Delete

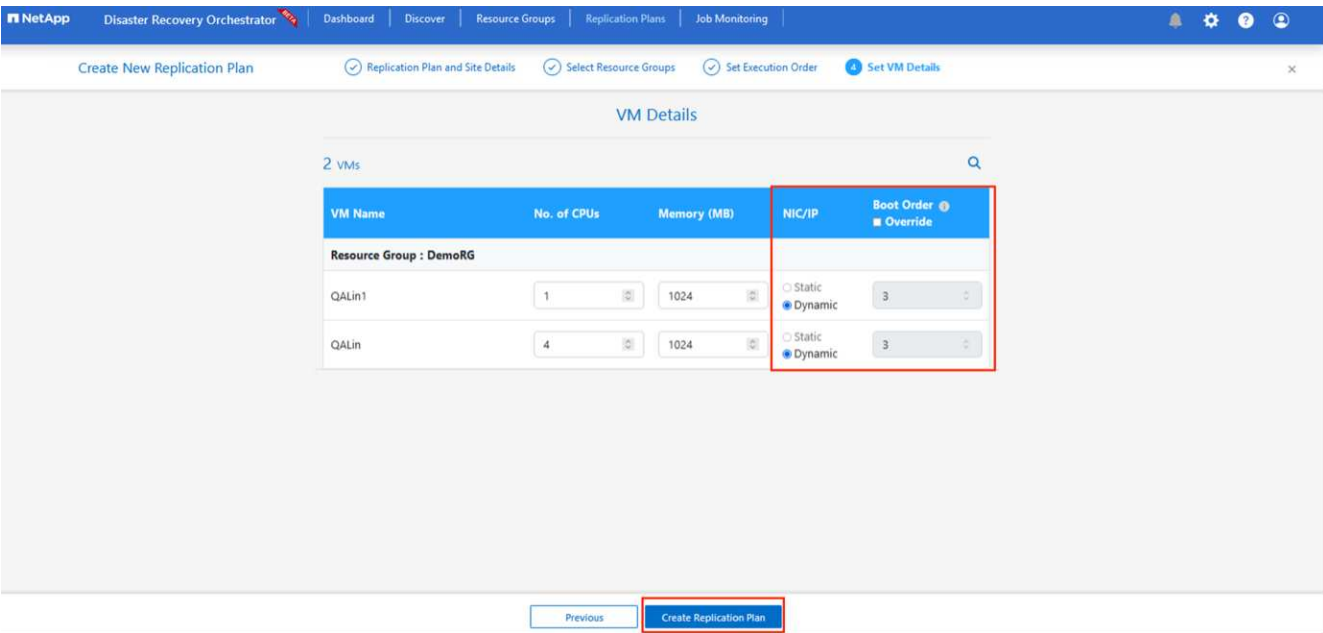
#### DataStore Mapping

Source DataStore	Destination Volume
TestSrc01	gwc_ntap_acct/gwc_DRO_cp/testsrc01copy

