



NetApp Disaster Recovery documentation

NetApp Disaster Recovery

NetApp
October 06, 2025

This PDF was generated from <https://docs.netapp.com/us-en/data-services-disaster-recovery/index.html> on October 06, 2025. Always check docs.netapp.com for the latest.

Table of Contents

NetApp Disaster Recovery documentation	1
Release notes	2
What's new in NetApp Disaster Recovery	2
06 October 2025	2
04 August 2025	2
14 July 2025	3
30 June 2025	4
23 June 2025	4
09 June 2025	4
13 May 2025	4
16 April 2025	6
10 March 2025	6
19 February 2025	7
30 October 2024	7
20 September 2024	9
02 August 2024	9
17 July 2024	10
05 July 2024	10
15 May 2024	11
05 March 2024	12
01 February 2024	12
11 January 2024	13
20 October 2023	13
27 September 2023	14
01 August 2023	14
18 May 2023	15
Limitations in NetApp Disaster Recovery	15
Wait until failback completes before running discovery	16
NetApp Console might not discover Amazon FSx for NetApp ONTAP	16
Get started	17
Learn about NetApp Disaster Recovery for VMware	17
NetApp Console	18
Benefits of using NetApp Disaster Recovery for VMware	18
What you can do with NetApp Disaster Recovery for VMware	19
Cost	20
Licensing	20
30-day free trial	21
How NetApp Disaster Recovery works	21
Supported protection targets and datastore types	23
Terms that might help you with NetApp Disaster Recovery	24
NetApp Disaster Recovery prerequisites	24
Software versions	24
ONTAP storage prerequisites	24

VMware vCenter clusters prerequisites	25
NetApp Console prerequisites	25
Workload prerequisites	26
Quick start for NetApp Disaster Recovery	26
Set up your infrastructure for NetApp Disaster Recovery	27
Hybrid cloud with VMware Cloud and Amazon FSx for NetApp ONTAP	27
Private cloud	29
Access NetApp Disaster Recovery	30
Set up licensing for NetApp Disaster Recovery	32
Try it out using a 30-day free trial.	33
After the trial ends, subscribe through one of the Marketplaces	34
After the trial ends, purchase a BYOL license through NetApp	35
Update your license when it expires	35
End the free trial.	35
Use NetApp Disaster Recovery	37
Use NetApp Disaster Recovery overview	37
View the health of your NetApp Disaster Recovery plans on the Dashboard	37
Add vCenters to a site in NetApp Disaster Recovery.	38
Add subnet mapping for a vCenter site	41
Edit the vCenter server site and customize the discovery schedule	44
Refresh discovery manually	45
Create a resource group to organize VMs together in NetApp Disaster Recovery	46
Create a replication plan in NetApp Disaster Recovery	49
Create the plan.	50
Edit schedules to test compliance and ensure failover tests work.	62
Replicate applications to another site with NetApp Disaster Recovery	64
Migrate applications to another site with NetApp Disaster Recovery	65
Fail over applications to a remote site with NetApp Disaster Recovery.	66
Test the failover process	66
Clean up the test environment after a failover test	67
Fail over the source site to a disaster recovery site.	67
Fail back applications to the original source with NetApp Disaster Recovery	69
Manage sites, resource groups, replication plans, datastores and virtual machines information with NetApp Disaster Recovery	70
Manage vCenter sites	70
Manage resource groups	70
Manage replication plans	71
View datastores information.	73
View virtual machines information	74
Monitor NetApp Disaster Recovery jobs	74
View jobs	74
Cancel a job	74
Create NetApp Disaster Recovery reports.	75
Reference	76
vCenter privileges needed for NetApp Disaster Recovery.	76

NetApp Disaster Recovery role-based access to features	77
Use NetApp Disaster Recovery with Amazon EVS	79
Introduction of NetApp Disaster Recovery using Amazon Elastic VMware Service and Amazon FSx for NetApp ONTAP	79
Solution overview of NetApp Disaster Recovery using Amazon EVS and Amazon FSs for NetApp ONTAP	79
Install the NetApp Console agent for NetApp Disaster Recovery	81
Configure NetApp Disaster Recovery for Amazon EVS	81
Create replication plans for Amazon EVS	93
Perform replication plan operations with NetApp Disaster Recovery	101
Frequently asked questions for NetApp Disaster Recovery	114
Knowledge and support	115
Register for support	115
Support registration overview	115
Register BlueXP for NetApp support	115
Associate NSS credentials for Cloud Volumes ONTAP support	117
Get help	119
Get support for a cloud provider file service	119
Use self-support options	119
Create a case with NetApp support	119
Manage your support cases (Preview)	122
Legal notices	125
Copyright	125
Trademarks	125
Patents	125
Privacy policy	125
Open source	125

NetApp Disaster Recovery documentation

Release notes

What's new in NetApp Disaster Recovery

Learn what's new in NetApp Disaster Recovery.

06 October 2025

BlueXP disaster recovery is now NetApp Disaster Recovery

BlueXP disaster recovery has been renamed to NetApp Disaster Recovery.

BlueXP is now NetApp Console

The NetApp Console, built on the enhanced and restructured BlueXP foundation, provides centralized management of NetApp storage and NetApp Data Services across on-premises and cloud environments at enterprise grade—delivering real-time insights, faster workflows, and simplified administration, that is highly secure and compliant.

For details on what has changed, see the [NetApp Console release notes](#).

Other updates

- Support for Amazon Elastic VMware Service (EVS) with Amazon FSx for NetApp ONTAP was in a public preview. With this release, it is now generally available. For details, refer to [Introduction of NetApp Disaster Recovery using Amazon Elastic VMware Service and Amazon FSx for NetApp ONTAP](#).
- Storage discovery improvements, including reduced discovery times for on-premises deployments
- Identity and Access Management (IAM) support, including role-based access control (RBAC) and enhanced user permissions
- Private Preview support for Azure VMware solution and Cloud Volumes ONTAP. With this support, you can now configure disaster recovery protection from on-premises to the Azure VMware solution using Cloud Volumes ONTAP storage.

04 August 2025

Version 4.2.5P2

NetApp Disaster Recovery updates

This release includes the following updates:

- Improved the VMFS support to handle the same LUN presented from multiple storage virtual machines.
- Improved the test teardown cleanup to handle the datastore already being unmounted and or deleted.
- Improved subnet mapping so that it now validates that the gateway entered is contained within the network provided.
- Corrected an issue that could cause the replication plan to fail if the VM name contains ".com".
- Removed a restriction preventing the destination volume from being the same as the source volume when creating the volume as part of the replication plan creation.

- Added support for a pay-as-you-go (PAYGO) subscription to NetApp Intelligent Services in the Azure Marketplace and added a link to the Azure Marketplace in the free trial dialog.

For details, see [NetApp Disaster Recovery licensing](#) and [Set up licensing for NetApp Disaster Recovery](#).

14 July 2025

Version 4.2.5

User roles in NetApp Disaster Recovery

NetApp Disaster Recovery now employs roles to govern the access that each user has to specific features and actions.

The service uses the following roles that are specific to NetApp Disaster Recovery.

- **Disaster recovery admin:** Perform any actions in NetApp Disaster Recovery.
- **Disaster recovery failover admin:** Perform failover and migrate actions in NetApp Disaster Recovery.
- **Disaster recovery application admin:** Create and modify replication plans and start test failovers.
- **Disaster recovery viewer:** View information in NetApp Disaster Recovery, but cannot perform any actions.

If you are clicking on the NetApp Disaster Recovery service and configuring it for the first time, you must have the **SnapCenterAdmin** permission or have the **Organization Admin** role.

For details, see [User roles and permissions in NetApp Disaster Recovery](#).

[Learn about access roles for all services.](#)

Other updates in NetApp Disaster Recovery

- Enhanced network discovery
- Scalability improvements:
 - Filtering for the required metadata instead of all the details
 - Discovery improvements to retrieve and update VM resources faster
 - Memory optimization and performance optimization for data retrieval and data updates
 - vCenter SDK client creation and pool management improvements
- Stale data management on the next scheduled or manual discovery:
 - When a VM is deleted in the vCenter, NetApp Disaster Recovery now automatically removes it from the replication plan.
 - When a datastore or network is deleted in the vCenter, NetApp Disaster Recovery now deletes it from the replication plan and resource group.
 - When a cluster, host, or datacenter is deleted in the vCenter, NetApp Disaster Recovery now deletes it from the replication plan and resource group.
- You can now access Swagger documentation in your browser's incognito mode. You can access it from within NetApp Disaster Recovery from the Settings option > API Documentation or directly at the following URL in your browser's incognito mode: [Swagger documentation](#).
- In some situations after a failback operation, the iGroup was left behind after the operation completed. This update removes the iGroup if it is stale.

- If the NFS FQDN was used in the replication plan, NetApp Disaster Recovery now resolves it to an IP address. This update is useful if the FQDN is not resolvable in the disaster recovery site.
- UI alignment improvements
- Log improvements to capture the vCenter sizing details after the successful discovery

30 June 2025

Version 4.2.4P2

Discovery improvements

This update improves the discovery process, which reduces the time needed for discovery.

23 June 2025

Version 4.2.4P1

Subnet mapping improvements

This update enhances the Add and Edit Subnet Mapping dialog with a new search functionality. You can now quickly find specific subnets by entering search terms, making it easier to manage subnet mappings.

09 June 2025

Version 4.2.4

Windows Local Administrator Password Solution (LAPS) support

Windows Local Administrator Password Solution (Windows LAPS) is a Windows feature that automatically manages and backs up the password of a local administrator account on Active directory.

You can now select subnet mapping options and check the LAPS option by providing the domain controller details. Using this option, you don't need to provide a password for each of your virtual machines.

For details, refer to [Create a replication plan](#).

13 May 2025

Version 4.2.3

Subnet mapping

With this release, you can manage IP addresses on failover in a new way using subnet mapping, which enables you to add subnets for each vCenter. When you do so, you define the IPv4 CIDR, the default gateway, and the DNS for each virtual network.

Upon failover, NetApp Disaster Recovery determines the appropriate IP address of each vNIC by looking at the CIDR provided for the mapped virtual network and uses it to derive the new IP address.

For example:

- NetworkA = 10.1.1.0/24

- NetworkB = 192.168.1.0/24

VM1 has a vNIC (10.1.1.50) that is connected to NetworkA.
NetworkA is mapped to NetworkB in the replication plan settings.

Upon failover, NetApp Disaster Recovery replaces the Network portion of the original IP address (10.1.1) and keeps the host address (.50) of the original IP address (10.1.1.50). For VM1, NetApp Disaster Recovery looks at the CIDR settings for NetworkB and uses that the NetworkB network portion 192.168.1 while keeping the host portion (.50) to create the new IP address for VM1. The new IP becomes 192.168.1.50.

In summary, the host address stays the same, while the network address is replaced with whatever is configured in the site subnet mapping. This enables you to manage IP address reassignment upon failover more easily, especially if you have hundreds of networks and thousands of VMs to manage.

For details about including subnet mapping in your sites, refer to [Add vCenter server sites](#).

Skip protection

You can now skip protection so that the service does not automatically create a reverse protection relationship after a replication plan failover. This is useful if you want to perform additional operations on the restored site before you bring it back online within NetApp Disaster Recovery.

When you initiate a failover, by default the service automatically creates a reverse protection relationship for each volume in the replication plan, if the original source site is online. This means that the service creates a SnapMirror relationship from the target site back to the source site. The service also automatically reverses the SnapMirror relationship when you initiate a failback.

When initiating a failover, you can now choose a **Skip protection** option. With this, the service does not automatically reverse the SnapMirror relationship. Instead, it leaves the writable volume on both sides of the replication plan.

After the original source site is back online, you can establish reverse protection by selecting **Protect resources** from the Replication plan Actions menu. This attempts to create a reverse replication relationship for each volume in the plan. You can run this job repeatedly until protection is restored. When protection is restored, you can initiate a failback in the usual way.

For details skipping protection, refer to [Fail over applications to a remote site](#).

SnapMirror schedule updates in the replication plan

NetApp Disaster Recovery now supports the use of external snapshot management solutions such as the native ONTAP SnapMirror policy scheduler or third-party integrations with ONTAP. If every datastore (volume) in the replication plan already has a SnapMirror relationship that is being managed elsewhere, you can use those snapshots as recovery points in NetApp Disaster Recovery.

To configure, in the Replication plan > Resource mapping section, check the **Use platform managed backups and retention schedules** checkbox when configuring the Datastores mapping.

When the option is selected, NetApp Disaster Recovery does not configure a backup schedule. However, you still need to configure a retention schedule because snapshots might still be taken for testing, failover, and failback operations.

After this is configured, the service doesn't take any regularly scheduled snapshots, but instead relies on the external entity to take and update those snapshots.

For details about using external snapshot solutions in the replication plan, refer to [Create a replication plan](#).

16 April 2025

Version 4.2.2

Scheduled discovery for VMs

NetApp Disaster Recovery performs discovery once every 24 hours. With this release, you can now customize the discovery schedule to meet your needs and reduce impact on performance when you need it. For example, if you have a large number of VMs, you can set the discovery schedule to run every 48 hours. If you have a small number of VMs, you can set the discovery schedule to run every 12 hours.

If you don't want to schedule discovery, you can disable the scheduled discovery option and refresh the discovery manually at any time.

For details, refer to [Add vCenter server sites](#).

Resource group datastore support

Previously, you could create resource groups only by VMs. With this release, you can create a resource group by datastores. When you're creating a replication plan and creating a resource group for that plan, all the VMs in a datastore will be listed. This is useful if you have a large number of VMs and want to group them by datastore.

You can create a resource group with a datastore in the following ways:

- When you're adding a resource group using datastores, you can see a list of datastores. You can select one or more datastores to create a resource group.
- When you're creating a replication plan and creating a resource group within the plan, you can see the VMs in the datastores.

For details, refer to [Create a replication plan](#).

Notifications of free trial or license expiration

This release provides notifications that the free trial will expire in 60 days to ensure you have time to get a license. This release also provides notifications on the day that the license expires.

Notification of service updates

With this release, a banner appears at the top to indicate that services are getting upgraded and that the service is placed in maintenance mode. The banner appears when the service is being upgraded and disappears when the upgrade is complete. While you can continue to work in the UI while the upgrade is in progress, you cannot submit new jobs. Scheduled jobs will run after the update is complete and the service returns to production mode.

10 March 2025

Version 4.2.1

Intelligent proxy support

The NetApp Console agent supports intelligent proxy. Intelligent proxy is a lightweight, secure, and efficient way to connect your on-premises system to NetApp Disaster Recovery. It provides a secure connection between your system and NetApp Disaster Recovery without requiring a VPN or direct internet access. This optimized proxy implementation offloads API traffic within the local network.

When a proxy is configured, NetApp Disaster Recovery attempts to communicate directly with VMware or ONTAP and uses the configured proxy if direct communication fails.

NetApp Disaster Recovery proxy implementation requires port 443 communication between the Console agent and any vCenter Servers and ONTAP arrays using an HTTPS protocol. The NetApp Disaster Recovery agent within the Console agent communicates directly with VMware vSphere, the VC, or ONTAP when performing any actions.

For more information about the intelligent proxy for NetApp Disaster Recovery, see [Set up your infrastructure for NetApp Disaster Recovery](#).

For more information about general proxy set up in the NetApp Console, see [Configure the Console agent to use a proxy server](#).

End the free trial any time

You can stop the free trial at any time or you can wait until it expires.

See [End the free trial](#).

19 February 2025

Version 4.2

ASA r2 support for VMs and datastores on VMFS storage

This release of NetApp Disaster Recovery provides support for ASA r2 for VMs and datastores on VMFS storage. On an ASA r2 system, ONTAP software supports essential SAN functionality while removes features not supported in SAN environments.

This release supports the following features for ASA r2:

- Consistency group provisioning for primary storage (flat consistency group only, meaning only one level without a hierarchical structure)
- Backup (consistency group) operations including SnapMirror automation

The support for ASA r2 in NetApp Disaster Recovery uses ONTAP 9.16.1.

While datastores can be mounted on an ONTAP volume or an ASA r2 storage unit, a resource group in NetApp Disaster Recovery cannot include both a datastore from ONTAP and one from ASA r2. You can select either a datastore from ONTAP or a datastore from ASA r2 in a resource group.

30 October 2024

Reporting

You can now generate and download reports to help you analyze your landscape. Predesigned reports

summarize failovers and failbacks, show replication details on all sites, and show job details for the past seven days.

Refer to [Create disaster recovery reports](#).

30-day free trial

You can now sign up for a 30-day free trial of NetApp Disaster Recovery. Previously, free trials were for 90 days.

Refer to [Set up licensing](#).

Disable and enable replication plans

A previous release included updates to the failover test schedule structure, which was needed to support daily and weekly schedules. This update required that you disable and re-enable all existing replication plans so that you will be able to use the new daily and weekly failover test schedules. This is a one-time requirement.

Here's how:

1. From the menu, select **Replication plans**.
2. Select a plan and select the Actions icon to show the drop-down menu.
3. Select **Disable**.
4. After a few minutes, select **Enable**.

Folder mapping

When you create a replication plan and map compute resources, you can now map folders so that VMs are recovered in a folder you specify for datacenter, cluster, and host.

For details, refer to [Create a replication plan](#).

VM details available for failover, failback, and test failover

When a failure occurs and you are starting a failover, performing a failback, or testing the failover, you can now see details of the VMs and identify which VMs did not restart.

Refer to [Fail over applications to a remote site](#).

VM boot delay with ordered boot sequence

When you create a replication plan, you can now set a boot delay for each VM in the plan. This enables you to set a sequence for the VMs to start to ensure that all your priority one VMs are running before subsequent priority VMs are started.

For details, refer to [Create a replication plan](#).

VM operating system information

When you create a replication plan, you can now see the operating system for each VM in the plan. This is helpful in deciding how to group VMs together in a resource group.

For details, refer to [Create a replication plan](#).

VM name aliasing

When you create a replication plan, you can now add a prefix and suffix to the VM names on the disaster recovery sit. This enables you to use a more descriptive name for the VMs in the plan.

For details, refer to [Create a replication plan](#).

Clean up old snapshots

You can delete any snapshots that are no longer needed beyond your specified retention count. Snapshots might accumulate over time when you lower your snapshot retention count, and you can now remove them to free up space. You can do this anytime on demand or when you delete a replication plan.

For details, refer to [Manage sites, resource groups, replication plans, datastores, and virtual machines information](#).

Reconcile snapshots

You can now reconcile snapshots that are out of sync between the source and target. This might occur if snapshots are deleted on a target outside of NetApp Disaster Recovery. The service deletes the snapshot on the source automatically every 24 hours. However, you can perform this on demand. This feature enables you to ensure that the snapshots are consistent across all sites.

For details, refer to [Manage replication plans](#).

20 September 2024

Support for on-premises to on-premises VMware VMFS datastores

This release includes support for VMs mounted on VMware vSphere virtual machine file system (VMFS) datastores for iSCSI and FC protected to on-premises storage. Previously, the service provided a *technology preview* supporting VMFS datastores for iSCSI and FC.

Here are some additional considerations regarding both iSCSI and FC protocols:

- FC support is for client front-end protocols, not for replication.
- NetApp Disaster Recovery supports only a single LUN per ONTAP volume. The volume should not have multiple LUNs.
- For any replication plan, the destination ONTAP volume should use the same protocols as the source ONTAP volume hosting the protected VMs. For example, if the source uses an FC protocol, the destination should also use FC.

02 August 2024

Support for on-premises to on-premises VMware VMFS datastores for FC

This release includes a *technology preview* of support for VMs mounted on VMware vSphere virtual machine file system (VMFS) datastores for FC protected to on-premises storage. Previously, the service provided a *technology preview* supporting VMFS datastores for iSCSI.



NetApp doesn't charge you for any previewed workload capacity.

Job cancel

With this release, you can now cancel a job in the Job Monitor UI.

Refer to [Monitor jobs](#).

17 July 2024

Failover test schedules

This release includes updates to the failover test schedule structure, which was needed to support daily and weekly schedules. This update requires that you disable and re-enable all existing replication plans so that you will be able to use the new daily and weekly failover test schedules. This is a one-time requirement.

Here's how:

1. From the menu, select **Replication plans**.
2. Select a plan and select the Actions icon to show the drop-down menu.
3. Select **Disable**.
4. After a few minutes, select **Enable**.

Replication plan updates

This release includes updates to replication plan data, which resolves a "snapshot not found" issue. This requires that you change the retention count in all replication plans to 1 and initiate an on-demand snapshot. This process creates a new backup and removes all older backups.

Here's how:

1. From the menu, select **Replication plans**.
2. Select the replication plan, click the **Failover mapping** tab, and click the **Edit** pencil icon.
3. Click the **Datastores** arrow to expand it.
4. Note the value of the retention count in the replication plan. You will need to reinstate this original value when you're finished with these steps.
5. Reduce the count to 1.
6. Initiate an on-demand snapshot. To do so, on the Replication plan page, select the plan, click the Actions icon, and select **Take snapshot now**.
7. After the snapshot job completes successfully, increase the count in the replication plan back to its original value that you noted in the first step.
8. Repeat these steps for all existing replication plans.

05 July 2024

This NetApp Disaster Recovery release includes the following updates:

Support for AFF A-series

This release supports the NetApp AFF A-series hardware platforms.

Support for on-premises to on-premises VMware VMFS datastores

This release includes a *technology preview* of support for VMs mounted on VMware vSphere virtual machine file system (VMFS) datastores protected to on-premises storage. With this release, disaster recovery is supported in a technology preview for on-premises VMware workloads to on-premises VMware environment with VMFS datastores.



NetApp doesn't charge you for any previewed workload capacity.

Replication plan updates

You can add a replication plan more easily by filtering VMs by datastore on the Applications page and by selecting more target details on the Resource mapping page.

Refer to [Create a replication plan](#).

Edit replication plans

With this release, the Failover mappings page has been enhanced for better clarity.

Refer to [Manage plans](#).

Edit VMs

With this release, the process for editing VMs in the plan included some minor UI improvements.

Refer to [Manage VMs](#).

Fail over updates

Before you initiate a failover, you can now determine the status of VMs and whether they are powered on or off. The failover process now enables you to take a snapshot now or choose the snapshots.

Refer to [Fail over applications to a remote site](#).

Failover test schedules

You can now edit the failover tests and set daily, weekly, and monthly schedules for the failover test.

Refer to [Manage plans](#).

Updates to prerequisite information

NetApp Disaster Recovery prerequisites information has been updated.

Refer to [NetApp Disaster Recovery prerequisites](#).

15 May 2024

This NetApp Disaster Recovery release includes the following updates:

Replicating VMware workloads from on-premises to on-premises

This is now released as a general availability feature. Previously, it was a technology preview with limited functionality.

Licensing updates

With NetApp Disaster Recovery, you can sign up for a 90-day free trial, purchase a pay-as-you-go (PAYGO) subscription with Amazon Marketplace, or Bring Your Own License (BYOL), which is a NetApp License File (NLF) that you obtain from your NetApp Sales Rep or from the NetApp Support Site (NSS).

For details about setting up licensing for NetApp Disaster Recovery, refer to [Set up licensing](#).

[Learn more about NetApp Disaster Recovery](#).

05 March 2024

This is the General Availability release of NetApp Disaster Recovery, which includes the following updates.

Licensing updates

With NetApp Disaster Recovery, you can sign up for a 90-day free trial or Bring Your Own License (BYOL), which is a NetApp License File (NLF) that you obtain from your NetApp Sales Rep. You can use the license serial number to get the BYOL activated in the NetApp Console subscriptions. NetApp Disaster Recovery charges are based on provisioned capacity of datastores.

For details about setting up licensing for NetApp Disaster Recovery, refer to [Set up licensing](#).

For details about managing licenses for **all** NetApp Console data services, refer to [Manage licenses for all NetApp Console data services](#).

Edit schedules

With this release, you can now set up schedules to test compliance and failover tests so that you ensure that they will work correctly should you need them.

For details, refer to [Create the replication plan](#).

01 February 2024

This NetApp Disaster Recovery preview release includes the following updates:

Network enhancement

With this release, you can now resize the VM CPU and RAM values. You can also now select a network DHCP or static IP address for the VM.

- DHCP: If you choose this option, you provide credentials for the VM.
- Static IP: You can select the same or different information from the source VM. If you choose the same as the source, you do not need to enter credentials. On the other hand, if you choose to use different information from the source, you can provide the credentials, IP address, subnet mask, DNS, and gateway information.

For details, refer to [Create a replication plan](#).

Custom scripts

Can now be included as post failover processes. With custom scripts, you can have NetApp Disaster Recovery run your script after a failover process. For example, you can use a custom script to resume all database

transactions after the failover is complete.

For details, refer to [Fail over to a remote site](#).

SnapMirror relationship

You can now create a SnapMirror relationship while developing the replication plan. Previously, you had to create the relationship outside of NetApp Disaster Recovery.

For details, refer to [Create a replication plan](#).

Consistency groups

When you create a replication plan, you can include VMs that are from different volumes and different SVMs. NetApp Disaster Recovery creates a Consistency Group Snapshot by including all the volumes and updates all the secondary locations.

For details, refer to [Create a replication plan](#).

VM power-on delay option

When you create a replication plan, you can add VMs to a Resource Group. With Resource Groups, you can set a delay on each VM so that they power up on a delayed sequence.

For details, refer to [Create a replication plan](#).

Application-consistent Snapshot copies

You can specify to create application-consistent Snapshot copies. The service will quiesce the application and then take a Snapshot to obtain a consistent state of the application.

For details, refer to [Create a replication plan](#).

11 January 2024

This preview release of NetApp Disaster Recovery includes the following updates:

Dashboard more quickly

With this release, you can access information on other pages from the Dashboard more quickly.

[Learn about NetApp Disaster Recovery](#).

20 October 2023

This preview release of NetApp Disaster Recovery includes the following updates.

Protect on-premises, NFS-based VMware workloads

Now with NetApp Disaster Recovery, you can protect your on-premises, NFS-based VMware workloads against disasters to another on-premises, NFS-based VMware environment in addition to the public cloud. NetApp Disaster Recovery orchestrates the completion of the disaster recovery plans.



With this preview offering, NetApp reserves the right to modify offering details, contents and timeline before General Availability.

[Learn more about NetApp Disaster Recovery.](#)

27 September 2023

This preview release of NetApp Disaster Recovery includes the following updates:

Dashboard updates

You can now click into the options on the Dashboard, making it easier for you to review the information quickly. Also, the Dashboard now shows the status of failovers and migrations.

Refer to [View the health of your disaster recovery plans on the Dashboard.](#)

Replication plan updates

- **RPO:** You can now enter the Recovery Point Objective (RPO) and Retention count in the Datastores section of the Replication plan. This indicates the amount of data that must exist that is not older than the set time. If, for example, you set it at 5 minutes, the system can lose up to 5 minutes of data if there's a disaster without impacting business critical needs.

Refer to [Create a replication plan.](#)

- **Networking enhancements:** When you are mapping networking between source and target locations in the virtual machines section of the replication plan, NetApp Disaster Recovery now offers two options: DHCP or static IP. Previously, just DHCP was supported. For static IPs, you configure the subnet, gateway, and DNS servers. Additionally, you can now enter credentials for virtual machines.

Refer to [Create a replication plan.](#)

- **Edit schedules:** You can now update replication plan schedules.

Refer to [Manage resources.](#)

- **SnapMirror automation:** While you are creating the replication plan in this release, you can define the SnapMirror relationship between source and target volumes in one of the following configurations:

- 1 to 1
- 1 to many in a fanout architecture
- Many to 1 as a Consistency Group
- Many to many

Refer to [Create a replication plan.](#)

01 August 2023

NetApp Disaster Recovery preview

NetApp Disaster Recovery preview is a cloud-based disaster recovery service that automates disaster recovery workflows. Initially, with the NetApp Disaster Recovery preview, you can protect your on-premises,

NFS-based VMware workloads running NetApp storage to VMware Cloud (VMC) on AWS with Amazon FSx for ONTAP.



With this preview offering, NetApp reserves the right to modify offering details, contents and timeline before General Availability.

[Learn more about NetApp Disaster Recovery.](#)

This release includes the following updates:

Resource groups update for boot order

When you create a disaster recovery or replication plan, you can add virtual machines into functional resource groups. Resource groups enable you to put a set of dependent virtual machines into logical groups that meet your requirements. For example, groups could contain boot order that can be executed upon recovery. With this release, each resource group can include one or more virtual machines. The virtual machines will power on based on the sequence in which you include them in the plan. Refer to [Select applications to replicate and assign resource groups](#).

Replication verification

After you create the disaster recovery or replication plan, identify the recurrence in the wizard, and initiate a replication to a disaster recovery site, every 30 minutes NetApp Disaster Recovery verifies that the replication is actually occurring according to the plan. You can monitor the progress in the Job Monitor page. Refer to [Replicate applications to another site](#).

Replication plan shows recovery point objective (RPO) transfer schedules

When you create a disaster recovery or replication plan, you select the VMs. In this release, you can now view the SnapMirror associated with each of the volumes that are associated with the datastore or VM. You can also see the RPO transfer schedules that are associated with the SnapMirror schedule. RPO helps you determine whether your backup schedule is enough to recover after a disaster. Refer to [Create a replication plan](#).

Job Monitor update

The Job Monitor page now includes a Refresh option so that you can get an up-to-date status of operations. Refer to [Monitor disaster recovery jobs](#).

18 May 2023

This is the initial release of NetApp Disaster Recovery.

Cloud-based disaster recovery service

NetApp Disaster Recovery is a cloud-based disaster recovery service that automates disaster recovery workflows. Initially, with the NetApp Disaster Recovery preview, you can protect your on-premises, NFS-based VMware workloads running NetApp storage to VMware Cloud (VMC) on AWS with Amazon FSx for ONTAP.

[Learn more about NetApp Disaster Recovery.](#)

Limitations in NetApp Disaster Recovery

Known limitations identify platforms, devices, or functions that are not supported by this

release of the service, or that do not interoperate correctly with it.

Wait until failback completes before running discovery

After a failover has finished, do not initiate discovery on the source vCenter manually. Wait until the failback has finished and then initiate discovery on the source vCenter.

NetApp Console might not discover Amazon FSx for NetApp ONTAP

Sometimes, the NetApp Console does not discover Amazon FSx for NetApp ONTAP clusters. This might be because the FSx credentials were not correct.

Workaround: Add the Amazon FSx for NetApp ONTAP cluster in the NetApp Console and periodically refresh the cluster to display any changes.


If you need to remove the ONTAP FSx cluster from NetApp Disaster Recovery, complete the following steps:

1. In the NetApp Console agent, use the connectivity options from your cloud provider, connect to the Linux VM that the Console agent runs on, restart the "occm" service using the `docker restart occm` command.

Refer to [Manage existing Console agents](#).

1. In the NetApp Console Systems page, add the Amazon FSx for ONTAP system again and provide the FSx credentials.

Refer to [Create an Amazon FSx for NetApp ONTAP file system](#).

2. From NetApp Disaster Recovery, select **Sites**, on the vCenter row select the **Actions** option , and from the Actions menu, select **Refresh** to refresh the FSx discovery in NetApp Disaster Recovery.

This rediscovers the datastore, its virtual machines, and its destination relationship.

Get started

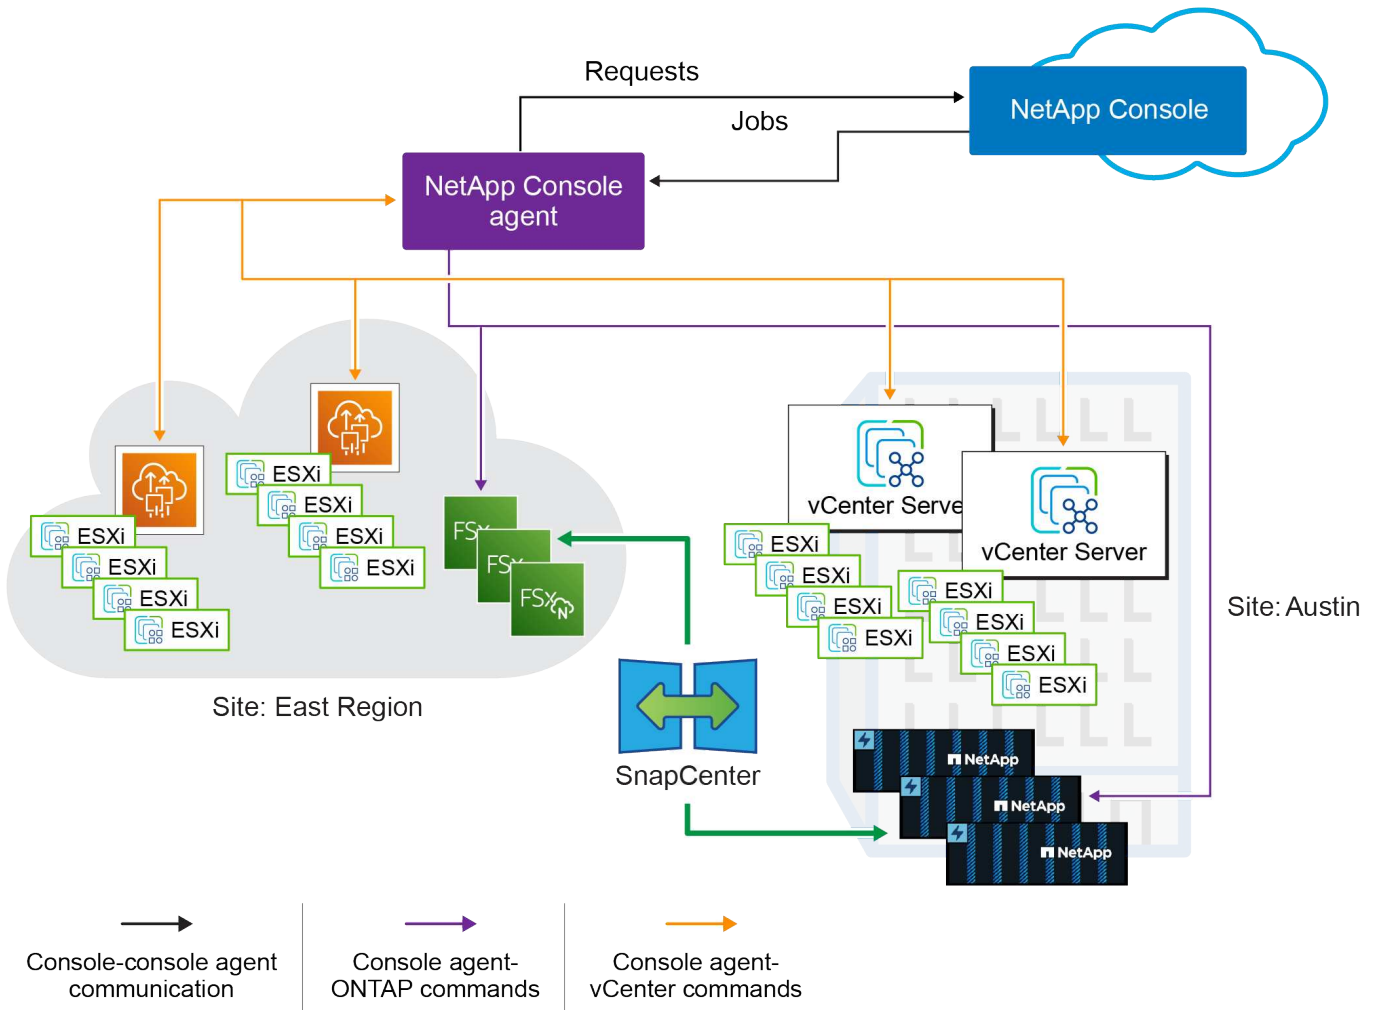
Learn about NetApp Disaster Recovery for VMware

Disaster recovery to the cloud is a resilient and cost-effective way of protecting workloads against site outages and data corruption events. With NetApp Disaster Recovery for VMware, you can replicate your on-premises VMware VM or datastore workloads running ONTAP storage to a VMware software-defined data center in a public cloud using NetApp cloud storage or to another on-premises VMware environment with ONTAP storage as a disaster recovery site. You can use Disaster Recovery also to migrate VM workloads from one site to another.

NetApp Disaster Recovery is a cloud-based disaster recovery service that automates disaster recovery workflows. With NetApp Disaster Recovery, you can protect your on-premises, NFS-based workloads and VMware vSphere virtual machine file system (VMFS) datastores for iSCSI and FC running NetApp storage to one of the following:

- VMware Cloud (VMC) on AWS with Amazon FSx for NetApp ONTAP
- Amazon Elastic VMware Service (EVS) with Amazon FSx for NetApp ONTAP For details, refer to [Introduction of NetApp Disaster Recovery using Amazon Elastic VMware Service and Amazon FSx for NetApp ONTAP](#).
- Azure VMware Solution (AVS) with NetApp Cloud Volumes ONTAP (iSCSI) (Private preview)
- Another on-premises NFS and or VMFS-based (iSCSI/FC) VMware environment with ONTAP storage

NetApp Disaster Recovery uses ONTAP SnapMirror technology with integrated native VMware orchestration to protect VMware VMs and their associated on-disk OS images, while retaining all storage efficiency benefits of ONTAP. Disaster Recovery uses these technologies as the replication transport to the disaster recovery site. This enables industry-best storage efficiency (compression and deduplication) on primary and secondary sites.



NetApp Console

NetApp Disaster Recovery is accessible through the NetApp Console.

The NetApp Console provides centralized management of NetApp storage and data services across on-premises and cloud environments at enterprise grade. The Console is required to access and use NetApp data services. As a management interface, it enables you to manage many storage resources from one interface. Console administrators can control access to storage and services for all systems within the enterprise.

You don't need a license or subscription to start using NetApp Console and you only incur charges when you need to deploy Console agents in your cloud to ensure connectivity to your storage systems or NetApp data services. However, some NetApp data services accessible from the Console are licensed or subscription-based.

Learn more about the [NetApp Console](#).

Benefits of using NetApp Disaster Recovery for VMware

NetApp Disaster Recovery offers the following benefits:


- Simplified user experience for vCenter discovery and recovery of applications with multiple point-in-time recovery operations.

- Lower total cost of ownership with reduced cost of operations and ability to create and adjust disaster recovery plans with minimal resources.
- Continuous disaster recovery readiness with virtual failover testing that does not disrupt operations. You can regularly test your DR failover plans without impacting production workloads.
- Faster time to value with dynamic changes in your IT environment and ability to address it in your disaster recovery plans.
- Ability to manage both the storage and virtual layers through backend orchestration of both ONTAP and VMware at the same time without the need for virtual server appliances (VSAs) that need to be deployed and maintained.
- DR solutions for VMware can be resource intensive. Many DR solutions replicate VMs at the VMware virtual layer using VSAs, which can consume more compute resources and lose the valuable storage efficiencies of ONTAP. Because Disaster Recovery uses ONTAP SnapMirror technology, it can replicate data from production datastores to the DR site using our incremental-forever replication model with all the native data compression and deduplication efficiencies of ONTAP.

What you can do with NetApp Disaster Recovery for VMware

NetApp Disaster Recovery provides you with full use of several NetApp technologies to accomplish the following goals:

- Replicate VMware apps on your on-premises production site to a disaster recovery remote site in the cloud or on-premises using SnapMirror replication.
- Migrate VMware workloads from your original site to another site.
- Conduct a failover test. When you do this, the service creates temporary virtual machines. Disaster Recovery makes a new FlexClone volume from the selected snapshot, and a temporary datastore, which is backed by the FlexClone volume, is mapped to the ESXi hosts. This process doesn't consume additional physical capacity on on-premises ONTAP storage or FSx for NetApp ONTAP storage in AWS. The original source volume is not modified and replica jobs can continue even during disaster recovery.
- In case of disaster, fail over your primary site on demand to the disaster recovery site, which can be VMware Cloud on AWS with Amazon FSx for NetApp ONTAP or an on-premises VMware environment with ONTAP.
- After the disaster has been resolved, fail back on demand from the disaster recovery site to the primary site.
- Group VMs or datastores into logical resource groups for efficient management.



Configuration of the vSphere server is done outside of NetApp Disaster Recovery in vSphere Server.

Cost

NetApp doesn't charge you for using the trial version of NetApp Disaster Recovery.

NetApp Disaster Recovery can be used either with a NetApp license or an annual subscription-based plan through Amazon Web Services.



Some releases include a technology preview. NetApp doesn't charge you for any previewed workload capacity. See [What's new in NetApp Disaster Recovery](#) for information about the latest technology previews.

Licensing

You can use the following license types:

- Sign up for a 30-day free trial.
- Purchase a pay-as-you-go (PAYGO) subscription with the Amazon Web Services (AWS) Marketplace or Microsoft Azure Marketplace. This license enables you to purchase a fixed protected capacity license without any long-term commitment.
- Bring your own license (BYOL), which is a NetApp License File (NLF) that you obtain from your NetApp Sales Rep. You can use the license serial number to get the BYOL activated in the NetApp Console.

Licenses for all NetApp data services are managed through subscriptions in the NetApp Console. After you set up your BYOL, you can see an active license for the service in the Console.

The service is licensed based on the amount of data hosted on protected ONTAP volumes. The service determines which volumes should be considered for licensing purposes by mapping protected VMs to their vCenter datastores. Each datastore is hosted on an ONTAP volume or LUN. The used capacity reported by ONTAP for that volume or LUN is used for licensing determinations.

Protected volumes can host many VMs. Some might not be part of a NetApp Disaster Recovery resource group. Regardless, the storage consumed by all VMs on that volume or LUN is used against the license

maximum capacity.



NetApp Disaster Recovery charges are based on used capacity of datastores on the source site when there is at least one VM that has a replication plan. Capacity for a failed over datastore is not included in the capacity allowance. For a BYOL, if the data exceeds the allowed capacity, operations in the service are limited until you obtain an additional capacity license or upgrade the license in the NetApp Console.

For details about setting up licensing for NetApp Disaster Recovery, refer to [Set up NetApp Disaster Recovery licensing](#).

30-day free trial

You can try out NetApp Disaster Recovery by using a 30-day free trial.


To continue after the 30-day trial, you'll need to obtain a Pay-as-you-go (PAYGO) subscription from your cloud provider or purchase a BYOL license from NetApp.

You can purchase a license at any time and you will not be charged until the 30-day trial ends.

How NetApp Disaster Recovery works

NetApp Disaster Recovery is a service hosted within the NetApp Console software as a service (SaaS) environment. Disaster Recovery can recover workloads replicated from an on-premises site to Amazon FSx for ONTAP or to another on-premises site. This service automates the recovery from the SnapMirror level, through virtual machine registration to VMware Cloud on AWS, and to network mappings directly on the VMware network virtualization and security platform, NSX-T. This feature is included with all Virtual Machine Cloud environments.

NetApp Disaster Recovery uses ONTAP SnapMirror technology, which provides highly efficient replication and preserves the ONTAP incremental-forever snapshot efficiencies. SnapMirror replication ensures that application-consistent snapshot copies are always in sync and the data is usable immediately after a failover.



When there is a disaster, this service helps you recover virtual machines in the other on-premises VMware environment or VMC by breaking the SnapMirror relationships and making the destination site active.

- The service also lets you fail back virtual machines to the original source location.
- You can test the disaster recovery failover process without disrupting the original virtual machines. The test recovers virtual machines to an isolated network by creating a FlexClone of the volume.
- For the failover or test failover process, you can choose the latest (default) or selected snapshot from which to recover your virtual machine.

Components of Disaster Recovery

Disaster Recovery uses the following components to provide disaster recovery for VMware workloads:

- **NetApp Console:** The user interface for managing your disaster recovery plans. You can use the NetApp Console to create and manage replication plans, resource groups, and failover operations across your on-premises and cloud environments.
- **Console agent:** A lightweight software component that runs in your cloud-hosted network or your on-premises VMware environment. It communicates with the NetApp Console and manages the replication of data between your on-premises environment and the disaster recovery site. The Console agent is installed on a virtual machine in your VMware environment.
- **ONTAP storage clusters:** The ONTAP storage clusters are the primary storage systems that host your VMware workloads. The ONTAP storage clusters provide the underlying storage infrastructure for your

disaster recovery plans. Disaster Recovery uses ONTAP storage APIs to manage ONTAP storage clusters such as on-premises arrays, and cloud-based solutions, such as Amazon FSx for NetApp ONTAP.

- **vCenter servers:** The VMware vCenter is the management server for your VMware environment. It manages the ESXi hosts and their associated datastores. The Console agent communicates with the VMware vCenter to manage the replication of data between your on-premises environment and the disaster recovery site. This includes registering ONTAP LUNs and volumes as datastores, reconfiguring VMs, and starting and stopping VMs.

The Disaster Recovery protection workflow

When a replication plan is assigned to a resource group, Disaster Recovery performs a discovery check of all the components in the resource group and plan to ensure that the plan can be activated.

If this check is successful, Disaster Recovery performs the following initialization steps:

1. For each VM in the target resource group, identify the hosting VMware datastore.
2. For each VMware datastore found, identify the hosting ONTAP FlexVol volume or LUN.
3. For each ONTAP volume and LUN found, determine if there is an existing SnapMirror relationship between the source volumes and a destination volume in the destination site.
 - a. If there is no pre-existing SnapMirror relationship, create any new destination volumes and create a new SnapMirror relationship between each unprotected source volume.
 - b. If there is a pre-existing SnapMirror relationship, use that relationship to perform all replication operations.

After Disaster Recovery creates and initializes all relationships, at each scheduled backup, the service performs the following data protection steps:

1. For each VM flagged as “application consistent,” use VMtools to place the supported application into a backup state.
2. Create a new snapshot of all ONTAP volumes hosting protected VMware datastores.
3. Perform a SnapMirror update operation to replicate those snapshots to the destination ONTAP cluster.
4. Determine if the number of retained snapshots has exceeded the maximum snapshot retention defined in the replication plan and delete any extraneous snapshots from both the source and destination volumes.

Supported protection targets and datastore types

Datastore types supported

NetApp Disaster Recovery supports the following datastore types:

- NFS datastores hosted on ONTAP FlexVol volumes residing on ONTAP clusters.
- VMware vSphere virtual machine file system (VMFS) datastores using the iSCSI or FC protocol

Protection targets supported

- VMware Cloud (VMC) on AWS with Amazon FSx for NetApp ONTAP
- Another on-premises, NFS-based VMware environment with ONTAP storage or an on-premises FC/iSCSI VMSF
- Amazon Elastic VMware Service
- Azure VMware Solution (AVS) with NetApp Cloud Volumes ONTAP (iSCSI) (Private preview)

Terms that might help you with NetApp Disaster Recovery

You might benefit by understanding some terminology related to disaster recovery.