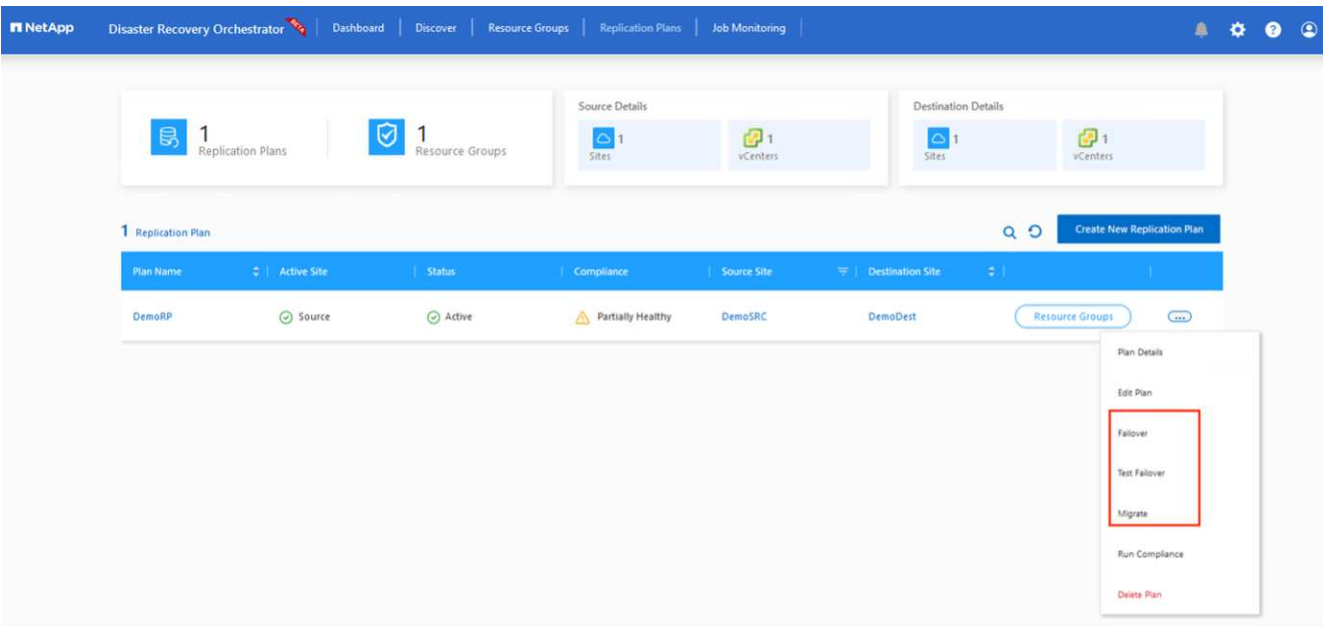
Previous | Continue

8. Under VM details, you can optionally resize the VMs CPU and RAM parameters. This can be very helpful

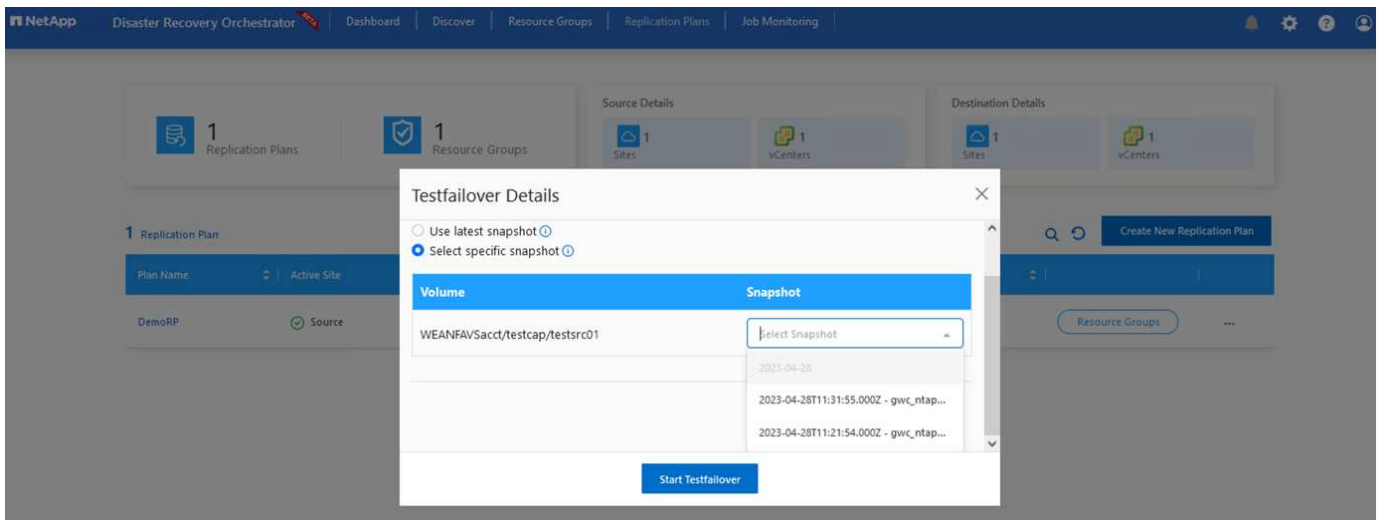
when you are recovering large environments to smaller target clusters or when you are conducting DR tests without having to provision a one-to-one physical VMware infrastructure. Also, modify the boot order and boot delay (secs) for all the selected VMs across the resource groups. There is an additional option to modify the boot order if any changes are required from what you selected during 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.