- **Datastore:** A VMware vCenter data container, which uses a file system to hold VMDK files. Typical datastore types are NFS, VMFS, vSAN or vVol. Disaster Recovery supports NFS and VMFS datastores. Each VMware datastore is hosted on a single ONTAP volume or LUN. Disaster Recovery supports NFS and VMFS datastores hosted on FlexVol volumes residing on ONTAP clusters.
- **Replication plan:** A set of rules about how often backups occur and how to handle failover events. Plans are assigned to one or more resource groups.
- **Recovery point objective (RPO):** The maximum amount of data loss that is acceptable in the event of a disaster. RPO is defined in the replication plan's frequency of data replication or replication schedule.
- **Recovery time objective (RTO):** The maximum amount of time that is acceptable to recover from a disaster. RTO is defined in the replication plan and is the time it takes to fail over to the DR site and restart all VMs.
- **Resource group:** A logical container that enables you to manage multiple VMs as a single unit. A VM can be in only one resource group at a time. You can create a resource group for each application or workload that you want to protect.
- **Site:** A logical container typically associated with a physical datacenter or cloud location hosting one or more vCenter clusters and ONTAP storage.

NetApp Disaster Recovery prerequisites

Before using NetApp Disaster Recovery, you should ensure that your environment meets the ONTAP storage, VMware vCenter cluster, and NetApp Console requirements.

Software versions

Component	Minimum version
ONTAP software	ONTAP 9.10.0 or later
VMware on-premises vCenter	7.0u3 or later
VMware Cloud for AWS	Latest available version
Amazon FSx for NetApp ONTAP	Latest available version

ONTAP storage prerequisites

These prerequisites apply to either ONTAP or Amazon FSX for NetApp ONTAP instances.

- Source and destination clusters must have a peer relationship.
- The SVM that will host the disaster recovery volumes must exist on the destination cluster.
- The source SVM and destination SVM must have a peer relationship.
- If deploying with Amazon FSx for NetApp ONTAP, the following prerequisite applies:

- An Amazon FSx for NetApp ONTAP instance to host VMware DR datastores must exist in your VPC. Refer to Amazon FSx for ONTAP documentation on [how to get started](#).

VMware vCenter clusters prerequisites

These prerequisites apply to both on-premises vCenter clusters and to VMware Cloud for AWS software-defined data center (SDDC).

- Review [vCenter privileges](#) required for NetApp Disaster Recovery.
- All VMware clusters that you want NetApp Disaster Recovery to manage use ONTAP volumes to host any VMs that you want to protect.
- All VMware datastores to be managed by NetApp Disaster Recovery must use one of the following protocols:
 - NFS
 - VMFS using the iSCSI or FC protocol
- VMware vSphere version 7.0 Update 3 (7.0v3) or later
- If you are using VMware Cloud SDDC, these prerequisites apply.
 - In the VMware Cloud Console, use the service roles of Administrator and NSX Cloud Administrator. Also use the organization owner for the Organization role. Refer to [Using VMware Cloud Foundations with AWS FSx for NetApp ONTAP documentation](#).
 - Link the VMware Cloud SDDC with Amazon FSx for NetApp ONTAP instance. Refer to [VMware Cloud on AWS integration with Amazon FSx for NetApp ONTAP deployment information](#).

NetApp Console prerequisites

Get started with the NetApp Console

If you haven't already done so, [sign up to the NetApp Console and create an organization](#).

Gather credentials for ONTAP and VMware

- Amazon FSx for ONTAP and AWS credentials must be added to the system within the NetApp Console project that will be used to manage NetApp Disaster Recovery.
- NetApp Disaster Recovery requires vCenter credentials. You enter the vCenter credentials when you're adding a site in NetApp Disaster Recovery.

For a list of vCenter privileges needed, refer to [vCenter privileges needed for NetApp Disaster Recovery](#). For instructions on how to add a site, refer to [Add a site](#).

Create the NetApp Console agent

The Console agent is a software component that enables the Console to communicate with your ONTAP storage and VMware vCenter clusters. It is required for Disaster Recovery to function properly. The agent resides in your private network (either in an on-premises data center or in a cloud VPC) and communicates with your ONTAP storage instances and any additional server and application components. For Disaster Recovery, this is access to your managed vCenter clusters.

A Console agent must be set up in the NetApp Console. When you use the agent, it will include the appropriate capabilities for the Disaster Recovery service.

- NetApp Disaster Recovery works only with the standard mode agent deployment. See [Getting started with the NetApp Console in standard mode](#).
- Ensure that both the source and destination vCenters use the same Console agent.
- Type of Console agent needed:
 - **On-premises to on-premises disaster recovery:** Install the Console on-premises agent in the disaster recovery site. Using this method, a failure of the primary site doesn't prevent the service from restarting your virtual resources at the DR site. Refer to [Install and set up the Console agent on premises](#).
 - **On-premises to AWS:** Install the Console agent for AWS in your AWS VPC. Refer to [Console agent installation options in AWS](#).



For on-premises to on-premises, use the on-premises Console agent. For on-premises to AWS, use the AWS Console agent, which has access to the source on-premises vCenter and the destination on-premises vCenter.

- The installed Console agent must be able to access any VMware cluster that NetApp Disaster Recovery will manage.
- All ONTAP arrays to be managed by NetApp Disaster Recovery must be added to any system within the NetApp Console project that will be used to manage NetApp Disaster Recovery.

See [Discover on-premises ONTAP clusters](#).

- For information about setting up an intelligent proxy for NetApp Disaster Recovery, see [Set up your infrastructure for NetApp Disaster Recovery](#).

Workload prerequisites

To ensure that application-consistency processes are successful, apply these prerequisites:

- Ensure that VMware tools (or Open VM tools) are running on the VMs that will be protected.
- For Windows VMs running Microsoft SQL Server or Oracle Database or both, the databases should have their VSS Writers enabled.
- Oracle databases that are running on a Linux operating system should have the operating system user authentication enabled for the Oracle database SYSDBA role.

Quick start for NetApp Disaster Recovery

Here's an overview of the steps needed to get started with NetApp Disaster Recovery. The links within each step take you to a page that provides more details.

1

Review prerequisites

[Ensure your system meets these requirements](#).

2

Set up NetApp Disaster Recovery

- [Set up the infrastructure for the service](#).

- [Set up licensing.](#)

3

What's next?

After you set up the service, here's what you might do next.

- [Add your vCenter sites to NetApp Disaster Recovery.](#)
- [Create your first resource group.](#)
- [Create your first replication plan.](#)
- [Replicate applications to another site.](#)
- [Fail over applications to a remote site.](#)
- [Fail back applications to the original source site.](#)
- [Manage sites, resource groups, and replication plans.](#)
- [Monitor disaster recovery operations.](#)

Set up your infrastructure for NetApp Disaster Recovery

To use NetApp Disaster Recovery, perform a few steps to set it up both in Amazon Web Services (AWS) and in the NetApp Console.



Review [prerequisites](#) to ensure that your system is ready.

You can use NetApp Disaster Recovery in the following infrastructures:

- Hybrid cloud DR that replicates an on-premises VMware plus ONTAP datacenter to an AWS DR infrastructure based on VMware Cloud on AWS and Amazon FSx for NetApp ONTAP.
- Private cloud DR that replicates an on-premises VMware plus ONTAP vCenter to another on-premises VMware plus ONTAP vCenter.

Hybrid cloud with VMware Cloud and Amazon FSx for NetApp ONTAP

This method consists of an on-premises production vCenter infrastructure using datastores hosted on ONTAP FlexVol volumes using an NFS protocol. The DR site consists of one or more VMware Cloud SDDC instances using datastores hosted on FlexVol volumes provided by one or more FSx for ONTAP instances using an NFS protocol.

The production and DR sites are connected by an AWS-compatible secure connection. Common connection types are a secure VPN (private or AWS provided), AWS Direct Connect, or other approved interconnect methods.

For Disaster Recovery involving AWS cloud infrastructure, you must use the Console agent for AWS. The agent should be installed in the same VPC as the FSx for ONTAP instance. If additional FSx for ONTAP instances were deployed in other VPCs, the VPC hosting the agent must have access to the other VPCs.

AWS availability zones

AWS supports deploying solutions in one or more availability zones (AZ) within a given region. Disaster Recovery uses two AWS hosted services: VMware Cloud for AWS and AWS FSx for NetApp ONTAP.

- **VMware Cloud for AWS:** Supports the deployment in a single-AZ or in a dual-AZ stretch-cluster SDDC environment. Disaster Recovery supports a single-AZ SDDC deployment only for Amazon VMware Cloud for AWS.
- **AWS FSx for NetApp ONTAP:** When this is deployed in a dual-AZ configuration, each volume is owned by a single FSx system. Each volume is owned by a single FSx system. The volume's data is mirrored to the second FSx system. The FSx for ONTAP systems can be deployed in either single- or dual-AZ deployments. Disaster Recovery supports both single- and multi-AZ FSx for ONTAP deployments.

BEST PRACTICE: For AWS DR site configuration, NetApp recommends using single-AZ deployments for both VMware Cloud and AWS FSx for ONTAP instances. Because AWS is being used for DR, there is no advantage to introducing multiple AZs. Multi-AZs can increase costs and complexity.

On-premises to AWS

AWS provides the following methods to connect private datacenters to the AWS cloud. Each solution has its benefits and cost considerations.

- **AWS Direct Connect:** This is an AWS cloud interconnect located in the same geographic area as your private datacenter and provided by an AWS partner. This solution provides a secure, private connection between your local datacenter and the AWS cloud without the need for a public internet connection. This is the most direct and efficient connection method offered by AWS.
- **AWS Internet Gateway:** This provides public connectivity between AWS cloud resources and external compute resources. This type of connection is typically used to provide service offerings to external customers, such as HTTP/HTTPS service where security is not a requirement. There is no quality-of-service control, security, or guarantee of connectivity. For this reason, this connection method is not recommended for connecting a production datacenter to the cloud.
- **AWS Site-Site VPN:** This virtual private network connection can be used to provide secure access connections along with a public internet service provider. The VPN encrypts and decrypts all data traveling to and from the AWS cloud. VPNs can be either software- or hardware-based. For enterprise applications, the public internet service provider (ISP) should offer quality-of-service guarantees to ensure that adequate bandwidth and latency are provided for DR replication.

BEST PRACTICE: For AWS DR site configuration, NetApp recommends using AWS Direct Connect. This solution provides the highest performance and security for enterprise applications. If it is not available, a high-performance public ISP connection along with a VPN should be used. Ensure that the ISP offers commercial QoS service levels to ensure adequate network performance.

VPC-to-VPC interconnections

AWS offers the following types of VPC-to-VPC interconnections. Each solution has its benefits and cost considerations.

- **VPC Peering:** This is a private connection between two VPCs. It is the most direct and efficient connection method offered by AWS. VPC peering can be used to connect VPCs in the same or different AWS regions.
- **AWS Internet Gateway:** This is typically used to provide connections between AWS VPC resources and non-AWS resources and endpoints. All traffic follows a "hair-pin" path where VPC traffic destined to another VPC exits the AWS infrastructure through the internet gateway and returns to the AWS infrastructure through the same or different gateway. This is not a suitable VPC connection type for enterprise VMware solutions.
- **AWS Transit Gateway:** This is a centralized router-based connection type that enables each VPC to connect to a single, central gateway, which acts as a central hub for all VPC-to-VPC traffic. This can also be connected to your VPN solution to enable on-premises datacenter resources to access AWS VPC-

hosted resources. This type of connection typically requires an additional cost to implement.

BEST PRACTICE: For DR solutions involving VMware Cloud and a single FSx for ONTAP VPC, NetApp recommends that you use the VPC peer connection. If multiple FSx for ONTAP VPCs are deployed, then we recommend using an AWS Transit Gateway to reduce the management overhead of multiple VPC peer connections.

Get ready for on-premises-to-cloud protection using AWS

To set up NetApp Disaster Recovery for on-premises-to-cloud protection using AWS, you need to set up the following:

- Set up AWS FSx for NetApp ONTAP
- Set up VMware Cloud on AWS SDDC

Set up AWS FSx for NetApp ONTAP

- Create an Amazon FSx for NetApp ONTAP file system.
 - Provision and configure FSx for ONTAP. Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage, built on the NetApp ONTAP file system.
 - Follow the steps in [Technical Report 4938: Mount Amazon FSx ONTAP as an NFS datastore with VMware Cloud on AWS](#) and [Quick start for Amazon FSx for NetApp ONTAP](#) to provision and configure FSx for ONTAP.
- Add Amazon FSx for ONTAP to the system, and add AWS credentials for FSx for ONTAP.
- Create or verify your destination ONTAP SVM in AWS FSx for ONTAP instance.
- Configure replication between your source on-premises ONTAP cluster and your FSx for ONTAP instance in the NetApp Console.

Refer to [how to set up an FSx for ONTAP system](#) for detailed steps.

Set up VMware Cloud on AWS SDDC

[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 VMware Cloud environment on AWS, follow the steps in [Deploy and configure the Virtualization Environment on AWS](#). A pilot-light cluster can also be used for disaster recovery purposes.

Private cloud

You can use NetApp Disaster Recovery to protect VMware VMs hosted on one or more vCenter clusters by replicating VM datastores to another vCenter cluster either in the same private datacenter or to a remote private or collocated datacenter.

For on-premises to on-premises situations, install the Console agent at one of the physical sites.

Disaster Recovery supports site-to-site replication using Ethernet and TCP/IP. Ensure that adequate bandwidth is available to support data change rates on the production site VMs so that all changes can be replicated to

the DR site within the Recovery Point Objective (RPO) time frame.

Get ready for on-premises-to-on-premises protection

Ensure that the following requirements are met before you set up NetApp Disaster Recovery for on-premises-to-on-premises protection:

- ONTAP storage
 - Ensure that you have ONTAP credentials.
 - Create or verify your disaster recovery site.
 - Create or verify your destination ONTAP SVM.
 - Ensure that your source and destination ONTAP SVMs are peered.
- vCenter clusters
 - Ensure that the VMs you want to protect are hosted on NFS datastores (using ONTAP NFS volumes) or VMFS datastores (using NetApp iSCSI LUNs).
 - Review [vCenter privileges](#) required for NetApp Disaster Recovery.
 - Create a disaster recovery user account (not the default vCenter admin account) and assign the vCenter privileges to the account.

Intelligent proxy support

The NetApp Console agent supports intelligent proxy. Intelligent proxy is a lightweight, secure, and efficient way to connect your on-premises environment to the NetApp Console. It provides a secure connection between your system and the Console service without requiring a VPN or direct internet access. This optimized proxy implementation offloads API traffic within the local network.

When a proxy is configured, NetApp Disaster Recovery attempts to communicate directly with VMware or ONTAP and uses the configured proxy if direct communication fails.

NetApp Disaster Recovery proxy implementation requires port 443 communication between the Console agent and any vCenter Servers and ONTAP arrays using an HTTPS protocol. The NetApp Disaster Recovery agent within the Console agent communicates directly with VMware vSphere, the VC, or ONTAP when performing any actions.

For more information about general proxy set up in the NetApp Console, see [Configure the Console agent to use a proxy server](#).

Access NetApp Disaster Recovery

You use the NetApp Console to log in to the NetApp Disaster Recovery service.

To log in, you can use your NetApp Support Site credentials or you can sign up for a NetApp cloud login using your email and a password. [Learn more about logging in](#).

Specific tasks require specific user roles.

[Learn about user roles and permissions in NetApp Disaster Recovery](#).

[Learn about NetApp Console access roles for all services](#).

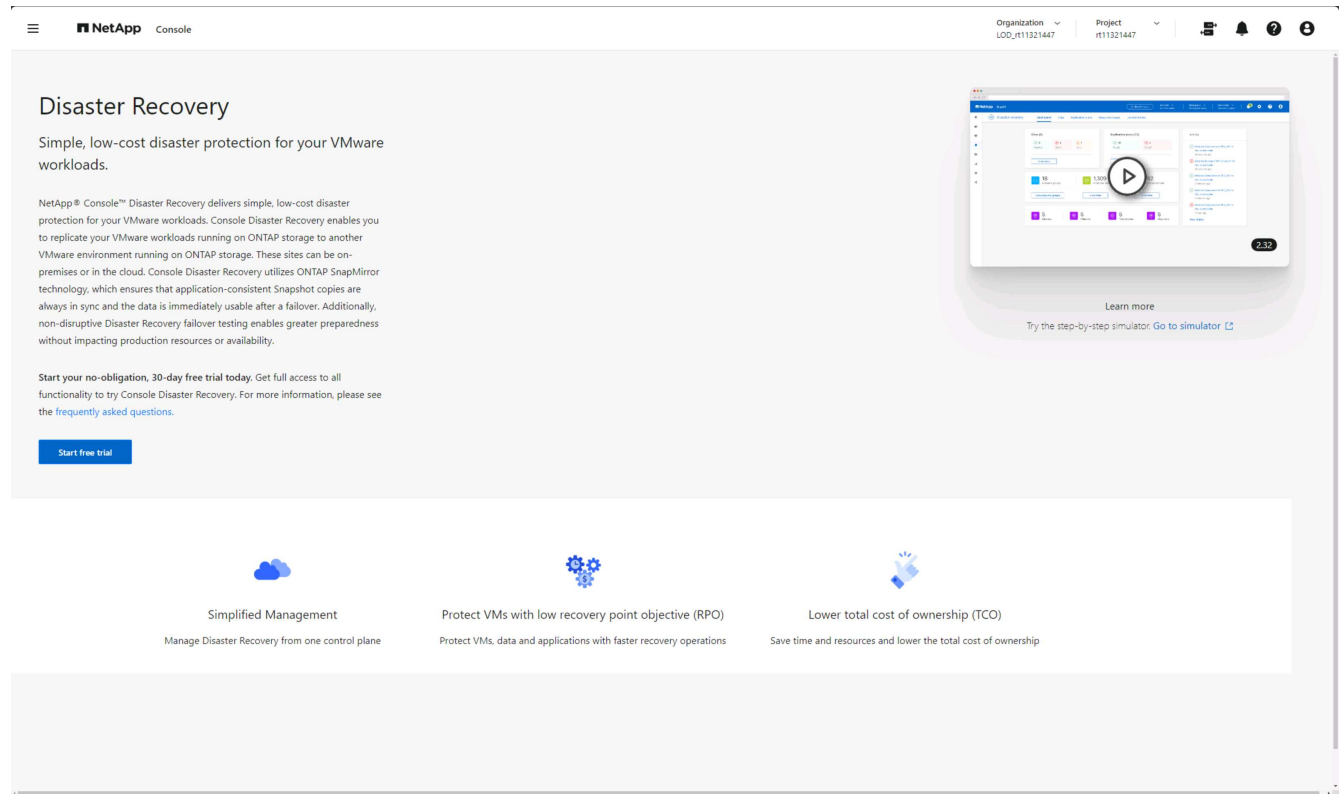
Steps

1. Open a web browser and go to the [NetApp Console](#).

The NetApp Console login page appears.


2. Log in to the NetApp Console.
3. From the NetApp Console left navigation, select **Protection > Disaster recovery**.

If this is your first time logging in to this service, the landing page appears and you can sign up for a free trial.



Otherwise, the NetApp Disaster Recovery Dashboard appears.

- If you haven't yet added a NetApp Console agent, you'll need to add one. To add the agent, refer to [Learn about Console agents](#).
- If you are a NetApp Console user with an existing agent, when you select "Disaster recovery," a message appears about signing up.
- If you are already using the service, when you select "Disaster recovery," the Dashboard appears.



Set up licensing for NetApp Disaster Recovery

With NetApp Disaster Recovery, you can use different licensing plans including a free trial, a pay-as-you-go subscription, or bring your own license.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, or Disaster recovery application admin role.

[Learn about user roles and permissions in NetApp Disaster Recovery.](#)

[Learn about access roles for all services.](#)

Licensing options

You can use the following licensing options:

- Sign up for a 30-day free trial.
- Purchase a pay-as-you-go (PAYGO) subscription to Amazon Web Services (AWS) Marketplace or to Microsoft Azure Marketplace.
- Bring your own license (BYOL), which is a NetApp License File (NLF) that you obtain from your NetApp Sales Rep. You can use the license serial number to get the BYOL activated in the NetApp Console.



NetApp Disaster Recovery charges are based on used capacity of datastores on the source site when there is at least one VM that has a replication plan. Capacity for a failed over datastore is not included in the capacity allowance. For a BYOL, if the data exceeds the allowed capacity, operations in the service are limited until you obtain an additional capacity license or upgrade the license in the NetApp Console.

[Learn more about subscriptions.](#)

After the free trial ends or the license expires, you can still do the following in the service:

- View any resource, such as a workload or replication plan.
- Delete any resource, such as a workload or replication plan.
- Run all scheduled operations that were created during the trial period or under the license.

Try it out using a 30-day free trial

You can try NetApp Disaster Recovery out by using a 30-day free trial.



No capacity limits are enforced during the trial.


To continue after the trial, you'll need to purchase a BYOL license or PAYGO AWS subscription. You can get a license at any time and you will not be charged until the trial ends.

During the trial, you have full functionality.

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.

If this is your first time logging in to this service, the landing page appears.



3. If you haven't already added a Console agent for other services, add one.

To add a Console agent, refer to [Learn about Console agents](#).

4. After you set up the agent, in the NetApp Disaster Recovery landing page, the button to add the agent changes to a button for starting a free trial. Select **Start free trial**.
5. Begin by adding vCenters.

For details, see [Add vCenter sites](#).