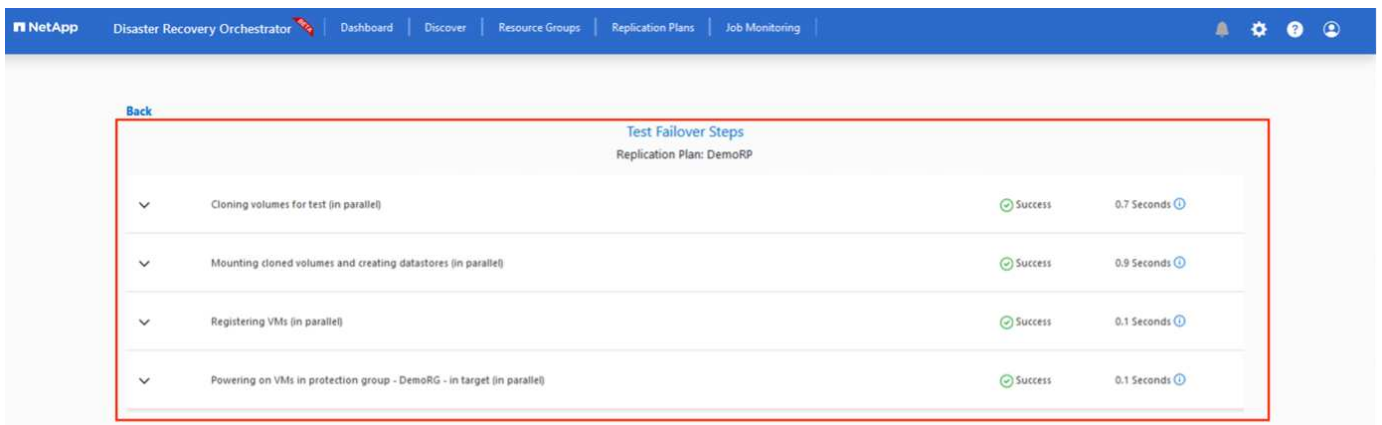
9. Click **Create Replication Plan**.After the replication plan is created, you can exercise the failover, test failover, or migrate options depending on your requirements.



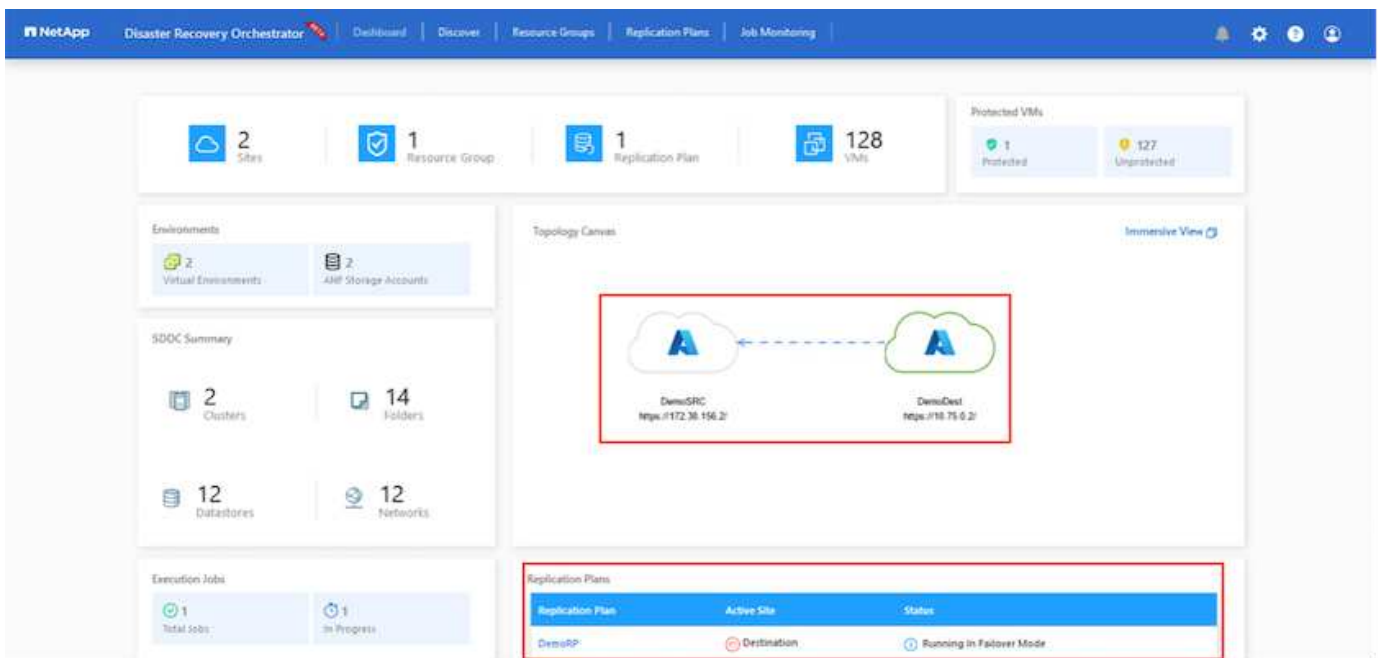
During the failover and test failover options, the most recent snapshot is used, or a specific snapshot can be selected from a point-in-time snapshot. The point-in-time option can be very beneficial if you are facing a corruption event like ransomware, where the most recent replicas are already compromised or encrypted. DRO shows all available time points.



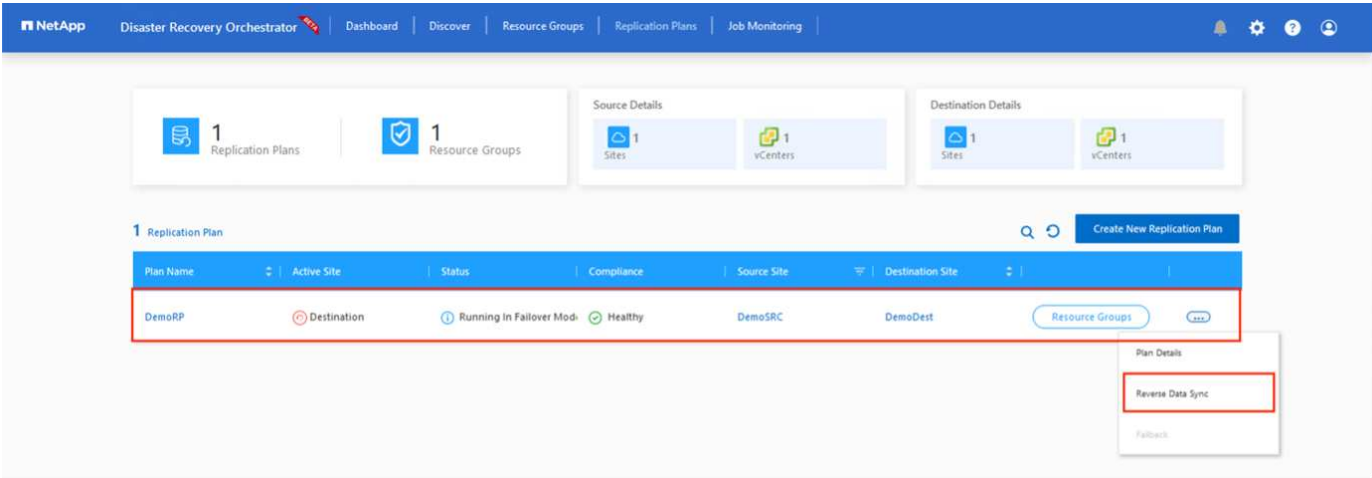
To trigger failover or test failover with the configuration specified in the replication plan, you can click **Failover** or **Test Failover**. You can monitor the replication plan in the task menu.