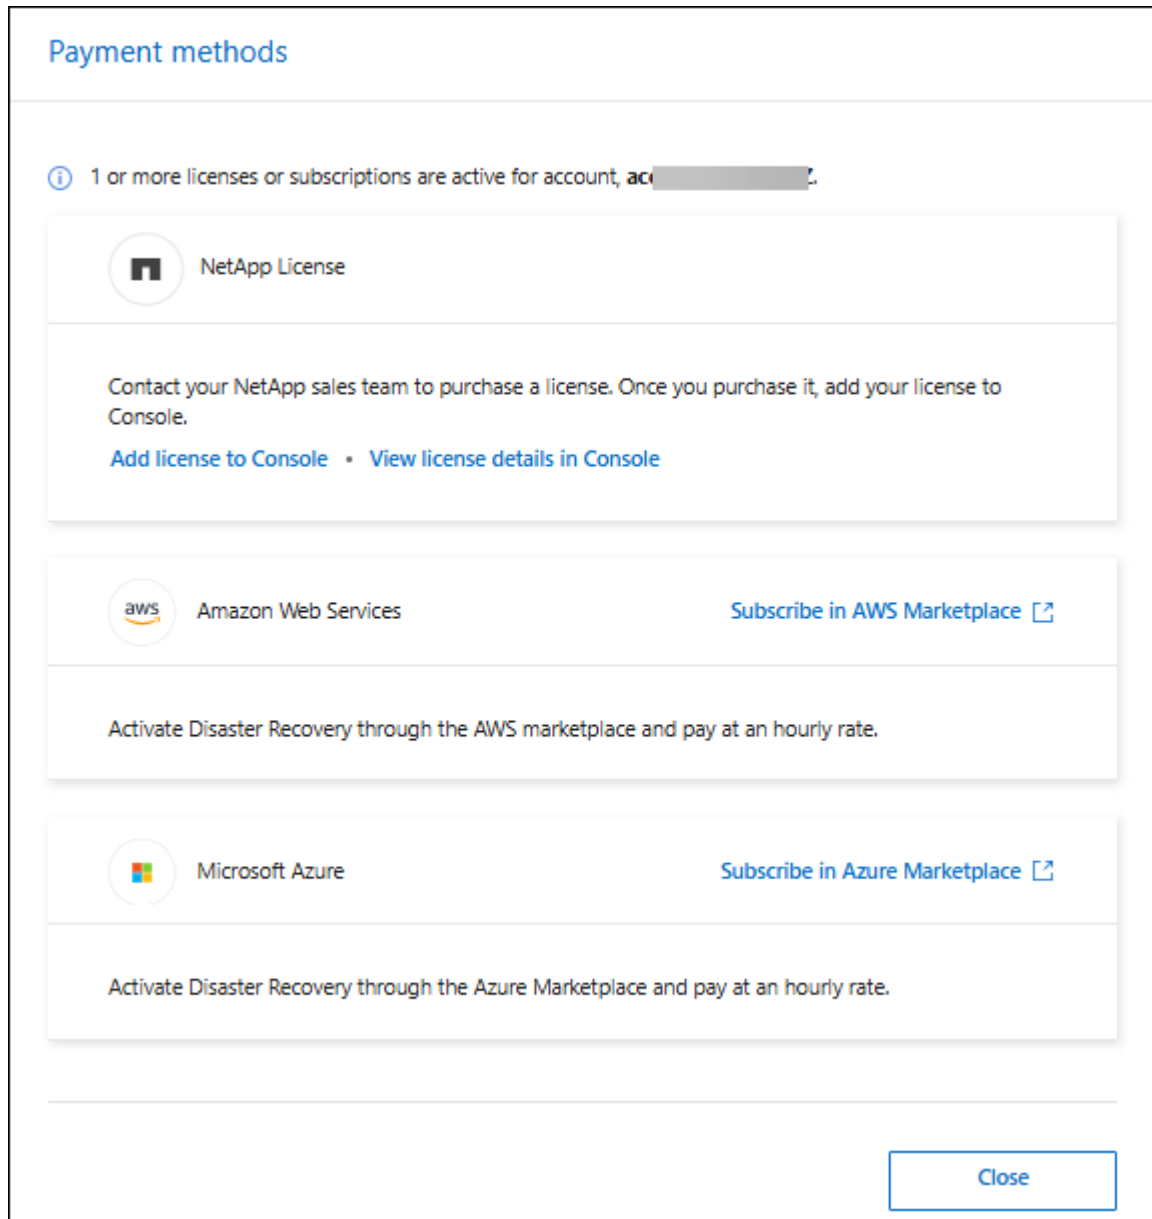
After the trial ends, subscribe through one of the Marketplaces

After the free trial ends, you can either purchase a license from NetApp or subscribe through AWS Marketplace or Microsoft Azure Marketplace. This procedure provides a high level overview of how to subscribe directly in one of the Marketplaces.

Steps

1. In the NetApp Disaster Recovery, you see a message that the free trial is expiring. In the message, select **Subscribe or purchase a license**.

Or, from the , select **View payment methods**.



2. Select **Subscribe in AWS Marketplace** or **Subscribe in Azure Marketplace**.
3. Use AWS Marketplace or Microsoft Azure Marketplace to subscribe to **NetApp Disaster Recovery**.
4. When you return to NetApp Disaster Recovery, a message states that you are subscribed.

You can view subscription details in the NetApp Console subscription page. [Learn more managing](#)

[subscriptions with the NetApp Console.](#)

After the trial ends, purchase a BYOL license through NetApp

After the trial ends, you can purchase a license through your NetApp Sales Rep.

If you bring your own license (BYOL), the set up includes purchasing the license, getting the NetApp License File (NLF), and adding the license to the NetApp Console.

Add the license to the NetApp Console*

After you've purchased your NetApp Disaster Recovery license from a NetApp Sales Rep, you can manage the license in the Console.

[Learn about adding licenses with the NetApp Console.](#)

Update your license when it expires

If your licensed term is nearing the expiration date, or if your licensed capacity is reaching the limit, you'll be notified in the NetApp Disaster Recovery UI. You can update your NetApp Disaster Recovery license before it expires so that there is no interruption in your ability to access your scanned data.



This message also appears in the NetApp Console and in [Notifications](#).

[Learn about updating licenses with the NetApp Console.](#)

End the free trial

You can stop the free trial at any time or you can wait until it expires.

Steps

1. In NetApp Disaster Recovery, select **Free trial - View details**.
2. In the drop-down details, select **End free trial**.

End free trial

Are you sure that you want to end your free trial on your account [redacted]to1? We will delete your data 60 days after you end your trial. If you subscribe or purchase a license within 60 days, we will retain your data. You may also delete your data immediately when you end your trial.

This action is not reversible.

☐ Delete data immediately after ending my free trial

Comments

Type "end trial" to end your free trial.

End

Cancel

3. If you want to delete all data, check **Delete data immediately after ending my free trial**.

This deletes all schedules, replication plans, resource groups, vCenters, and sites. Audit data, operation logs, and jobs history are retained until the end of the life of the product.



If you end the free trial, did not request to delete data, and don't purchase a license or subscription, then NetApp Disaster Recovery deletes all of your data 60 days after the free trial ends.

4. Type "end trial" in the text box.
5. Select **End**.

Use NetApp Disaster Recovery

Use NetApp Disaster Recovery overview

Using NetApp Disaster Recovery, you can accomplish the following goals:

- [View the health of your disaster recovery plans.](#)
- [Add vCenter sites.](#)
- [Create resource groups to organize VMs together](#)
- [Create a disaster recovery plan.](#)
- [Replicate VMware apps](#) on your primary site to a disaster recovery remote site in the cloud using SnapMirror replication.
- [Migrate VMware apps](#) on your primary site to another site.
- [Test the fail over](#) without disrupting the original virtual machines.
- In case of disaster, [fail over your primary site](#) to VMware Cloud on AWS with FSx for NetApp ONTAP.
- After the disaster has been resolved, [fail back](#) from the disaster recovery site to the primary site.
- [Monitor disaster recovery operations](#) on the Job Monitoring page.

View the health of your NetApp Disaster Recovery plans on the Dashboard

Using the NetApp Disaster Recovery Dashboard, you can determine the health of your disaster recovery sites and replication plans. You can quickly ascertain which sites and plans are healthy, disconnected, or degraded.

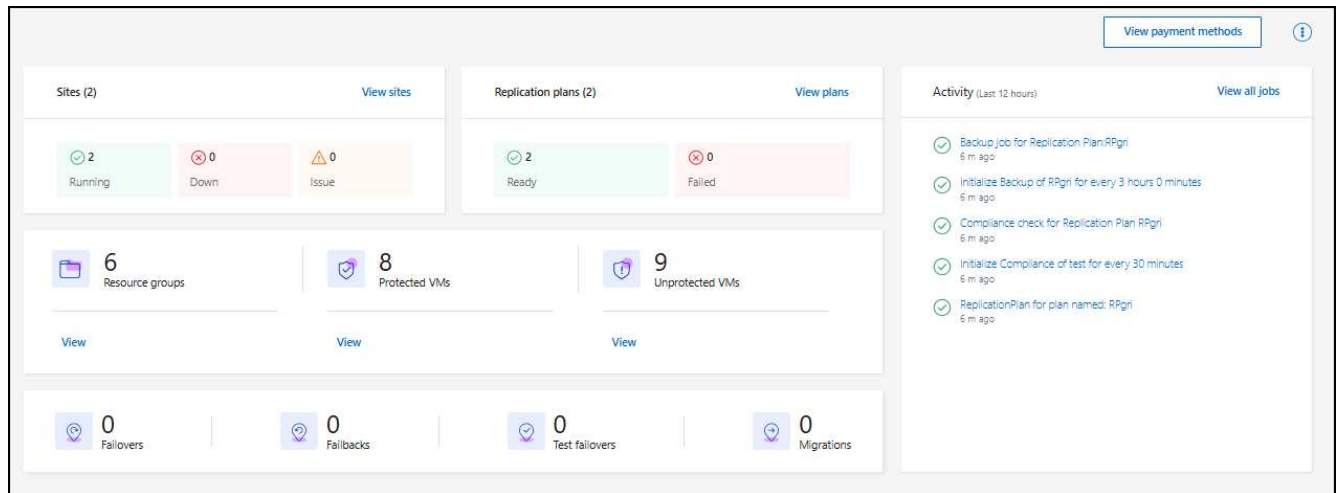
Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, Disaster recovery application admin, or Disaster recovery viewer role.

[Learn about user roles and permissions in NetApp Disaster Recovery.](#)
[Learn about NetApp Console access roles for all services.](#)

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.
3. From the NetApp Disaster Recovery menu, select **Dashboard**.



4. Review the following information on the Dashboard:

- **Sites:** View the health of your sites. A site can have one of the following statuses:

- **Running:** The vCenter is connected, healthy, and running.
- **Down:** The vCenter is not reachable or having connectivity issues.
- **Issue:** The vCenter is not reachable or having connectivity issues.

To see site details, select **View all** for a status or **View sites** to see them all.

- **Replication plans:** View the health of your plans. A plan can have one of the following statuses:

- **Ready**
- **Failed**

To review replication plan details, select **View all** for a status or **View replication plans** to see them all.

- **Resource groups:** View the health of your resource groups. A resource group can have one of the following statuses:
- **Protected VMs:** The VMs are part of a resource group.
- **Unprotected VMs:** The VMs are not part of a resource group.

To review details, select the **View** link below each.

- The number of failovers, test failovers, and migrations. For example, if you created two plans and migrated to the destinations, the migration count appears as "2."

5. Review all operations in the Activity pane. To view all operations on the Job Monitor, select **View all jobs**.

Add vCenters to a site in NetApp Disaster Recovery

Before you can create a disaster recovery plan, you need to add a primary vCenter server to a site and a target vCenter disaster recovery site in the NetApp Console.



Ensure that both the source and destination vCenters use the same NetApp Console agent.

After vCenters are added, NetApp Disaster Recovery performs a deep discovery of the vCenter environments,

including vCenter clusters, ESXi hosts, datastores, storage foot print, virtual machine details, SnapMirror replicas, and virtual machine networks.

Required NetApp Console role

Organization admin, Folder or project admin, or Disaster recovery admin.

[Learn about user roles and permissions in NetApp Disaster Recovery.](#)

[Learn about NetApp Console access roles for all services.](#)

About this task

If you added vCenters in previous releases and want to customize the discovery schedule, you must edit the vCenter server site and set the schedule.



NetApp Disaster Recovery performs discovery once every 24 hours. After setting up a site, you can later edit the vCenter to customize the discovery schedule that meets your needs. For example, if you have a large number of VMs, you can set the discovery schedule to run every 23 hours and 59 minutes. If you have a small number of VMs, you can set the discovery schedule to run every 12 hours. The minimum interval is 30 minutes and the maximum is 24 hours.

You should first perform a few manual discoveries to get the most up-to-date information about your environment. After that, you can set the schedule to run automatically.

If you have vCenters from earlier versions and want to change when discovery runs, edit the vCenter server site and set the schedule.

Newly added or deleted VMs are recognized in the next scheduled discovery or during an immediate manual discovery.

VMs can be protected only if the replication plan is in one of the following states:

- Ready
- Failback committed
- Test failover committed

vCenter clusters in a site

Each site contains one or more vCenters. These vCenters use one or more ONTAP storage clusters to host NFS or VMFS datastores.

A vCenter cluster can reside in only one site. You need the following information to add a vCenter cluster to a site:

- The vCenter management IP address or FQDN
- Credentials for a vCenter account with the required privileges to perform operations. See [required vCenter privileges](#) for more information.
- For cloud-hosted VMware sites, the required cloud access keys
- A security certificate to access your vCenter.



The service supports self-signed security certificates or certificates from a central certificate authority (CA).

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.

You'll land on NetApp Disaster Recovery Dashboard page. When you first start with the service, you need to add vCenter information. Later, the Dashboard displays data about your sites and replication plans.



Different fields appear depending on the type of site you are adding.

3. If some vCenter sites already exist and you want to add more, from the menu, select **Sites** and then select **Add**.
4. In the Sites page, select the site, and select **Add vCenter**.
5. **Source:** Select **Discover vCenter servers** to enter information about the source vCenter site.



If some vCenter sites already exist and you want to add more, from the top menu, select **Sites** and then select **Add**.

Add vCenter server

Enter connection details for the vCenter server that is accessible from the Console Agent.

Site	Console Agent
<input type="text" value="sit .gri2"/>	<input type="text" value="DRaaSTest"/>
vCenter IP address	Port
<input type="text" value=""/>	<input type="text" value="443"/>
vCenter user name	vCenter password
<input type="text" value="admin"/>	<input type="password" value="....."/>

☒ Use self-signed certificates

By default, vCenter discovery will run automatically once every 24 hours. This can be edited later. Discovery can also be triggered manually at any time.

- Select a site, select the NetApp Console agent, and provide vCenter credentials.
- (Applies only to on-premises sites) To accept self-signed certificates for the source vCenter, check the box.



Self-signed certificates are not as secure as other certificates. If your vCenter is **NOT** configured with certificate authority (CA) certificates, you should check this box; otherwise, the connection to the vCenter will not work.

6. Select **Add**.

Next, you will add a target vCenter.

7. Add a site again for the target vCenter.

8. Again, select **Add vCenter** and add target vCenter information.

9. **Target:**

a. Choose the target site and the location. If the target is cloud, select **AWS**.

- (Applies only to cloud sites) **API token**: Enter the API token to authorize service access for your organization. Create the API token by providing specific organization and service roles.
- (Applies only to cloud sites) **Long organization ID**: Enter the unique ID for the organization. You can identify this ID by clicking on the username in the Account section of the NetApp Console.

b. Select **Add**.

The source and target vCenters appear on the list of sites.

Sites (4)					
DemoOnPremSite_1					
a30C	17	5	6	h	
Healthy	VMs	Datastores	Resource groups	Agent	
DemoCloudSite_1					
vcntersd	11	3	0	hm	
Healthy	VMs	Datastores	Resource groups	Agent	

10. To see the progress of the operation, from the menu, select **Job monitoring**.

Add subnet mapping for a vCenter site

You can manage IP addresses on failover operations using subnet mapping, which enables you to add subnets for each vCenter. When you do so, you define the IPv4 CIDR, the default gateway, and the DNS for each virtual network.

Upon failover, NetApp Disaster Recovery uses the mapped network's CIDR to assign each vNIC a new IP address.

For example:

- NetworkA = 10.1.1.0/24
- NetworkB = 192.168.1.0/24

VM1 has a vNIC (10.1.1.50) that is connected to NetworkA.
NetworkA is mapped to NetworkB in the replication plan settings.

Upon failover, NetApp Disaster Recovery replaces the Network portion of the original IP address (10.1.1) and


keeps the host address (.50) of the original IP address (10.1.1.50). For VM1, NetApp Disaster Recovery looks at the CIDR settings for NetworkB and uses that the NetworkB network portion 192.168.1 while keeping the host portion (.50) to create the new IP address for VM1. The new IP becomes 192.168.1.50.

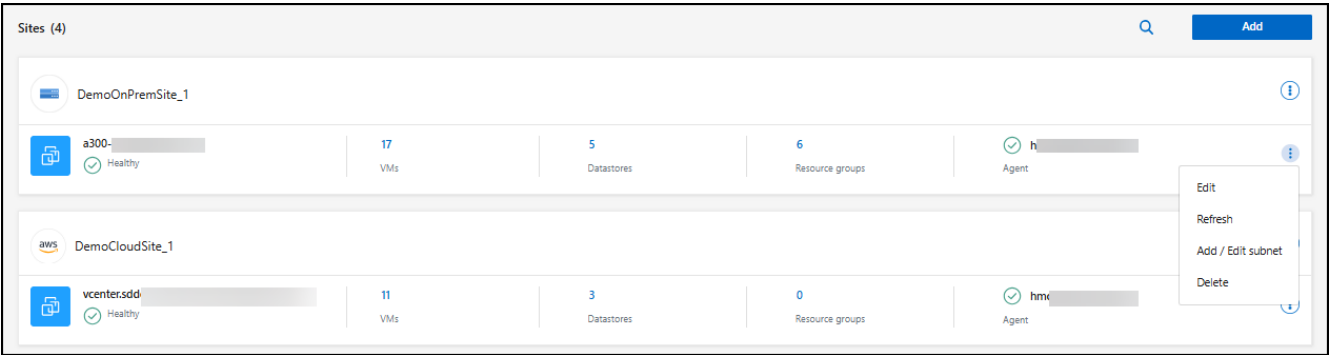
In summary, the host address stays the same, while the network address is replaced with whatever is configured in the site subnet mapping. This enables you to manage IP address reassignment upon failover more easily, especially if you have hundreds of networks and thousands of VMs to manage.

Using subnet mapping is an optional two-step process:

- First, add the subnet mapping for each vCenter site.
- Second, in the replication plan, indicate that you want to use subnet mapping in the Virtual Machines tab and Target IP field.

Steps

1. From the NetApp Disaster Recovery menu, select **Sites**.
2. From the Actions  icon on the right, select **Add subnet**.



The Configure subnet page appears:

Configure subnet

Network Name	Datacenter Name	Subnet	Gateway	DNS
mgmt_1_esxi98	Datacenter90_1	<input type="text" value="Enter CIDR format"/>	<input type="text" value="Enter Gateway"/>	<input type="text" value="Enter DNS"/>
mgmt_1_esxi92	Datacenter90_1	<input type="text" value="Enter CIDR format"/>	<input type="text" value="Enter Gateway"/>	<input type="text" value="Enter DNS"/>
VM Network	Datacenter90_1	<input type="text" value="Enter CIDR format"/>	<input type="text" value="Enter Gateway"/>	<input type="text" value="Enter DNS"/>
mgmt_1_esxi94	Datacenter90_1	<input type="text" value="Enter CIDR format"/>	<input type="text" value="Enter Gateway"/>	<input type="text" value="Enter DNS"/>
Mgmt_1_esxi91	Datacenter90_1	<input type="text" value="Enter CIDR format"/>	<input type="text" value="Enter Gateway"/>	<input type="text" value="Enter DNS"/>

1 - 5 of 12 << < 1 > >>

Add subnet mapping

Cancel

3. In the Configure subnet page, enter the following information:

- a. Subnet: Enter the IPv4 CIDR for the subnet up to /32.



CIDR notation is a method of specifying IP addresses and their network masks. The /24 denotes the netmask. The number consists of an IP address with the number after the "/" indicating how many bits of the IP address denote the network. For example, 192.168.0.50/24, the IP address is 192.168.0.50 and the total number of bits in the network address is 24. 192.168.0.50 255.255.255.0 becomes 192.168.0.0/24.

- b. Gateway: Enter the default gateway for the subnet.

- c. DNS: Enter the DNS for the subnet.

4. Select **Add subnet mapping**.

Select subnet mapping for a replication plan

When you create a replication plan, you can select the subnet mapping for the replication plan.

Using subnet mapping is an optional two-step process:

- First, add the subnet mapping for each vCenter site.
- Second, in the replication plan, indicate that you want to use subnet mapping.

Steps

1. From the NetApp Disaster Recovery menu, select **Replication plans**.
2. Select **Add** to add a replication plan.
3. Complete the fields in the usual way by adding the vCenter servers, selecting the resource groups or applications, and completing the mappings.
4. In the Replication plan > Resource mapping page, select the **Virtual machines** section.

Virtual machines

IP address type

Static

Target IP

Use subnet mapping

When a subnet exhausts its IP addresses, you cannot add more VMs to it. New VMs must connect to a different subnet with available IP addresses, which can be an existing alternative subnet or a newly created one.

☐

Use the same credentials for all VMs

☐

Use Windows LAPS

☐

Use the same script for all VMs

Target VM prefix

Optional

Target VM suffix

Optional

Preview: Sample VM name

5. In the **Target IP** field, select **Use subnet mapping** from the drop-down list.



If there are two VMs (for example, one is Linux and the other is Windows), credentials are needed only for Windows.

6. Continue with the creating the replication plan.


Edit the vCenter server site and customize the discovery schedule

You can edit the vCenter server site to customize the discovery schedule. For example, if you have a large number of VMs, you can set the discovery schedule to run every 23 hours and 59 minutes. If you have a small number of VMs, you can set the discovery schedule to run every 12 hours.

If you have vCenters from earlier versions and want to change when discovery runs, edit the vCenter server site and set the schedule.

If you don't want to schedule discovery, you can disable the scheduled discovery option and refresh the discovery manually at any time.

Steps

1. From the NetApp Disaster Recovery menu, select **Sites**.
2. Select the site you want to edit.
3. Select the Actions  icon on the right and select **Edit**.
4. In the Edit vCenter server page, edit the fields as needed.
5. To customize the discovery schedule, check the **Enable scheduled discovery** box and select the date and time interval you want.

Edit vCenter server

Enter connection details for the vCenter server that is accessible from the BlueXP Connector.

Site

Source

BlueXP Connector

SecLab_Connector_4

vCenter IP address

172.26.212.218

port

443

vCenter user name

vCenter password

☒ Use self-signed certificates ⓘ

☒ Enable scheduled discovery

Start discovery from

2025-04-02

 ⓘ

12

 :

00

AM

 ⓘ

Run discovery once every

23

 Hour(s)

59

 Minute(s)

Save

Cancel


6. Select **Save**.

Refresh discovery manually

You can refresh the discovery manually at any time. This is useful if you have added or removed VMs and want to update the information in NetApp Disaster Recovery.

Steps

1. From the NetApp Disaster Recovery menu, select **Sites**.
2. Select the site you want to refresh.
- 3.

Select the Actions  icon on the right and select **Refresh**.

Create a resource group to organize VMs together in NetApp Disaster Recovery

After adding vCenter sites, you can create resource groups to protect VMs by VM or datastore as a single unit. Resource groups enable you to organize a set of dependent VMs into logical groups that meet your requirements. For example, you might group VMs associated with one application or you might group applications that have similar tiers. As another example, groups could contain delayed boot orders that can be run upon recovery.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, or Disaster recovery application admin role.

[Learn about user roles and permissions in NetApp Disaster Recovery.](#)

[Learn about NetApp Console access roles for all services.](#)

About this task

You can group VMs themselves or VMs in datastores.

You can create resource groups using the following methods:

- From the Resource groups option
- While you're creating a disaster recovery or *replication plan*. If you have a lot of VMs hosted by a source vCenter cluster, it might be easier for you to create the resource groups while you're creating the replication plan. For instructions on creating resource groups while you're creating a replication plan, see [Create a replication plan](#).



Each resource group can include one or more VMs or datastores. The VMs will power on based on the sequence in which you include them in the replication plan. You can change the order by dragging the VMs or datastores up or down the resource group list.

About resource groups

Resource groups let you combine VMs or datastores to them as a single unit.

For example, a point-of-sale application might use several VMs for databases, business logic, and storefronts. You can manage all these VMs with one resource group. Set up resource groups to apply replication plan rules for VM startup order, network connection, and recovery of all VMs needed for the application.

How does it work?

NetApp Disaster Recovery protects VMs by replicating the underlying ONTAP volumes and LUNs hosting the VMs in the resource group. To do this, the system queries vCenter for the name of each data store hosting VMs in a resource group. NetApp Disaster Recovery then identifies the source ONTAP volume or LUN hosting that data store. All protection is performed at the ONTAP volume level using SnapMirror replication.