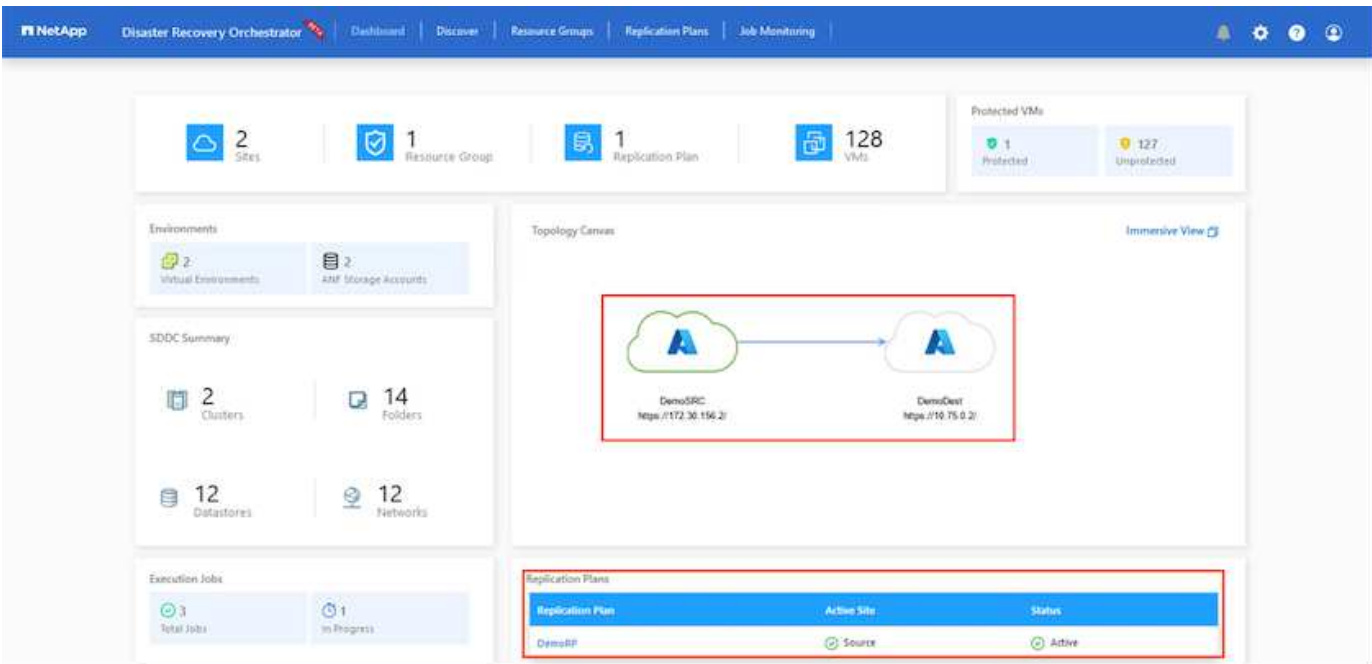
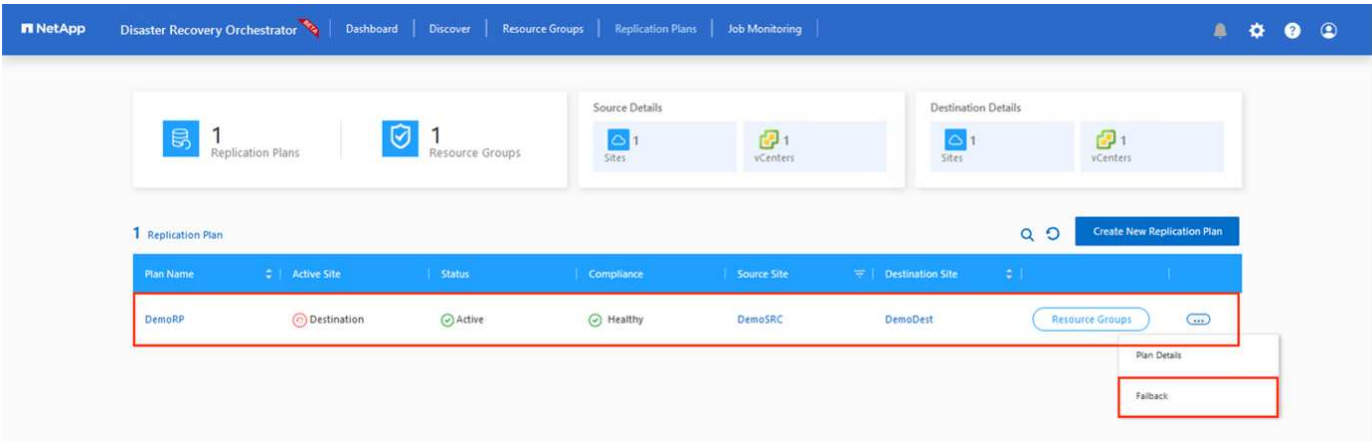
After failover is triggered, the recovered items can be seen in the secondary site AVS SDDC vCenter (VMs, networks, and datastores). By default, the VMs are recovered to Workload folder.