If VMs in the resource group are hosted on different data stores, NetApp Disaster Recovery uses one of the following methods to create a data-consistent snapshot of the ONTAP volumes or LUNs.

Relative location of FlexVol volumes	Snapshot replica process
Multiple data stores - FlexVol volumes in the same SVM	<ul style="list-style-type: none"> • ONTAP consistency group created • Snapshots of the consistency group taken • Volume-scoped SnapMirror replication performed
Multiple data stores - FlexVol volumes in multiple SVMs	<ul style="list-style-type: none"> • ONTAP API: <code>cg_start</code>. Quiesces all volumes so snapshots can be taken and initiates volume-scoped snapshots of all resource group volumes. • ONTAP API: <code>cg_end</code>. Resumes I/O on all volumes and enables volume-scoped SnapMirror replication after snapshots are taken.

When you create resource groups, consider the following issues:

- Before you add datastores to resource groups, start a manual discovery or a scheduled discovery of the VMs first. This ensures that the VMs are discovered and listed in the resource group. If you do not start a manual discovery, the VMs might not be listed in the resource group.
- Ensure that there is at least one VM in the datastore. If there are no VMs in the datastore, Disaster Recovery does not discover the datastore.
- A single datastore should not host VMs protected by more than one replication plan.
- Do not host protected and unprotected VMs on the same datastore. If protected and unprotected VMs are hosted on the same datastore, the following issues could arise:
 - Because NetApp Disaster Recovery uses SnapMirror and the system replicates entire ONTAP volumes, the used capacity of that volume is used for licensing considerations. In this case, the volume space consumed by both protected and unprotected VMs would be included in this calculation.
 - If the resource group and its associated datastores need to be failed over to the disaster recovery site, any unprotected VMs (VMs not part of the resource group, but hosted on the ONTAP volume) will no longer exist on the source site from the failover process, resulting in failure of unprotected VMs at the source site. Also, NetApp Disaster Recovery will not start those unprotected VMs at the failover vCenter site.
- To have a VM protected, it must be included in a resource group.

BEST PRACTICE: Organize your VMs before deploying NetApp Disaster Recovery to minimize “datastore sprawl.” Place VMs that need protection on a subset of datastores and place VMs that are not going to be protected on a different subset of datastores. Ensure that the VMs on any given datastore are not protected by different replication plans.

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.
3. From the NetApp Disaster Recovery menu, select **Resource groups**.
4. Select **Add**.
5. Enter a name for the resource group.
6. Select the source vCenter cluster where the VMs are located.
7. Select either **Virtual machines** or **Datastores** depending on how you want to search.

8. Select the **Add resource groups** tab. The system lists all datastores or VMs in the selected vCenter cluster. If you selected **Datastores**, the system lists all datastores in the selected vCenter cluster. If you selected **Virtual machines**, the system lists all VMs in the selected vCenter cluster.
9. On the left side of the Add resource groups page, select the VMs that you want to protect.

Add resource group

Name

DemoRG

vCenter

☒ Virtual machines

☐ Datastores

Select virtual machines

Search all datastores

☒ VMFS_Centos_vm1_ds4

☒ VMFS_Centos_vm1_ds5

☒ VMFS_RHEL_vm2_ds1

☐ VMFS_RHEL_vm2_ds2

☐ VMFS_RHEL_vm2_ds3

☐ VMFS_RHEL_vm2_ds4

☐ VMFS_RHEL_vm2_ds5

Selected VMs (3)

VMFS_Centos_vm1_ds4

VMFS_Centos_vm1_ds5

VMFS_RHEL_vm2_ds1

Add

Cancel

Add resource group

Name:

vCenter:

☐ Virtual machines ☒ Datastores

Select datastores

Search datastores

- ☐ DS4_auto_vmfs_6d7
- ☐ DS2_auto_vmfs_6d7
- ☐ DS1_surya_nfs_scale
- ☒ DS4_auto_nfs_450
- ☒ DS3_auto_nfs_450
- ☐ DS1_auto_nfs_450
- ☐ DS2_auto_nfs_450

Selected datastores (2)

- DS4_auto_nfs_450 X
- DS3_auto_nfs_450 X

10. Optionally, change the order of the VMs on the right by dragging each VM up or down the list. The VMs will power on based on the sequence in which you include them.

11. Select **Add**.

Create a replication plan in NetApp Disaster Recovery

After you've added vCenter sites, you're ready to create a disaster recovery or *replication plan*. Replication plans manage the data protection of the VMware infrastructure. Select the source and destination vCenters, pick the resource groups, and group how applications should be restored and powered on. For example, you might group virtual machines (VMs) associated with one application or you might group applications that have similar tiers. Such plans are sometimes called *blueprints*.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, Disaster recovery failover admin, or Disaster recovery application admin role.

[Learn about user roles and permissions in NetApp Disaster Recovery.](#)

[Learn about NetApp Console access roles for all services.](#)

About this task

You can create a replication plan and also edit schedules for compliance and testing. Run test failovers of VMs without affecting production workloads.

You can protect multiple VMs on multiple datastores. NetApp Disaster Recovery creates ONTAP Consistency Groups for all ONTAP volumes hosting protected VM datastores.

VMs can be protected only if the replication plan is in one of the following states:


- Ready
- Failback committed
- Test failover committed

Replication plan snapshots

Disaster Recovery maintains the same number of snapshots on the source and destination clusters. By default, the service performs a snapshot reconciliation process every 24 hours to ensure that the number of snapshots on the source and destination clusters is the same.

The following situations can cause the number of snapshots to differ between the source and destination clusters:

- Some situations can cause ONTAP operations outside of Disaster Recovery to add or remove snapshots from the volume:
 - If there are missing snapshots on the source site, the corresponding snapshots on the destination site might be deleted, depending on the default SnapMirror policy for the relationship.
 - If there are missing snapshots on the destination site, the service might delete the corresponding snapshots on the source site during the next scheduled snapshot reconciliation process, depending on the default SnapMirror policy for the relationship.
- A reduction of the replication plan's snapshot retention count can cause the service to delete the oldest snapshots on both the source and destination sites to meet the newly reduced retention number.

In these cases, Disaster Recovery removes older snapshots from the source and destination clusters upon the next consistency check. Or, the administrator can perform an immediate snapshot cleanup by selecting the **Actions**  icon on the replication plan and selecting **Clean up snapshots**.

The service performs snapshot symmetry checks every 24 hours.

Before you begin

Before creating a SnapMirror relationship, set up the cluster and SVM peering outside of Disaster Recovery.

BEST PRACTICE: Organize your VMs before deploying NetApp Disaster Recovery to minimize “datastore sprawl.” Place VMs that need protection on a subset of datastores and place VMs that are not going to be protected on a different subset of datastores. Use datastore-based protection to ensure that the VMs on any given datastore are protected.

Create the plan

A wizard takes you through these steps:

- Select vCenter servers.
- Select the VMs or datastores that you want to replicate and assign resource groups.
- Map how resources from the source environment map to the destination.
- Set how often the plan runs, run a guest-hosted script, set the boot order, and select the recovery point

objective.

- Review the plan.

When you create the plan, you should follow these guidelines:

- Use the same credentials for all VMs in the plan.
- Use the same script for all VMs in the plan.
- Use the same subnet, DNS, and gateway for all VMs in the plan.

Select vCenter servers

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

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.
3. From the NetApp Disaster Recovery menu, select **Replication plans** and select **Add**. Or, if you are just beginning to use the service, from the Dashboard, select **Add replication plan**.

Add replication plan 1 vCenter servers 2 Applications 3 Resource mapping 4 Review

Replication plan > Add plan

vCenter servers
Provide the plan name and select the source and target vCenter servers.

Replication plan name
RPgr4

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

Source vCenter
a3C

Target vCenter
vcenter.sdd

Replicate

Cancel Next

4. Create a name for the replication plan.

5. Select the source and target vCenters from the Source and Target vCenter lists.
6. Select **Next**.

Select applications to replicate and assign resource groups

The next step is to group the required VMs or datastores into functional resource groups. Resource groups enable you to protect a set of VMs or datastores with a common snapshot.

When you select applications in the replication plan, you can see the operating system for each VM or datastore in the plan. This is helpful in deciding how to group VMs or datastores together in a resource group.



Each resource group can include one or more VMs or datastores.

When you create resource groups, consider the following issues:

- Before you add datastores to resource groups, start a manual discovery or a scheduled discovery of the VMs first. This ensures that the VMs are discovered and listed in the resource group. If you do not trigger a manual discovery, the VMs might not be listed in the resource group.
- Ensure that there is at least one VM in the datastore. If there are no VMs in the datastore, the datastore will not be discovered.
- A single datastore should not host VMs protected by more than one replication plan.
- Do not host protected and unprotected VMs on the same datastore. If protected and unprotected VMs are hosted on the same datastore, the following issues could arise:
 - Because NetApp Disaster Recovery uses SnapMirror and the system replicates entire ONTAP volumes, the used capacity of that volume is used for licensing considerations. In this case, the volume space consumed by both protected and unprotected VMs would be included in this calculation.
 - If the resource group and its associated datastores need to be failed over to the disaster recovery site, any unprotected VMs (VMs not part of the resource group, but hosted on the ONTAP volume) will no longer exist on the source site from the failover process, resulting in failure of unprotected VMs at the source site. Also, NetApp Disaster Recovery will not start those unprotected VMs at the failover vCenter site.
- To have a VM protected, it must be included in a resource group.

BEST PRACTICE: Create a separate dedicated set of mappings for your failover tests to prevent VMS from being connected to production networks using the same IP addresses.

Steps

1. Select **Virtual machines** or **Datastores**.
2. Optionally search for specific VM or datastore by name.
3. On the left side of the Applications page, select the VMs or datastores that you want to protect and assign to the selected group.

The source vCenter must reside on the on-premises vCenter. The target vCenter can be a second on-premises vCenter in the same site or a remote site, or a cloud-based software defined data center (SDDC) such as VMware Cloud on AWS. Both vCenters should already be added to your BlueXP disaster recovery working environment.

The selected resource is automatically added to group 1 and a new group 2 is started. Each time you add a resource to the last group, another group is added.

☐ Resource groups
☒ Virtual machines
☐ Datastores

Datastore

All datastores

☐ Select all VMs in view (100)

VMs in view: 100/703

☐ Pavan_windows19_vm3_vmfs_DS3

☐ Pavan_windows19_vm3_vmfs_ds4

☐ SQLServer

☒ VMFS_Centos_vm1_ds2

☒ VMFS_Centos_vm1_ds3

☒ VMFS_Centos_vm1_ds4

View more VMs

Selected VMs to replicate.

Selected VMs (3)

DemoPlan_ResourceGroup1 (2)

VMFS_Centos_vm1_ds2

VMFS_Centos_vm1_ds3

DemoPlan_ResourceGroup2 (1)

VMFS_Centos_vm1_ds4

DemoPlan_ResourceGroup3 (0)

Previous

Next

Or, for datastores:

☐ Resource groups
☐ Virtual machines
☒ Datastores

☐ DS3_auto_vmfs_6d7

☐ DS1_auto_vmfs_6d7

☒ DS4_auto_vmfs_6d7

☐ DS2_auto_vmfs_6d7

☐ DS1_surya_nfs_scale

☒ DS4_auto_nfs_450

☐ DS3_auto_nfs_450

Selected datastores to replicate.

Selected datastores (2)

DemoPlan_ResourceGroup1 (1)

DS4_auto_nfs_450

DemoPlan_ResourceGroup2

DS4_auto_vmfs_6d7


DemoPlan_ResourceGroup4 (0)

Drag datastores to regroup.

Previous

Next

4. Optionally, do any of the following:

- To change the group's name, click the group **Edit**  icon.
- To remove a resource from a group, select **X** next to the resource.
- To move a resource to a different group, drag and drop it into the new group.

To move a datastore to a different resource group, unselect the unwanted datastore and submit the replication plan. Then, create or edit the other replication plan and reselect the datastore.

5. Select **Next**.

53

Map source resources to the target

In the Resource mapping step, specify how the resources from the source environment should map to the target. When you create a replication plan, you can set a boot delay and order for each VM in the plan. This enables you to set a sequence for the VMs to start.

If you plan to perform test failovers as part of your DR plan, you should provide a set of test failover mappings to ensure that VMs started during the failover test don't interfere with production VMs. You can accomplish this by either providing test VMs with different IP addresses, or by mapping the virtual NICs of the test VMs to a different network that is isolated from production yet has the same IP configuration (referred to as a *bubble* or *test network*).

Before you begin

If you want to create a SnapMirror relationship in this service, the cluster and its SVM peering should have already been set up outside of NetApp Disaster Recovery.

Steps

1. In the Resource mapping page, to use the same mappings for both failover and test operations, check the box.
2. In the Failover mappings tab, select the down arrow to the right of each resource and map the resources in each section:
 - Compute resources
 - Virtual networks
 - Virtual machines
 - Datastores

Map resources > Compute resources section

The Compute resources section defines where VMs will be restored after a failover. Map the source vCenter data center and cluster to a target data center and cluster.

Optionally, VMs can be restarted on a specific vCenter ESXi host. If VMWare DRS is enabled, you can move the VM to an alternate host automatically if needed to meet the DR configured policy.

Optionally, you can place all VMs in this replication plan into a unique folder with the vCenter. This provides an easy way to quickly organize failed over VMs within the vCenter.

Select the down arrow next to **Compute resources**.

- **Source and target datacenters**
- **Target cluster**
- **Target host** (optional): After you select the cluster, you can then set this information.



If a vCenter has a Distributed Resource Scheduler (DRS) configured to manage multiple hosts in a cluster, you don't need to select a host. If you select a host, NetApp Disaster Recovery will place all the VMs on the selected host.

* **Target VM folder** (optional): Create a new root folder to store the selected VMs.

Map resources > Virtual networks section

VMs use virtual NICs connected to virtual networks. In the failover process, the service connects these virtual NICs to virtual networks defined in the destination VMware environment. For each source virtual network used by the VMs in the resource group, the service requires a destination virtual network assignment.



You can assign multiple source virtual networks to the same target virtual network. This might however create IP network configuration conflicts. You can map multiple source networks to a single target network to ensure that all source networks have the same configuration.

In the Failover mappings tab, select the down arrow next to **Virtual networks**. Select the source virtual LAN and target virtual LAN.

Select the network mapping to the appropriate virtual LAN. The virtual LANs should already be provisioned, so select the appropriate virtual LAN to map the VM.

Map resources > Virtual machines section

You can configure each VM in the resource group protected by the replication plan to suit the destination vCenter virtual environment by setting any of the following options:

- The number of virtual CPUs
- The amount of virtual DRAM
- The IP address configuration
- The ability to execute guest-OS shell scripts as part of the failover process
- The ability to change failed over VM names by using a unique prefix and suffix
- The ability to set the restart order during VM failover

In the Failover mappings tab, select the down arrow next to **Virtual machines**.

The default for the VMs is mapped. Default mapping uses the same settings that the VMs use in the production environment (same IP address, subnet mask, and gateway).

If you make any changes from the default settings, you must change the Target IP field to "Different from source."



If you change settings to "Different from source," you need to provide VM guest OS credentials.

This section might display different fields depending on your selection.

You can increase or decrease the number of virtual CPUs assigned to each failed over VM. However, each VM requires at least one virtual CPU. You can change the number of virtual CPUs and virtual DRAM assigned to each VM. The most common reason why you might want to change the default virtual CPU and virtual DRAM settings is if the target vCenter cluster nodes do not have as many available resources as the source vCenter cluster.

Network settings

Disaster Recovery supports an extensive set of configuration options for VM networks. Changing these might be required if the target site has virtual networks that use different TCP/IP settings as the production virtual networks on the source site.

At the most basic (and default) level, the settings simply use the same TCP/IP network settings for each VM on

the destination site as used on the source site. This requires that you configure the same TCP/IP settings on the source and destination virtual networks.

The service supports network settings of static or Dynamic Host Configuration Protocol (DHCP) IP configuration for VMs. DHCP provides a standards-based method of dynamically configuring the TCP/IP settings of a host network port. DHCP must provide, at a minimum, a TCP/IP address, and can also provide a default gateway address (for routing to an external internet connection), a subnet mask, and a DNS server address. DHCP is commonly used for end-user computing devices such as employee desktop, laptop, and mobile phone connections, while it can also be used for any networking computing device such as servers.

- **Use the same subnet mask, DNS, and gateway settings** option: Because these settings are typically the same for all VMs connected to the same virtual networks, you might find it easier to configure these once and let Disaster Recovery use the settings for all VMs in the resource group protected by the replication plan. If some VMs use different settings, you need to uncheck this box and provide those settings for each VM.
- **IP address type**: Reconfigure the VMs configuration to match the target virtual network requirements. NetApp Disaster Recovery offers two options: DHCP or static IP. For static IPs, configure the subnet mask, gateway, and DNS servers. Additionally, enter credentials for VMs.
 - **DHCP**: Select this setting if you want your VMs to obtain network configuration information from a DHCP server. If you choose this option, you provide just the credentials for the VM.
 - **Static IP**: Select this setting if you want to specify IP configuration information manually. You can select one of the following: same as source, different from source, or subnet mapping. If you choose the same as the source, you do not need to enter credentials. On the other hand, if you choose to use different information from the source, you can provide the credentials, IP address of the VM, subnet mask, DNS, and gateway information. VM guest OS credentials should be supplied to either the global level or at each VM level.

This can be very helpful when recovering large environments to smaller target clusters or for conducting disaster recovery tests without having to provision a one-to-one physical VMware infrastructure.

Virtual machines

IP address type: Static Target IP: Use subnet mapping

① When a subnet exhausts its IP addresses, you cannot add more VMs to it. New VMs must connect to a different subnet with available IP addresses, which can be an existing alternative subnet or a newly created one.

☐ Use the same credentials for all VMs

☒ Use Windows LAPS ①

Domain controller: WIN-DLF9SSVRCR3 Account name: draasanf\administrator Password: Required

Domain: draasanf.csjad.com

☐ Use the same script for all VMs

Target VM prefix: Optional Target VM suffix: Optional Preview: Sample VM name

- **Scripts:** You can include custom guest-OS hosted scripts in .sh, .bat, or .ps1 format as post processes. With custom scripts, BlueXP disaster recovery can run your script after a failover, failback, and migrate processes. For example, you can use a custom script to resume all database transactions after the failover is complete. The service can run scripts within VMs running Microsoft Windows or any supported Linux variant with command-line parameters supported. You can assign a script to individual VMs or to all VMs in the replication plan.

To enable script execution with the VM guest OS, the following conditons must be met:

- VMware Tools must be installed on the VM.
- Appropriate user credentials must be provided with adequate guest OS privileges to run the script.
- Optionally, include a timeout value in seconds for the script.

VMs running Microsoft Windows: can run either Windows batch (.bat) or PowerShell (ps1) scripts. Windows scripts can use command-line arguments. Format each argument in the `arg_name$value` format, where `arg_name` is the name of the argument and `$value` is the value of the argument and a semi-colon separates each `argument$value` pair.

VMs running Linux: can run any shell script (.sh) supported by the version of Linux used by the VM. Linux scripts can use command-line arguments. Provide arguments in a list of values separated by semi-colons. Named arguments are not supported. Add each argument to the `Arg[x]` argument list and reference each value using a pointer into the `Arg[x]` array, for example, `value1;value2;value3`.

- **Target VM prefix and suffix:** Under the virtual machines details, you can optionally add a prefix and suffix to each failed over VM name. This can be helpful in differentiating the failed over VMs from the production VMs running on the same vCenter cluster. For example, you can add a prefix of "DR-" and a suffix of "-failover" to the VM name. Some people add a second production vCenter to host VMs temporarily at a

different site in the event of a disaster. Adding a prefix or suffix can help you quickly identify failed over VMs. You can also use the prefix or suffix in custom scripts.

You can use the alternative method of setting the Target VM folder in the Compute resources section.

- **Source VM CPU and RAM:** Under the virtual machines details, you can optionally resize the VM CPU and RAM parameters.



You can configure DRAM either in gigabytes (GiB) or megabytes (MiB). While each VM requires at least one MiB of RAM, the actual amount must ensure that the VM guest OS and any running applications can operate efficiently.

Source VM	Operating system	CPUs	RAM (GB)	Boot order	Boot delay (mins)	Create application-consistent replicas	Scripts	Credentials
Resource group 1								
SQL_PRD_1	Linux	4	16	1	0	<input checked="" type="checkbox"/>	None	Required
Resource group 2								
SQL_PRD_2	Linux	4	32	2	0	<input checked="" type="checkbox"/>	file.py, +2	Required
SQL_PRD_3	Linux	8	64	3	0	<input checked="" type="checkbox"/>	sql_dr_prod.py	Provided
SQL_PRD_4	Linux	8	64	4	0	<input checked="" type="checkbox"/>	sql_dr_prod.py	Provided
SQL_PRD_5	Linux	8	64	5	0	<input checked="" type="checkbox"/>	sql_dr_prod.py	Provided
SQL_PRD_6	Linux	8	64	6	0	<input checked="" type="checkbox"/>	sql_dr_prod.py	Provided
Datstores		Mapped						

- **Boot order:** You can modify the boot order after a failover for all the selected virtual machines across the resource groups. By default, all VMs boot together in parallel; however, you can make changes at this stage. This is helpful to ensure that all your priority one VMs are running before subsequent priority VMs are started.

BlueXP disaster recovery boots any VMs with the same boot order number in parallel.

- Sequential boot: Assign each VM a unique number to boot the in the assigned order, for example, 1,2,3,4,5.
- Simultaneous boot: Assign the same number to any VMs to boot them at the same time, for example, 1,1,1,1,2,2,3,4,4.
- **Boot delay:** Adjust the delay in minutes of the boot up action, indicating the amount of time that the VM will wait before it starts the power-on process. Enter a value from 0 to 10 minutes.



To reset the boot order to the default, select **Reset VM settings to default** and then choose which settings you want to change back to the default.

- **Create application-consistent replicas:** Indicate whether to create application-consistent snapshot copies. The service will quiesce the application and then take a snapshot to get a consistent state of the application. This feature is supported with Oracle running on Windows and Linux and SQL Server running on Windows. See more details next.
- **Use Windows LAPS:** If you are using Windows Local Administrator Password Solution (Windows LAPS), check this box. This option is available only if you have selected the **Static IP** option. When you check this box, you do not need to provide a password for each of your virtual machines. Instead, you provide the domain controller details.

If you do not use Windows LAPS, then the VM is a Windows VM and the credentials option on the VM row is enabled. You can provide the credentials for the VM.

Disaster recovery
Add replication plan

✓ vCenter servers ✓ Applications 3 Resource mapping 4 Recurrence 5 Review

DHCP

☐ Use the same credentials for all VMs
☐ Use the same scripts for all VMs

Q

Source VM	Operating system	CPUs	RAM (GB)	Boot order	Boot delay (mins)	Create application-consistent replicas	Scripts	Credentials
Resource group 1								
SQL_PRD_1	Linux	4	16	1	0	<input checked="" type="checkbox"/>	None	Required
Resource group 2								
SQL_PRD_2	Linux	4	32	2	0	<input checked="" type="checkbox"/>	file.py, +2	Required
SQL_PRD_3	Linux	8	64	3	0	<input checked="" type="checkbox"/>	sql_dr_prod.py	Provided
SQL_PRD_4	Linux	8	64	4	0	<input checked="" type="checkbox"/>	sql_dr_prod.py	Provided
SQL_PRD_5	Linux	8	64	5	0	<input checked="" type="checkbox"/>	sql_dr_prod.py	Provided
SQL_PRD_6	Linux	8	64	6	0	<input checked="" type="checkbox"/>	sql_dr_prod.py	Provided

Datastores ☒ Mapped

Previous Next

Create application-consistent replicas

Many VMs host database servers such as Oracle or Microsoft SQL Server. These database servers require application-consistent snapshots to ensure that the database is in a consistent state when the snapshot is taken.