Failback can be triggered at the replication plan level. In case of test failover, the tear down option can be used to roll back the changes and remove the newly created volume. Failbacks related to failover are a two- step process. Select the replication plan and select **Reverse Data sync**.



After this step is complete, trigger failback to move back to the primary AVS site.





From the Azure portal, we can see that the replication health has been broken off for the appropriate volumes that were mapped to the secondary site AVS SDDC as read/write volumes. During test failover, DRO does not map the destination or replica volume. Instead, it creates a new volume of the required cross- region replication snapshot and exposes the volume as a datastore, which consumes additional physical capacity from the capacity pool and ensures that the source volume is not modified. Notably, replication jobs can continue during DR tests or triage workflows. Additionally, this process makes sure that the recovery can be cleaned up without the risk of the replica being destroyed if errors occur or corrupted data is recovered.

## Ransomware recovery

Recovering from ransomware can be a daunting task. Specifically, it can be difficult for IT organizations to pinpoint what the safe point of return is, and, once that's determined, how to ensure that recovered workloads are safeguarded from the attacks reoccurring (for example, from sleeping malware or through vulnerable applications).

DRO addresses these concerns by allowing organizations to recover from any available point-in-time. Workloads are then recovered to functional and yet isolated networks, so that applications can function and communicate with each other but are not exposed to any north- south traffic. This process gives security teams a safe place to conduct forensics and identify any hidden or sleeping malware.

## Conclusion

The Azure NetApp Files and Azure VMware disaster recovery solution provide you with the following benefits:

- Leverage efficient and resilient Azure NetApp Files cross- region replication.
- Recover to any available point-in-time with snapshot retention.
- Fully automate all required steps to recover hundreds to thousands of VMs from the storage, compute, network, and application validation steps.
- Workload recovery leverages the "Create new volumes from the most recent snapshots" process, which doesn't manipulate the replicated volume.
- Avoid any risk of data corruption on the volumes or snapshots.
- Avoid replication interruptions during DR test workflows.
- Leverage DR data and cloud compute resources for workflows beyond DR, such as dev/test, security testing, patch and upgrade testing, and remediation testing.
- CPU and RAM optimization can help lower cloud costs by allowing recovery to smaller compute clusters.

## Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- Create volume replication for Azure NetApp Files

<https://learn.microsoft.com/en-us/azure/azure-netapp-files/cross-region-replication-create-peering>

- Cross-region replication of Azure NetApp Files volumes

<https://learn.microsoft.com/en-us/azure/azure-netapp-files/cross-region-replication-introduction#service-level-objectives>



- Azure VMware Solution

<https://learn.microsoft.com/en-us/azure/azure-vmware/introduction>

- Deploy and configure the Virtualization Environment on Azure

[Setup AVS on Azure](#)

- Deploy and configure Azure VMware Solution

<https://learn.microsoft.com/en-us/azure/azure-vmware/deploy-azure-vmware-solution?tabs=azure-portal>

## Copyright information

Copyright © 2025 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.