Application-consistent snapshots ensure that the database is in a consistent state when the snapshot is taken. This is important because it ensures that the database can be restored to a consistent state after a failover or fallback operation.

The data managed by the database server might be hosted on the same datastore as the VM hosting the database server, or it might be hosted on a different datastore. The following table shows the supported configurations for application-consistent snapshots in Disaster Recovery:

Data location	Supported	Notes
Within the same vCenter datastore as the VM	Yes	Because the database server and database both reside on the same datastore, both the server and the data will be in sync upon failover.
Within a different vCenter datastore from the VM	No	<p>Disaster Recovery can't identify when a database server's data is on a different vCenter datastore. The service can't replicate the data, but can replicate the database server VM.</p> <p>While the database data cannot be replicated, the service ensures that the database server performs all necessary steps to ensure that the database is quiesced at the time of the VM backup.</p>
Within an external data source	No	<p>If the data resides on a guest-mounted LUN or NFS share, Disaster Recovery can't replicate the data, but can replicate the database server VM.</p> <p>While the database data cannot be replicated, the service ensures that the database server performs all necessary steps to ensure that the database is quiesced at the time of the VM backup.</p>

During a scheduled backup, Disaster Recovery quiesces the database server and then takes a snapshot of the VM hosting the database server. This ensures that the database is in a consistent state when the snapshot is taken.

- For Windows VMs, the service uses the Microsoft Volume Shadow Copy Service (VSS) to coordinate with either database server.
- For Linux VMs, the service uses a set of scripts to place the Oracle server in backup mode.

To enable application-consistent replicas of the VMs and their hosting datastores, check the box next to **Create application-consistent replicas** for each VM and provide guest login credentials with the appropriate privileges.

Map resources > Datastores section

VMware datastores are hosted on ONTAP FlexVol volumes, or ONTAP iSCSI or FC LUNs using VMware VMFS. Use the Datastores section to define the target ONTAP cluster, storage virtual machine (SVM), and volume or LUN to replicate the on-disk data to the destination.

Select the down arrow next to **Datastores**. Based on the selection of VMs, datastore mappings are automatically selected.

This section might be enabled or disabled depending on your selection.

Datastores

☒ Use platform managed backups and retention schedules ⓘ

Start running retention from
2025-05-13
12 : 00 AM ⓘ

Run retention once every
03 Hour(s) 00 Minute(s)

Retention count for all datastores ⓘ
30

Source datastore
DS_Testing_Staging (Temp_3510_N1:DR_Vol_Staging)

Target datastore
DS_Testing_Staging (test:DR_Vol_Staging_dest)

Preferred NFS LIF
Select preferred NFS LIF

Export policy
Select export policy

- **Use platform managed backups and retention schedules:** If you are using an external snapshot management solution, check this box. NetApp Disaster Recovery supports the use of external snapshot management solutions such as the native ONTAP SnapMirror policy scheduler or third-party integrations. If every datastore (volume) in the replication plan already has a SnapMirror relationship that is being managed elsewhere, you can use those snapshots as recovery points in NetApp Disaster Recovery.

When this option selected, NetApp Disaster Recovery does not configure a backup schedule. However, you still need to configure a retention schedule because snapshots might still be taken for testing, failover, and failback operations.

After this is configured, the service doesn't take any regularly scheduled snapshots, but instead relies on the external entity to take and update those snapshots.

- **Start time:** Enter the date and time when you want backups and retention to start running.
- **Run interval:** Enter the time interval in hours and minutes. For example, if you enter 1 hour, the service will take a snapshot every hour.
- **Retention count:** Enter the number of snapshots you want to retain.



The number of snapshots retained along with the data change rate between each snapshot determines the amount of storage space consumed on both the source and destination. The more snapshots you retain, the more storage space is consumed.

- **Source and Target datastores:** If multiple (fan-out) SnapMirror relationships exist, you can select the destination to use. If a volume has a SnapMirror relationship already established, the corresponding source and target datastores appear. If a volume that does not have a SnapMirror relationship, you can create one now by selecting a target cluster, selecting a target SVM, and providing a volume name. The service will create the volume and SnapMirror relationship.



If you want to create a SnapMirror relationship in this service, the cluster and its SVM peering should have already been set up outside of NetApp Disaster Recovery.

- If the VMs are from same volume and same SVM, then the service performs a standard ONTAP snapshot and updates the secondary destinations.
- If the VMs are from different volume and same SVM, the service creates a consistency group snapshot by including all the volumes and updates the secondary destinations.

- If the VMs are from different volume and different SVM, the service performs a consistency group start phase and commit phase snapshot by including all the volumes in the same or different cluster and updates the secondary destinations.
- During the failover, you can select any snapshot. If you select the latest snapshot, the service creates an on-demand backup, updates the destination, and uses that snapshot for the failover.
- **Preferred NFS LIF and Export policy:** Typically, let the service select the preferred NFS LIF and export policy. If you want to use a specific NFS LIF or export policy, select the down arrow next to each field and select the appropriate option.

You can optionally use specific data interfaces (LIFs) for a volume after a failover event. This is useful for data traffic balancing if the target SVM has multiple LIFs.

For additional control over NAS data access security, the service can assign different datastore volumes specific NAS export policies. Export policies define the access control rules for NFS clients that access the datastore volumes. If you don't specify an export policy, the service uses the default export policy for the SVM.

BEST PRACTICE: We strongly recommend that you create a dedicated export policy that limits volume access only to the source and destination vCenter ESXi hosts that will host the protected VMs. This helps to ensure that external entities can't gain access to the NFS export.

Add test failover mappings

Steps

1. To set different mappings for the test environment, uncheck the box and select the **Test mappings** tab.
2. Go through each tab as before, but this time for the test environment.

On the Test mappings tab, the Virtual machines and Datastores mappings are disabled.



You can later test the entire plan. Right now, you are setting up the mappings for the test environment.

Review the replication plan

Finally, take a few moments to review the replication plan.



You can later disable or delete the replication plan.

Steps

1. Review information in each tab: Plan Details, Failover Mapping, and VMs.
2. Select **Add plan**.

The plan is added to the list of plans.

Edit schedules to test compliance and ensure failover tests work

You might want to set up schedules to test compliance and failover tests so that you ensure that they will work correctly should you need them.

- **Compliance time impact:** When a replication plan is created, the service creates a compliance schedule by default. The default compliance time is 30 minutes. To change this time, you can use edit the schedule in the replication plan.
- **Test failover impact:** You can test a failover process on demand or by a schedule. This lets you test the failover of virtual machines to a destination that is specified in a replication plan.

A test failover creates a FlexClone volume, mounts the datastore, and moves the workload on that datastore. A test failover operation does *not* impact production workloads, the SnapMirror relationship used on the test site, and protected workloads that must continue to operate normally.

Based on the schedule, the failover test runs and ensures that workloads are moving to the destination specified by the replication plan.

Steps

1. From the NetApp Disaster Recovery menu, select **Replication plans**.

Name	Compliance status	Plan status	Protected site	Resource groups	Failover site
RP_DRAAS	Healthy	Ready	DemoOnPremSite_1	RG2, RG1, RG4	DemoCloudSite_1
RPgri	Healthy	Ready	DemoOnPremSite_1	rggri1	DemoCloudSite_1
rpgr3	Healthy	Ready	site-onprem-gri2	rpgr3_ResourceGroup1	DemoOnPremSite_1

2. Select the **Actions** icon and select **Edit schedules**.
3. Enter how frequently in minutes that you want NetApp Disaster Recovery to check test compliance.
4. To check that your failover tests are healthy, check **Run failovers on a monthly schedule**.
 - a. Select the day of the month and time you want these tests to run.
 - b. Enter the date in yyyy-mm-dd format when you want the test to start.

Edit schedules: RP_DRAAS

Compliance checks and test failovers run on a recurring basis. Enter how often these actions should occur.

Compliance check

Frequency (min) ?

30

Test failover

☒ Run test failovers on a schedule ?

☒ Use on-demand snapshot for scheduled test failover

Repeat

Daily

Hour : Minute AM/PM Start date ?

12 : 00 AM 2025-05-13

☒ Automatically cleanup minutes after test failover ?

Save **Cancel**

5. **Use ondemand snapshot for scheduled test failover:** To take a new snapshot before initiating the automated test failover, check this box.
6. To clean up the test environment after the failover test finishes, check **Automatically clean up after test failover** and enter the number of minutes you want to wait before the cleanup starts.



This process unregisters the temporary VMs from the test location, deletes the FlexClone volume that was created, and unmounts the temporary datastores.

7. Select **Save**.

Replicate applications to another site with NetApp Disaster Recovery

Using NetApp Disaster Recovery, you can replicate VMware apps on your source site to a disaster recovery remote site in the cloud using SnapMirror replication.



After you create the disaster recovery plan, identify the recurrence in the wizard, and initiate a replication to a disaster recovery site, every 30 minutes NetApp Disaster Recovery verifies that the replication is actually occurring according to the plan. You can monitor the progress in the Job Monitor page.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, or Disaster recovery failover admin role.


[Learn about user roles and permissions in NetApp Disaster Recovery.](#)

[Learn about NetApp Console access roles for all services.](#)

Before you begin

Before you initiate the replication, you should have created a replication plan and selected to replicate the apps. Then, the **Replicate** option appears in the Actions menu.

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.
3. From the menu, select **Replication plans**.
4. Select the replication plan.
5. On the right, select the **Actions** option  and select **Replicate**.

Migrate applications to another site with NetApp Disaster Recovery

Using NetApp Disaster Recovery, you can migrate VMware apps on your source site to another site.




After you create the replication plan, identify the recurrence in the wizard, and initiate the migration, every 30 minutes NetApp Disaster Recovery verifies that the migration is actually occurring according to the plan. You can monitor the progress in the Job Monitor page.

Before you begin

Before you initiate the migration, you should have created a replication plan and selected to migrate the apps. Then, the **Migrate** option appears in the Actions menu.

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.
3. From the menu, select **Replication plans**.
4. Select the replication plan.
5. On the right, select the **Actions** option  and select **Migrate**.

Fail over applications to a remote site with NetApp Disaster Recovery

In case of a disaster, fail over your primary on-premises VMware site to another on-premises VMware site or VMware Cloud on AWS. You can test the failover process to ensure success when you need it.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, or Disaster recovery failover admin role.

[Learn about user roles and permissions in NetApp Disaster Recovery.](#)

[Learn about NetApp Console access roles for all services.](#)

About this task

During a failover, Disaster Recovery uses the most recent SnapMirror snapshot copy. Or, you can select a specific snapshot from a point-in-time snapshot (per the retention policy of SnapMirror).

Use the point-in-time option if the most recent replicas are compromised, such as during a ransomware attack. BlueXP disaster recovery shows all available points in time.

This process differs depending on whether the production site is healthy and you are performing a failover to the disaster recovery site for reasons other than a critical infrastructure failure:

- Critical production site failure where the source vCenter or ONTAP cluster is not accessible: NetApp Disaster Recovery lets you select any available snapshot from which to restore.
- Production environment is healthy: You can either "Take a snapshot now" or select a previously created snapshot.

This procedure breaks the replication relationship, places the vCenter source VMs offline, registers the volumes as datastores in the disaster recovery vCenter, restarts the protected VMs using the failover rules in the plan, and enables read/write on the target site.

Test the failover process

Before you start the failover, you can test the process. The test does not place the virtual machines offline.

During a failover test, BlueXP disaster recovery temporarily creates virtual machines. BlueXP disaster recovery maps a temporary datastore backing the FlexClone volume to the ESXi hosts.

This process doesn't consume additional physical capacity on on-premises ONTAP storage or FSx for NetApp ONTAP storage in AWS. The original source volume is not modified and replica jobs can continue even during disaster recovery.

When you finish the test, you should reset the virtual machines with the **Clean up test** option. While this is recommended, it is not required.


A test failover operation does *not* impact production workloads, the SnapMirror relationship used on the test site, and protected workloads that must continue to operate normally.

For a test failover, Disaster Recovery performs the following operations:

- Perform prechecks on the destination cluster and the SnapMirror relationship.

- Create a new FlexClone volume from the selected snapshot for each protected ONTAP volume on the target site ONTAP cluster.
- If any datastores are VMFS, create and map an iGroup to each LUN.
- Register the target virtual machines within vCenter as new datastores.
- Power on the target virtual machines based on the boot order captured in the Resource groups page.
- Unquiesce any supported database applications in VMs indicated as "application consistent."
- If the source vCenter and ONTAP clusters are still active, create a reverse direction SnapMirror relationship to replicate any changes while in failover state back to the original source site.


Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.
3. From the NetApp Disaster Recovery menu, select **Replication plans**.
4. Select the replication plan.
5. On the right, select the **Actions** option  and select **Test failover**.
6. In the Test failover page, enter "Test failover" and select **Test fail over**.
7. After the test is complete, clean up the test environment.

Clean up the test environment after a failover test

After the failover test finishes, you should clean up the test environment. This process removes the temporary VMs from the test location, the FlexClones, and the temporary datastores.

Steps

1. From the NetApp Disaster Recovery menu, select **Replication plans**.
2. Select the replication plan.
3. On the right, select the **Actions** option  and select **Clean up failover test**.
4. In the Test failover page, enter "Clean up failover" and select **Clean up failover test**.

Fail over the source site to a disaster recovery site

In case of a disaster, fail over your primary on-premises VMware site on demand to another on-premises VMware site or VMware Cloud on AWS with FSx for NetApp ONTAP.

The failover process involves in the following operations:


- Disaster Recovery performs prechecks on the destination cluster and SnapMirror relationship.
- If you selected the latest snapshot, the SnapMirror update is performed to replicate the latest changes.
- The source virtual machines are powered down.
- The SnapMirror relationship is broken and the target volume is made read/write.
- Based on the selection of the snapshot, the active file system is restored to the specified snapshot (latest or selected).
- Datastores are created and mounted to the VMware or VMC cluster or host based on the information captured in the replication plan. If any datastores are VMFS, create and map an iGroup to each LUN.

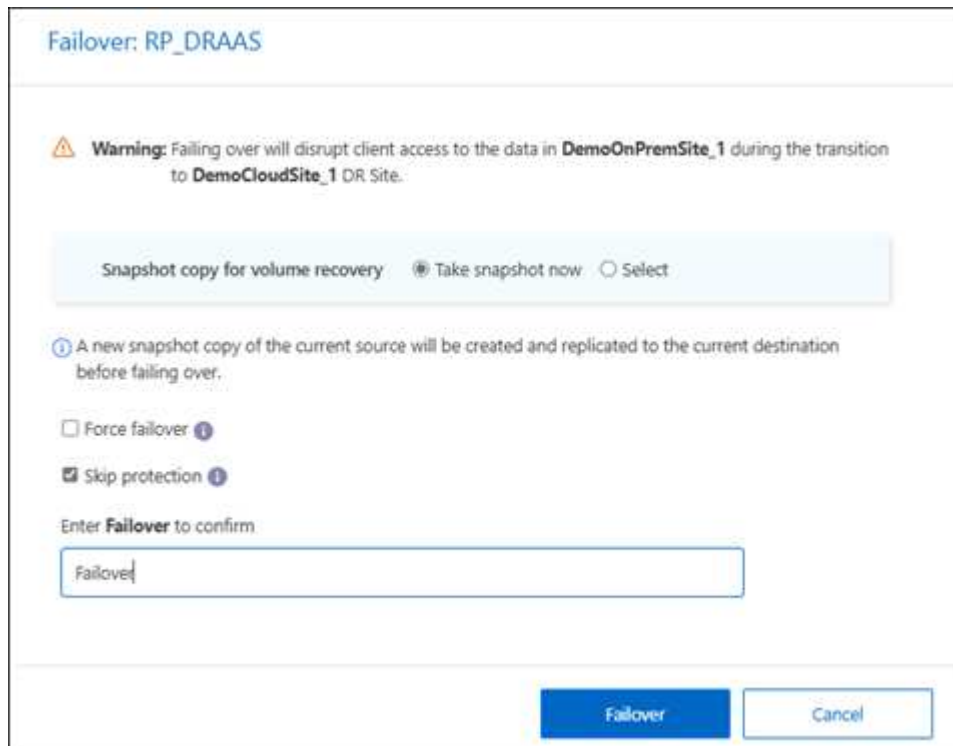
- The target virtual machines are registered within vCenter as new datastores.
- The target virtual machines are powered on based on the boot order captured in the Resource groups page.
- If the source vCenter is still active, power off all source side VMs that are being failed over.
- Unquiesce any supported database applications in VMs indicated as "application consistent."
- If source vCenter and ONTAP clusters are still active, create a reverse direction SnapMirror relationship to replicate any changes while in failover state back to the original source site. The SnapMirror relationship is reversed from target to source virtual machine.



After the failover starts, you can see the recovered VMs in the vCenter of the disaster recovery site (virtual machines, networks, and datastores). By default, the virtual machines are recovered to the Workload folder.

Steps

1. From the NetApp Disaster Recovery menu, select **Replication plans**.
2. Select the replication plan.
3. On the right, select the **Actions** option  and select **Fail over**.



Failover: RP_DRAAS

Warning: Failing over will disrupt client access to the data in **DemoOnPremSite_1** during the transition to **DemoCloudSite_1** DR Site.

Snapshot copy for volume recovery ☒ Take snapshot now ☐ Select

i A new snapshot copy of the current source will be created and replicated to the current destination before failing over.

☐ Force failover **i**

☒ Skip protection **i**

Enter **Failover** to confirm

Failover

Failover Cancel

4. In the Fail over page, either initiate a snapshot now or choose the snapshot for the datastore from which to recover. The default is the latest.

A snapshot of the current source will be taken and replicated to the current destination before the fail over occurs.

5. Optionally, select **Force failover** if you want the failover to occur even if an error is detected that would normally prevent the failover from occurring.
6. Optionally, select **Skip protection** if you want the service to not automatically create a reverse SnapMirror protection relationship after a replication plan failover. This is useful if you want to perform additional

operations on the restored site before you bring it back online within NetApp Disaster Recovery.



You can establish reverse protection by selecting **Protect resources** from the Replication plan Actions menu. This attempts to create a reverse replication relationship for each volume in the plan. You can run this job repeatedly until protection is restored. When protection is restored, you can initiate a failback in the usual way.

7. Type "failover" in the box.
8. Select **Fail over**.
9. To check the progress, in the menu, select **Job monitoring**.

Fail back applications to the original source with NetApp Disaster Recovery

After a disaster has been resolved, fail back from the disaster recovery site to the source site to return to normal operations. You can select the snapshot from which to recover.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, or Disaster recovery failover admin role.

[Learn about user roles and permissions in NetApp Disaster Recovery.](#)


[Learn about NetApp Console access roles for all services.](#)

About this task

In this workflow, NetApp Disaster Recovery 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:

- Perform a compliance check on the recovered site.
- Refresh the vCenter information for each vCenter cluster identified as located in the recovered site.
- On the target site, power off and unregister the virtual machines, and unmount volumes.
- Break the SnapMirror relationship on the original source to make it read/write.
- Resynchronize the SnapMirror relationship to reverse the replication.
- Power on and register the source virtual machines, and mount the volumes on the source.

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.
3. From the NetApp Disaster Recovery menu, select **Replication plans**.
4. Select the replication plan.
5. On the right, select the **Actions** option  and select **Fail back**.
6. Enter the replication plan name to confirm and start the failback.
7. Choose the snapshot for the datastore from which to recover. The default is the latest.
8. To check the progress, in the menu, select **Job monitoring**.

Manage sites, resource groups, replication plans, datastores and virtual machines information with NetApp Disaster Recovery

You can get a quick glance of all your NetApp Disaster Recovery resources or look at each in detail:

- Sites
- Resource groups
- Replication plans
- Datastores
- Virtual machines

Tasks require different NetApp Console roles. For details, see the **Required NetApp Console role** section in each task.

[Learn about user roles and permissions in NetApp Disaster Recovery.](#)

[Learn about NetApp Console access roles for all services.](#)


Manage vCenter sites

You can edit the vCenter site name and the site type (on-premises or AWS).

Required NetApp Console role

Organization admin, Folder or project admin, or Disaster recovery admin role.

Steps

1. From the menu, select **Sites**.
2. Select the **Actions** option  on the right of the vCenter name and select **Edit**.
3. Edit the vCenter site name and location.

Manage resource groups

While you can add a resource group as part of creating a replication plan, you might find it more convenient to add the groups separately and later use those groups in the plan. You create resource groups by VMs or by datastores.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, or Disaster recovery application admin role.

You can create a resource group by datastores in the following ways:


- When you're adding a resource group using datastores, you can see a list of datastores. You can select one or more datastores to create a resource group.
- When you're creating a replication plan and creating a resource group within the plan, you can see the VMs in the datastores.

You can do the following tasks with resource groups:

- Change the resource group name.
- Add VMs to the resource group.
- Remove VMs from the resource group.
- Delete resource groups.

For details about creating a resource group, refer to [Create a resource group to organize VMs together](#).

Steps

1. From the menu, select **Resource groups**.
2. To add a resource group, select **Add group**.
3. To perform actions with the resource group, select the **Actions** option  at the right and select any of the options, such as **Edit resource group** or **Delete resource group**.

Manage replication plans

You can disable, enable and delete replication plans. You can change schedules.

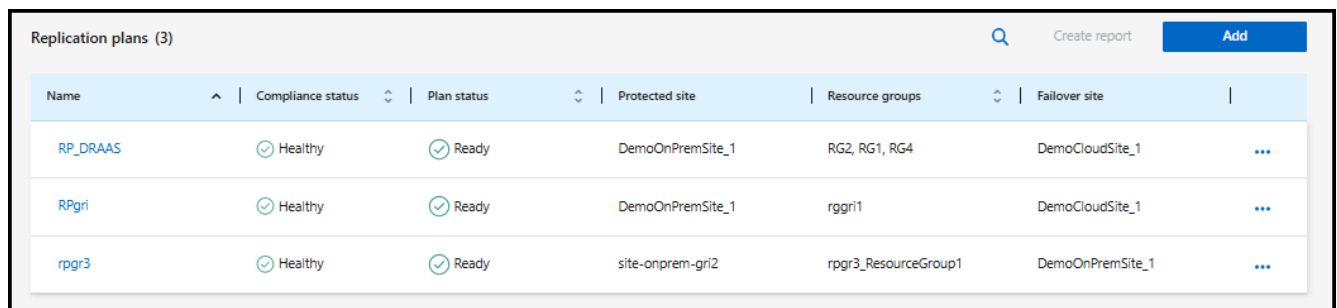
Required NetApp Console role

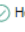
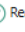

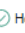


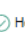
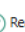

Organization admin, Folder or project admin, Disaster recovery admin, Disaster recovery failover admin, or Disaster recovery application admin role.



- If you want to pause a replication plan temporarily, you can disable it and later enable it.
- If you no longer need the plan, you can delete it.

Steps


1. From the menu, select **Replication plans**.



Name	Compliance status	Plan status	Protected site	Resource groups	Failover site	
RP_DRAAS	 Healthy	 Ready	DemoOnPremSite_1	RG2, RG1, RG4	DemoCloudSite_1	
RPgr1	 Healthy	 Ready	DemoOnPremSite_1	rggri1	DemoCloudSite_1	
rpgr3	 Healthy	 Ready	site-onprem-gr12	rpgr3_ResourceGroup1	DemoOnPremSite_1	

2. To view the plan details, select the **Actions** option  and select **View plan details**.
3. Do any of the following:
 - To edit the plan details (change the recurrence), select the **Plan details** tab and select the **Edit** icon to the right.
 - To edit the resource mappings, select the **Failover mapping** tab and select the **Edit** icon.
 - To add or edit the virtual machines, select the **Virtual machines** tab and select the **Add VMs** option or **Edit** icon.
4. Return to the list of plans by selecting "Replication plans" in the breadcrumbs at the left.
5. To perform actions with the plan, from the list of replication plans, select the **Actions** option  to the right of the plan and select any of the options, such as **Edit schedules**, **Test failover**, **Fail over**, **Fail back**,

Migrate, Take snapshot now, Clean up old snapshots, Disable, Enable, or Delete.

6. To set or change a test failover schedule or set the compliance frequency check, select the **Actions** option  to the right of the plan and select **Edit schedules**.
 - a. In the Edit schedules page, enter how often in minutes you want the failover compliance check to occur.
 - b. Check **Run test failovers on a schedule**.
 - c. In the Repeat option, select the daily, weekly, or monthly schedule.
 - d. Select **Save**.

Reconcile snapshots on demand

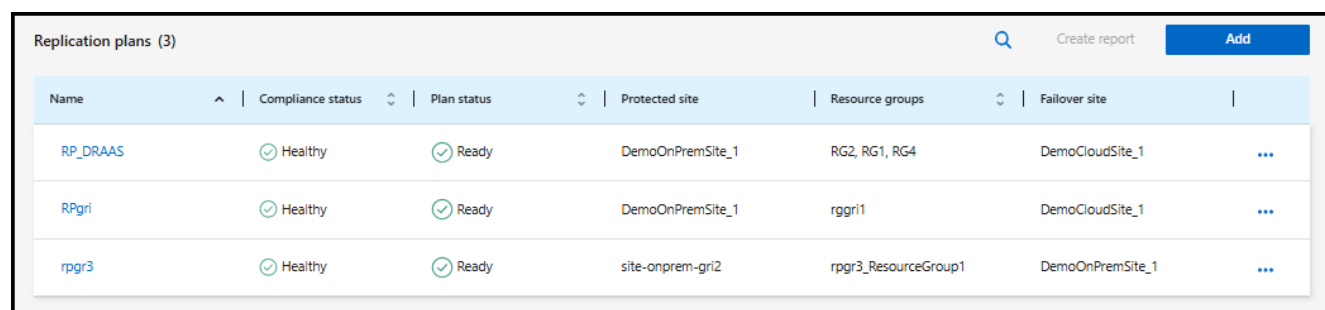
You can reconcile snapshots that are out of sync between the source and target. This might occur if snapshots are deleted on a target outside of NetApp Disaster Recovery. The service automatically deletes the snapshot on the source automatically every 24 hours. However, you can perform this on demand. This feature enables you to ensure that the snapshots are consistent across all sites.

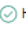



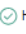

Required NetApp Console role


Organization admin, Folder or project admin, Disaster recovery admin, Disaster recovery failover admin, or Disaster recovery application admin role.

Steps

1. From the menu, select **Replication plans**.



Name	Compliance status	Plan status	Protected site	Resource groups	Failover site	
RP_DRAAS	 Healthy	 Ready	DemoOnPremSite_1	RG2, RG1, RG4	DemoCloudSite_1	...
RPgr1	 Healthy	 Ready	DemoOnPremSite_1	rggr1	DemoCloudSite_1	...
rpgr3	 Healthy	 Ready	site-onprem-gr12	rpgr3_ResourceGroup1	DemoOnPremSite_1	...

2. From the list of replication plans, select the **Actions** option  to the right of the plan and select **Reconcile snapshots**.
3. Review the reconciliation information.
4. Select **Reconcile**.


Delete a replication plan

You can delete a replication plan if you no longer need it. If you delete a replication plan, you can also delete the primary and secondary snapshots created by the plan.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, Disaster recovery failover admin, or Disaster recovery application admin role.

Steps

1. From the menu, select **Replication plans**.
2. Select the **Actions** option  to the right of the plan and select **Delete**.

3. Select whether you want to delete the primary snapshots, secondary snapshots, or just the metadata created by the plan.
4. Type "delete" to confirm the deletion.
5. Select **Delete**.

Change retention count for failover schedules

You can change how many datastores are retained.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, Disaster recovery failover admin, or Disaster recovery application admin role.

Steps

1. From the menu, select **Replication plans**.
2. Select the replication plan, select the **Failover mapping** tab, and select the **Edit** pencil icon.
3. Select the **Datastores** arrow to expand it.

Disaster recovery | Dashboard | Sites | **Replication plans** | Resource groups | Job monitoring | View payment methods

Replication plan > Plan details > Edit failover mappings

☒ Use same mappings for failover and test mappings

Failover mappings | Test mappings

Compute resources ✓ Mapped

Virtual networks ✓ Mapped

Datastores ⌵

The selected virtual machines are from different volumes. Once the plan is created, disaster recovery will create a consistency group snapshot of the source that spans multiple volumes.

RPO for all datastores in minutes 132 | Retention count for all datastores 30

Source datastore	Target datastore
BizAppDatastore (Temp_3510_N1:DR_Prod_Source)	BizAppDatastore_dest (test:DR_Prod_dest)
DS_SFO (Temp_3510_N1:DR_SFO)	DS_SFO (test:DR_SFO_dest)
DS_Testing_Staging	DS_Testing_Staging_dest (test:DR_Vol_Staging_dest)

Transfer schedule(RPO) : hourly

Cancel Save

4. Change the value of the retention count in the replication plan.
5. With the replication plan selected, select the Actions menu, then select **Clean up old snapshots** to remove old snapshots on the target to match the new retention count.

View datastores information

You can view information about how many datastores exist on the source and on the target.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, Disaster recovery failover admin, Disaster recovery application admin, or Disaster recovery viewer role.

Steps

1. From the menu, select **Dashboard**.
2. Select the vCenter in the site row.
3. Select **Datastores**.
4. View the datastores information.

View virtual machines information

You can view information about how many virtual machines exist on the source and on the target along with CPU, memory, and available capacity.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, Disaster recovery failover admin, Disaster recovery application admin, or Disaster recovery viewer role.

Steps

1. From the menu, select **Dashboard**.
2. Select the vCenter in the site row.
3. Select **Virtual machines**.
4. View the virtual machines information.

Monitor NetApp Disaster Recovery jobs

You can monitor all NetApp Disaster Recovery jobs and determine their progress.

View jobs

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, Disaster recovery application admin, or Disaster recovery viewer role.

[Learn about user roles and permissions in NetApp Disaster Recovery.](#)

[Learn about NetApp Console access roles for all services.](#)

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.
3. From the menu, select **Job monitoring**.
4. Explore all jobs related to operations and review their timestamps and status.
5. To view details of a particular job, select that row.
6. To refresh information, select **Refresh**.

Cancel a job

If a job is in progress or in a queued state and you don't want it to continue, you can cancel it. You might want to cancel a job if it is stuck in the same state and you want to free up the next operation in the queue. You might want cancel a job before it times out.

Required NetApp Console role

Organization admin, Folder or project admin, Disaster recovery admin, Disaster recovery failover admin, or Disaster recovery application admin role.

[Learn about user roles and permissions in NetApp Disaster Recovery.](#)

[Learn about NetApp Console access roles for all services.](#)

Steps

1. From the NetApp Console left navigation bar, select **Protection > Disaster recovery**.
2. From the menu, select **Job monitoring**.
3. In the Job monitor page, note the ID of the job you want to cancel.

The job must be in an "In progress" or "Queued" state.

4. In the Actions column, select **Cancel job**.

Create NetApp Disaster Recovery reports

Reviewing NetApp Disaster Recovery reports can help you analyze your disaster recovery preparedness. Predesigned reports include a summary of test failovers, replication plan details, and job details on all sites within an account for the past seven days.

You can download reports in PDF, HTML, or JSON format.

The Download link is valid for six hours.

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console left navigation, select **Protection > Disaster recovery**.
3. From the NetApp Console left navigation bar, select **Replication plans**.
4. Select **Create report**.
5. Select the type of file format and the time period within the last 7 days.
6. Select **Create**.



The report might take a few minutes to display.

7. To download a report, select **Download report** and select it in the administrator's Download folder.

Reference

vCenter privileges needed for NetApp Disaster Recovery

The vCenter account must have a minimum set of vCenter privileges to allow NetApp Disaster Recovery to perform its services, such as registering and deregistering datastores, starting and stopping VMs, and reconfiguring virtual machines (VMs). The following table lists all privileges required for NetApp Disaster Recovery to interface with a vCenter cluster.

Type	Privilege name	Description
Datastore	Datastore.Configure datastore	Use to configure a datastore.
	Datastore.Remove datastore	Use to remove a datastore.
Virtual Machine	Virtual machine.Configuration.Change Settings	Use to change general VM settings.
	Virtual machine.Configuration.Modify device settings	Use to change the properties of an existing device.
	Virtual machine.Configuration.Reload from path	Use to change a VM configuration patch while preserving the identity of the VM. Solutions such as VMware vCenter Site Recovery Manager use this operation to maintain VM identity during failover and failback.
	Virtual machine.Configuration.Rename	Use to rename a VM or modify the associated nodes of a VM.
	Virtual machine.Configuration.Reset guest information	Use to edit the guest operating system information for a VM.
	Virtual machine.Configuration.Change Memory	Use to change the amount of memory allocated to the VM.
	Virtual machine.Configuration.Change CPU count	Use to change the number of virtual CPUs.
Virtual Machine Guest	Virtual machine.Guest Operations.Guest Operation Modifications	Enables VM guest operations that involve changes to a guest operating system in a VM, such as transferring a file to the VM.

Type	Privilege name	Description
Virtual Machine Interaction	Virtual machine.Interaction.Power Off	Use to power off a powered-on VM. This operation powers down the guest operating system.
	Virtual machine.Interaction.Power on	Use to power on a powered-off VM and resume a suspended VM.
	Virtual machine.Interaction.VMware Tools install	Use to mount and unmount the VMware Tools CD installer as a CD-ROM for the guest operating system.
Virtual Machine Inventory	Virtual machine.Inventory.Create new	Use to create a VM and allocate resources for its execution.
	Virtual machine.Inventory.Register	Use to add an existing VM to a vCenter Server or host inventory.
	Virtual machine.Inventory.Unregister	Use to unregister a VM from a vCenter Server or host inventory.
Virtual Machine State	Virtual machine.Snapshot management.Create snapshot	Use to create a snapshot from the VM's current state.
	Virtual machine.Snapshot management.Remove Snapshot	Use to remove a snapshot from the snapshot history.
	Virtual machine.Snapshot management.Revert to snapshot	Use to set the VM to the state it was in at a given snapshot.

NetApp Disaster Recovery role-based access to features

NetApp Disaster Recovery employs roles to govern the access that each user has to specific features and actions.

The service uses the following roles that are specific to NetApp Disaster Recovery.

- **Disaster recovery admin:** Perform any actions in NetApp Disaster Recovery.
- **Disaster recovery failover admin:** Perform failover and migrate actions in NetApp Disaster Recovery.
- **Disaster recovery application admin:** Create and modify replication plans and start test failovers.
- **Disaster recovery viewer:** View information in NetApp Disaster Recovery, but cannot perform any actions.

These roles are specific to NetApp Disaster Recovery and are not the same as the platform roles that are used in the NetApp Console. For details about all NetApp Console platform roles, see [the NetApp Console setup and administration documentation](#).

The following table indicates the actions that each NetApp Disaster Recovery role can perform.

Feature and action	Disaster recovery admin	Disaster recovery failover admin	Disaster recovery application admin	Disaster recovery viewer
View dashboard and all tabs	Yes	Yes	Yes	Yes
Start free trial	Yes	No	No	No
Initiate discovery of workloads	Yes	No	No	No
View license information	Yes	Yes	Yes	Yes
Activate license	Yes	No	Yes	No
On the Sites option:				
View sites	Yes	Yes	Yes	Yes
Add, modify, or delete sites	Yes	No	No	No
On the Replication plans option:				
View replication plans	Yes	Yes	Yes	Yes
View replication plan details	Yes	Yes	Yes	Yes
Create or modify replication plans	Yes	Yes	Yes	No
Create reports	Yes	No	No	No
View snapshots	Yes	Yes	Yes	Yes
Perform failover tests	Yes	Yes	Yes	No
Perform failovers	Yes	Yes	No	No
Perform failbacks	Yes	Yes	No	No
Perform migrations	Yes	Yes	No	No
On the Resource groups option:				
View resource groups	Yes	Yes	Yes	Yes
Create, modify, or delete resource groups	Yes	No	Yes	No
On the Job Monitoring option:				

Feature and action	Disaster recovery admin	Disaster recovery failover admin	Disaster recovery application admin	Disaster recovery viewer
View jobs	Yes	No	Yes	Yes
Cancel jobs	Yes	Yes	Yes	No

Use NetApp Disaster Recovery with Amazon EVS

Introduction of NetApp Disaster Recovery using Amazon Elastic VMware Service and Amazon FSx for NetApp ONTAP

Increasingly, customers have become more dependent on virtualized infrastructures for production compute workloads such as those based on VMware vSphere. As these virtual machines (VMs) have become more critical to their businesses, customers need to protect these VMs from the same types of disasters as their physical compute resources. Disaster recovery (DR) solutions currently offered are complex, expensive, and resource intensive. NetApp, the largest storage provider used for virtualized infrastructures, has a vested interest in ensuring its customers' VMs are protected in the same way that we protect ONTAP storage-hosted data of any type. To meet this goal, NetApp created the NetApp Disaster Recovery service.

One of the primary challenges with any DR solution is managing the incremental cost of purchasing, configuring, and maintaining additional compute, network, and storage resources just to provide a DR replication and recovery infrastructure. One popular option for protecting critical on-premises virtual resources is to use cloud-hosted virtual resources as the DR replication and recovery infrastructure. Amazon is one example of such a solution that can provide cost-effective resources that are compatible with NetApp ONTAP hosted VM infrastructures.

Amazon introduced its Amazon Elastic VMware Service (Amazon EVS) that enables VMware Cloud Foundation within your virtual private cloud (VPC). Amazon EVS provides the resilience and performance of AWS along with the familiar VMware software and tools enabling Amazon EVS vCenters to be integrated as an extension of your on-premises virtualized infrastructure.

While Amazon EVS comes with included storage resources, using native storage can reduce its effectiveness for organizations with storage-heavy workloads. In these cases, teaming Amazon EVS with Amazon FSx for NetApp ONTAP storage (Amazon FSxN) can provide a more flexible storage solution. In addition, when you are using NetApp ONTAP storage solutions on-premises to host your VMware infrastructure, using Amazon EVS with FSx for ONTAP means you get best-in-class data interoperability and protection features between your on-premises and cloud-hosted infrastructures.

For information about Amazon FSx for NetApp ONTAP, see [Getting started with Amazon FSx for NetApp ONTAP](#).

Solution overview of NetApp Disaster Recovery using Amazon EVS and Amazon FSs for NetApp ONTAP

NetApp Disaster Recovery is a value-added service hosted within the NetApp Console

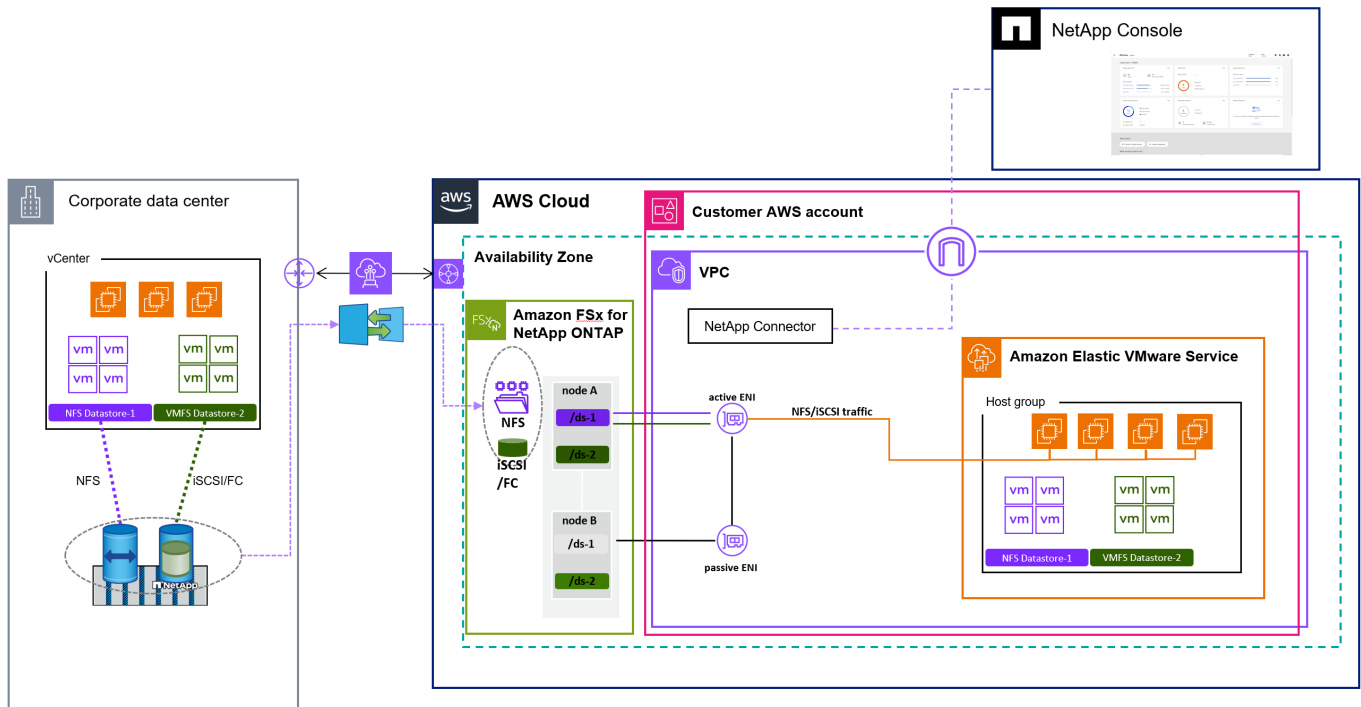
software-as-a-service environment, which depends on the core NetApp Console architecture. Several main components comprise the DR service for VMware protection within the Console.

For a complete overview of the NetApp Disaster Recovery solution, see [Learn about NetApp Disaster Recovery for VMware](#).

If you want to protect your on-premises VMware hosted virtual machines to Amazon AWS, use the service to back up to Amazon EVS with Amazon FSx for NetApp ONTAP storage hosted datastores.

The following figure shows how the service works to protect your VMs with Amazon EVS.

Overview of NetApp Disaster Recovery using Amazon EVS and FSx for ONTAP



1. Amazon EVS is deployed in your account in a single Availability Zone (AZ) configuration and within the your Virtual Private Cloud (VPC).
2. An FSx for ONTAP file system is deployed in the same AZ as the Amazon EVS deployment. The file system connects to Amazon EVS either directly through an Elastic Network Interface (ENI), a VPC peer connection, or an AmazonTransit Gateway.
3. The NetApp Console agent is installed in your VPC. The NetApp Console agent hosts multiple data management services (called agents), including the NetApp Disaster Recovery agent that manages DR of the VMware infrastructure on both your local physical datacenters and your Amazon AWS hosted resources.
4. The NetApp Disaster Recovery agent securely communicates with the NetApp Console cloud-hosted service to receive tasks and distributes those tasks to the appropriate on-premises and AWS hosted vCenter and ONTAP storage instances.
5. You create a replication plan by using the NetApp Console cloud-hosted UI console indicating the VMs that should be protected, the frequency those VMs should be protected, and the procedures that need to be performed to restart those VMs in the event of a failover from the on-premises site.
6. The replication plan determines which vCenter datastores are hosting the protected VMs and the ONTAP volumes that are hosting those datastores. If volumes do not yet exist on the FSx for ONTAP cluster,

NetApp Disaster Recovery automatically creates them.

7. A SnapMirror relationship is created for each identified source ONTAP volume to each destination FSx for ONTAP hosted ONTAP volume and a replication schedule is created based on the user-provided RPO in the replication plan.
8. In the event of the primary site failure, an administrator initiates a manual failover process within the NetApp Console and selects a backup to use as the restore point.
9. The NetApp Disaster Recovery agent activates the FSx for ONTAP hosted data protection volumes.
10. The agent registers each activated FSx for ONTAP volume with the Amazon EVS vCenter, registers each protected VM with the Amazon EVS vCenter, and starts each according to the predefined rules contained in the replication plan.

Install the NetApp Console agent for NetApp Disaster Recovery

A NetApp Console agent is NetApp software running in your cloud or on-premises network. It executes the actions that the NetApp Console needs to perform to manage your data infrastructure. The Console agent constantly polls the NetApp Disaster Recovery software as a service layer for any actions that it needs to take.

For NetApp Disaster Recovery, the actions that are performed orchestrate VMware vCenter clusters and ONTAP storage instances using native APIs for each respective service to provide protection for production VMs running in an on-premises location. While the Console agent can be installed in any of your network locations, for NetApp Disaster Recovery we recommend that you install the Console agent in the DR site. This ensures that in the event of a failure of the primary site, the NetApp cloud-based console UI continues to have contact with the Console agent and can orchestrate the recovery process within that DR site.

To use the service, install the Console agent in standard mode. To learn more about the types of Console agent installations, visit [Learn about NetApp Console deployment modes | NetApp Documentation](#).

While the Console agent is critical to using the service, the installation steps to install the Console agent depend on your needs and network configuration. It is beyond the scope of this information to provide specific instructions for installation.

The simplest method for installing the Console agent with Amazon AWS is to use the AWS Marketplace. For details about Console agent installation using the AWS Marketplace, see [Create a Console agent from the AWS Marketplace | NetApp Documentation](#).

Configure NetApp Disaster Recovery for Amazon EVS

Configure NetApp Disaster Recovery for Amazon EVS overview

After you install the NetApp Console agent, you need to integrate all the ONTAP storage and VMware vCenter resources that will participate in the disaster recovery process with NetApp Disaster Recovery.

- [Prerequisites for Amazon EVS with NetApp Disaster Recovery](#)
- [Add ONTAP storage arrays to NetApp Disaster Recovery](#)
- [Enable NetApp Disaster Recovery for Amazon EVS](#)
- [Add vCenter sites to NetApp Disaster Recovery](#)
- [Add vCenter clusters to NetApp Disaster Recovery](#)

Prerequisites for Amazon EVS with NetApp Disaster Recovery

You should ensure that several prerequisites are met before you continue to configure Amazon EVS with NetApp Disaster Recovery.

Specifically, do the following:

- Create a vCenter user account with the specific VMware privileges required for NetApp Disaster Recovery to perform the necessary operations.



We do not recommend using the default "administrator@vsphere.com" administrator account. Instead, you should create a NetApp Disaster Recovery specific user account on all vCenter clusters that will participate in the DR process. For a list of specific privileges required, see [vCenter privileges needed for NetApp Disaster Recovery](#).

- Ensure that all vCenter datastores that will host VMs protected by NetApp Disaster Recovery are located on NetApp ONTAP storage resources.

The service supports NFS and VMFS on iSCSI (and not FC) when using Amazon FSx on NetApp ONTAP. While the service supports FC, Amazon FSx for NetApp ONTAP does not.

- Ensure that your Amazon EVS vCenter is connected to an Amazon FSx for NetApp ONTAP storage cluster.
- Ensure that VMware tools are installed on all protected VMs.
- Ensure that your on-premises network is connected to your AWS VPC network using an Amazon approved connection method. We recommend that you use AWS Direct Connect, AWS Private Link, or an AWS Site-to-Site VPN.

Add on-premises arrays to the NetApp Console system for Amazon EVS with NetApp Disaster Recovery

Before using NetApp Disaster Recovery, you must add on-premises and cloud-hosted storage instances to the NetApp Console system.


You need to do the following:

- Add on-premises arrays to your NetApp Console system.
- Add Amazon FSx for NetApp ONTAP (FSx for ONTAP) instances to your NetApp Console system.

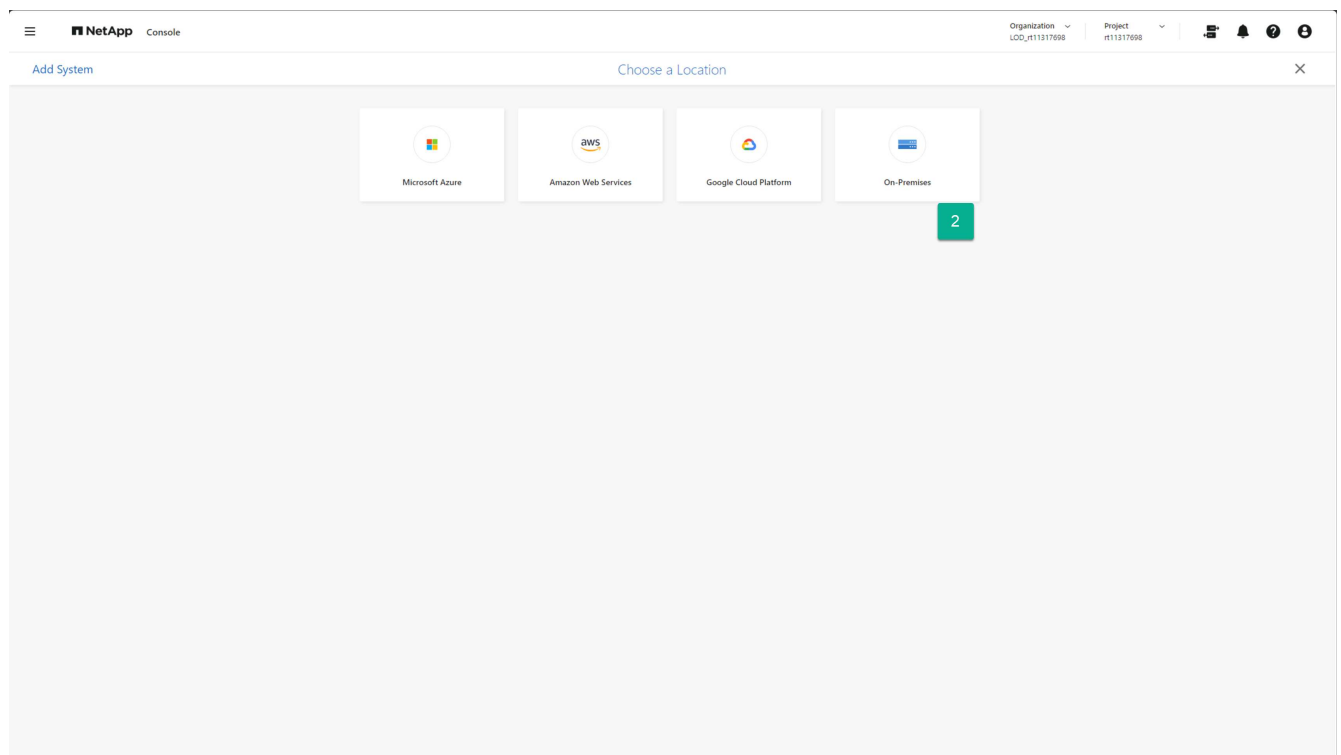
Add on-premises storage arrays to NetApp Console system

Add on-premises ONTAP storage resources to your NetApp Console system.


1. From the NetApp Console Systems page, select **Add System**.



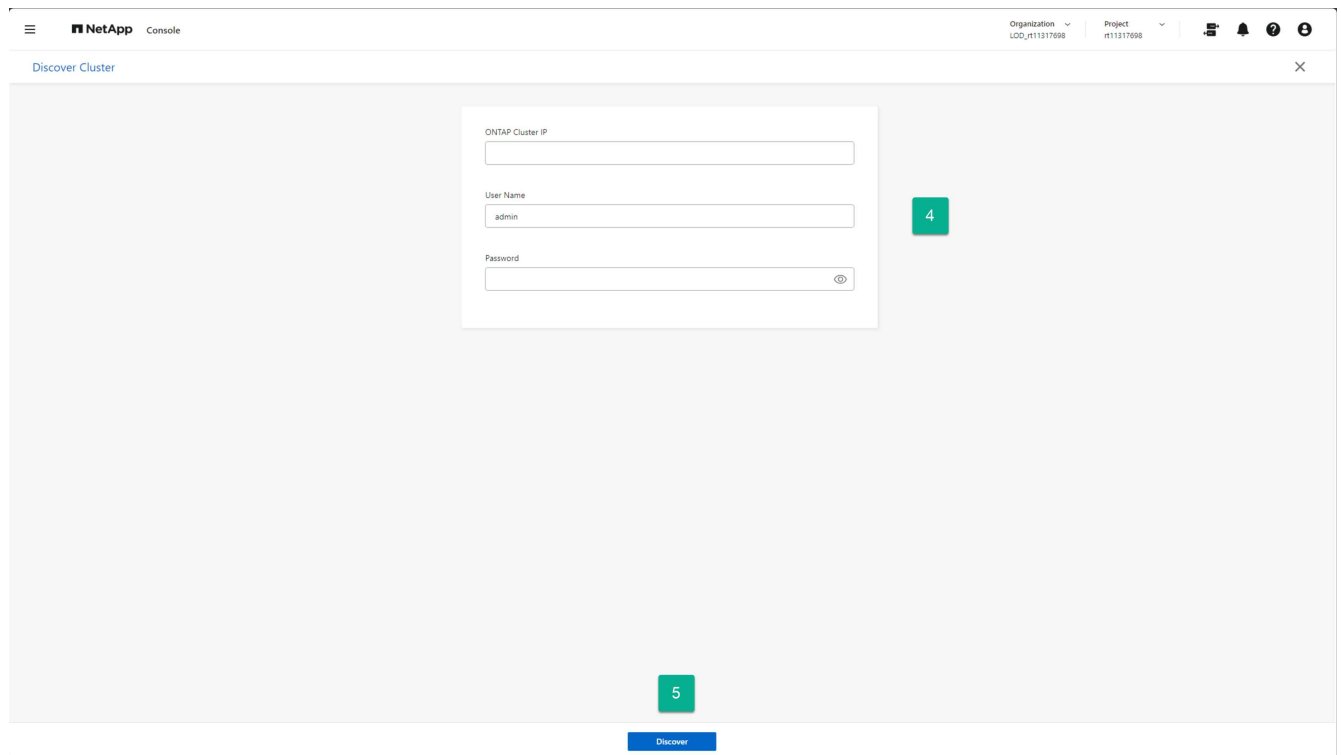
2. From the Add System page, select the **On-Premises** card.



3. Select **Discover** on the On-Premises ONTAP card.



4. On the Discover Cluster page, enter the following information:
 - a. The IP address of the ONTAP array cluster management port
 - b. The administrator username
 - c. The administrator password
5. Select **Discover** at the bottom of the page.

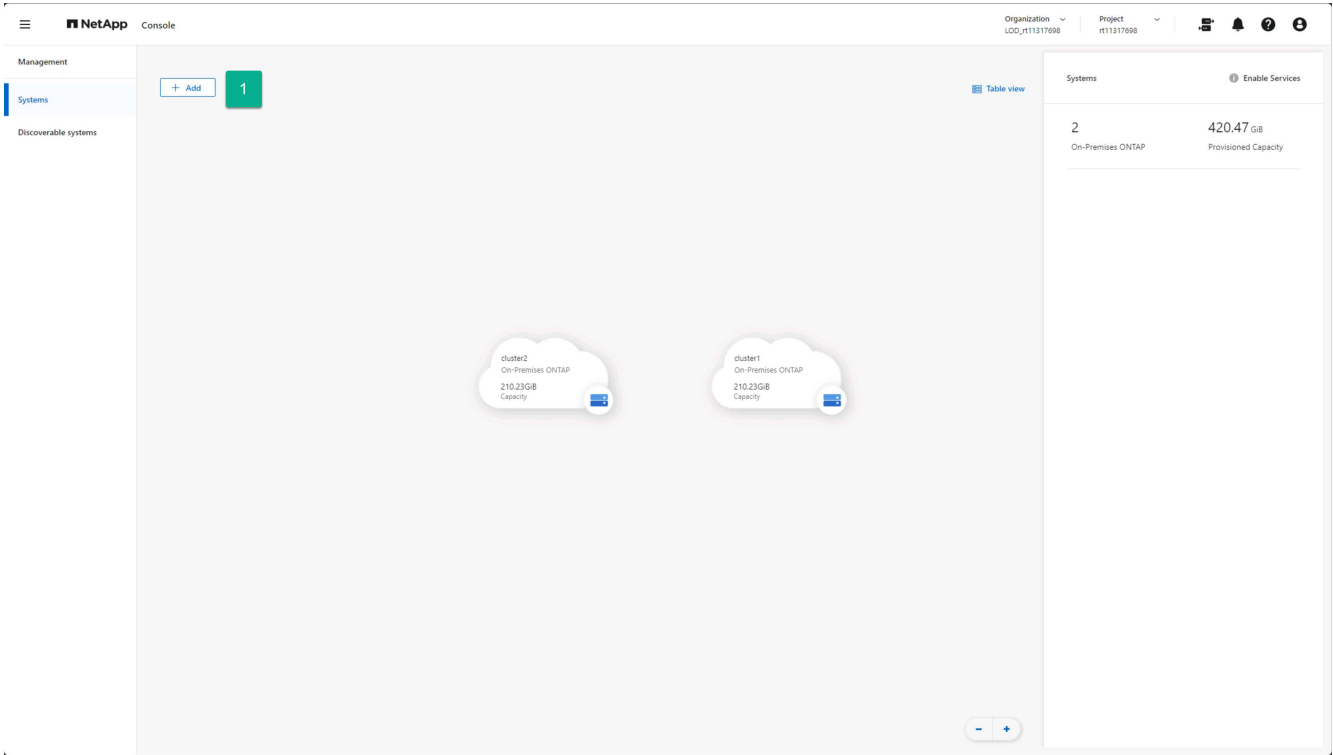


6. Repeat steps 1-5 for each ONTAP array that will host vCenter datastores.

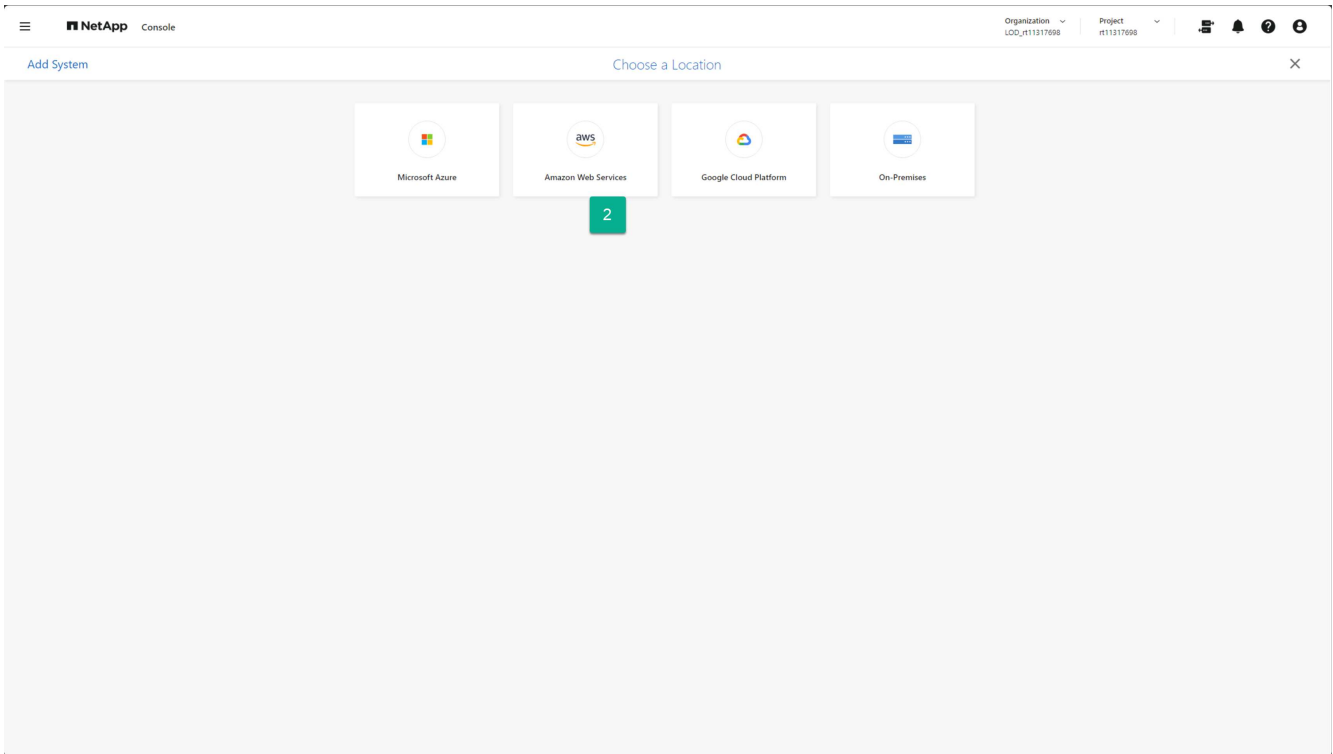
Add Amazon FSx for NetApp ONTAP storage instances to NetApp Console system

Next, add an Amazon FSx for NetApp ONTAP storage resources to your NetApp Console system.

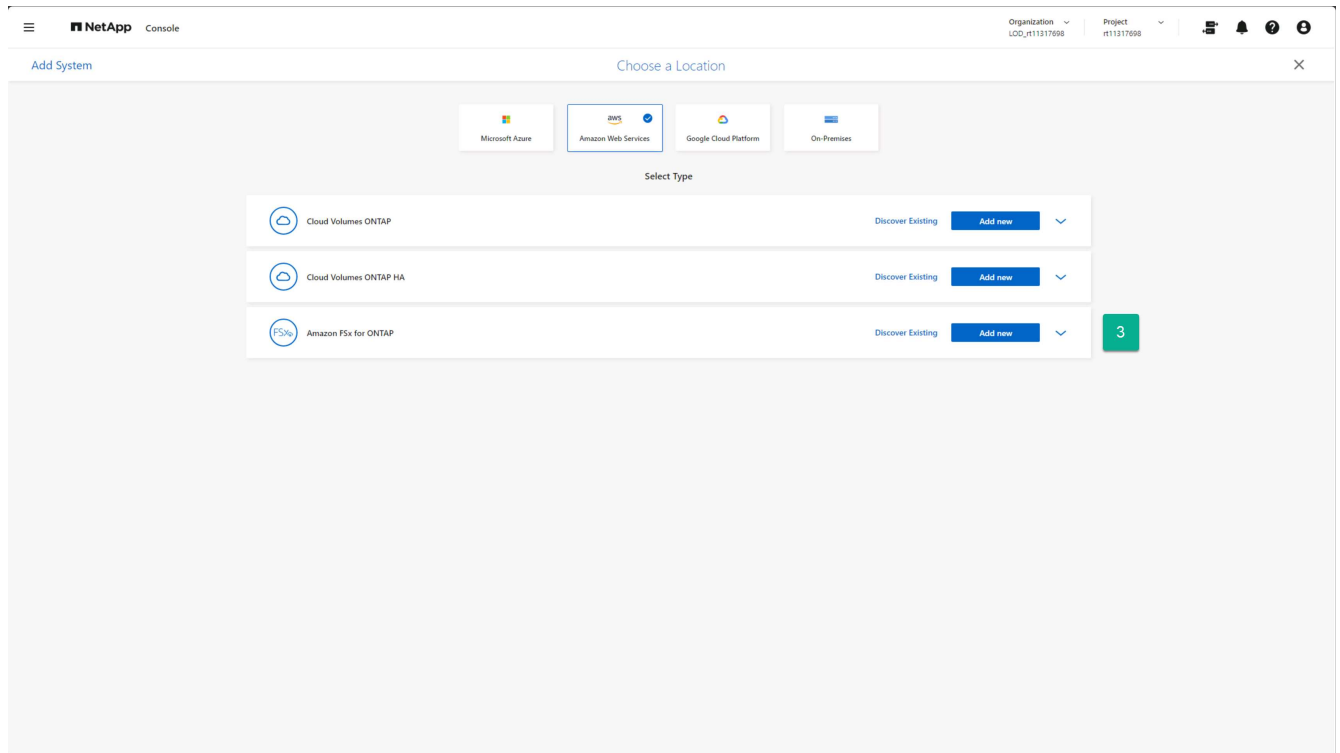
- 1. From the NetApp Console Systems page, select **Add System**.



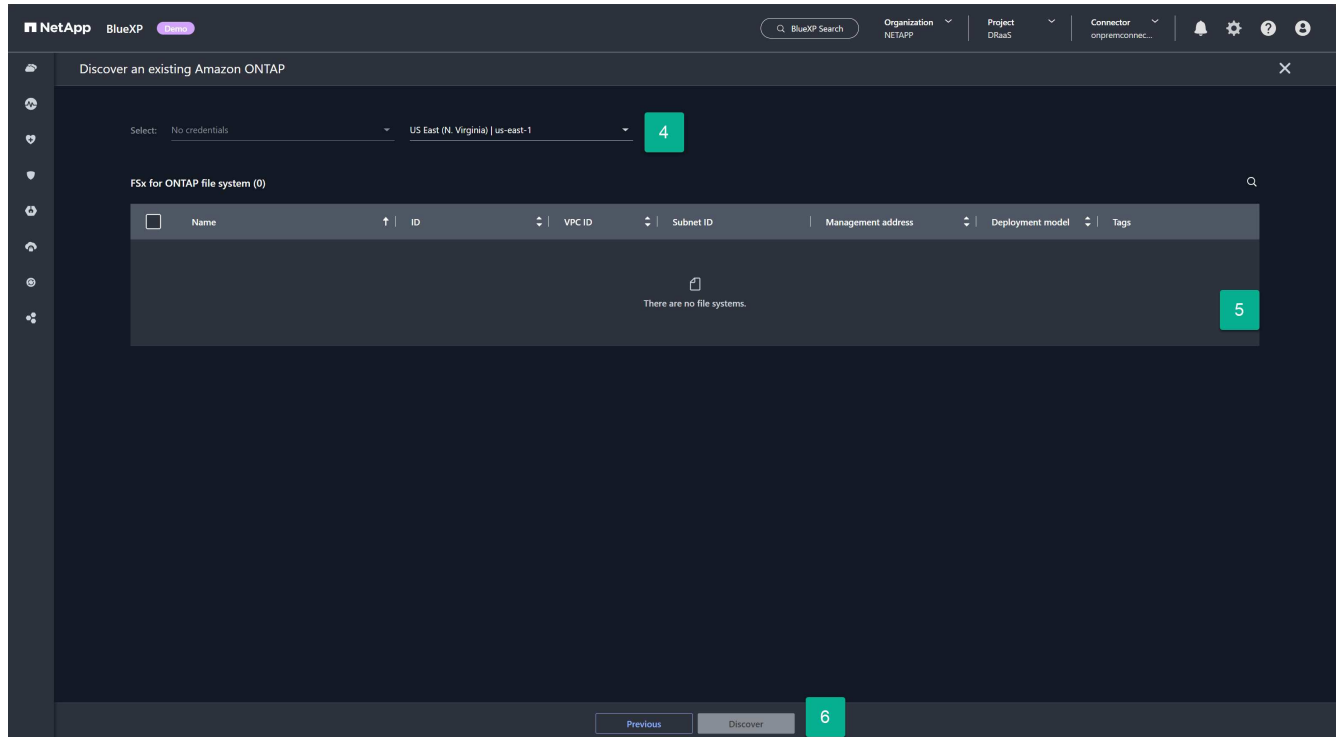
- 2. From the Add System page, select the **Amazon Web Services** card.



- 3. Select the **Discover Existing** link on the Amazon FSx for ONTAP card.



4. Select the credentials and AWS region hosting the FSx for ONTAP instance.
5. Select one or more FSx for ONTAP file systems to be added.
6. Select **Discover** at the bottom of the page.



7. Repeat steps 1-6 for each FSx for ONTAP instance that will host vCenter datastores.

Add NetApp Disaster Recovery service to your NetApp Console account for Amazon EVS

NetApp Disaster Recovery is a licensed product offering that must be purchased before it can be used. There are several types of licenses and several ways that you can purchase licenses. A license entitles you to protect a specific amount data for a specific span of time.

For more information about NetApp Disaster Recovery licenses, see [Set up licensing for NetApp Disaster Recovery](#).

License types

There are two primary license types:

- NetApp offers a [30-day trial license](#) that you can use to evaluate NetApp Disaster Recovery using your ONTAP and VMware resources. This license provides 30 days of use for an unlimited amount of protected capacity.
- Purchase a production license if you want DR protection beyond the 30-day trial period. This license can be purchased through the marketplaces of any of NetApp's cloud partners, but for this guide, we recommend that you purchase your marketplace license for NetApp Disaster Recovery using the Amazon AWS Marketplace. To learn more about purchasing a license through the Amazon Marketplace, see [Subscribe through AWS Marketplace](#).

Size your disaster recovery capacity needs

Before you purchase your license, you should understand how much ONTAP storage capacity you need to protect. One of the advantages of using NetApp ONTAP storage is the high efficiency with which NetApp stores your data. All data stored in an ONTAP volume — such as VMware datastore hosting VMs — is that the data is stored in a highly efficient manner. ONTAP defaults to three types of storage efficiency when writing data to physical storage: compaction, deduplication, and compression. The net result is storage efficiencies of between 1.5:1 and 4:1 depending on the types of data being stored. In fact, NetApp offers a [storage efficiency guarantee](#) for certain workloads.

This can benefit you because NetApp Disaster Recovery computes capacity for the purposes of licensing after all ONTAP storage efficiencies are applied. For example, let's say you have provisioned a 100 terabyte (TiB) NFS datastore within vCenter to host 100 VMs that you want to protect using the service. Additionally, let's assume when the data is written to the ONTAP volume, automatically applied storage efficiency techniques result in those VMs consuming only 33TiB (3:1 storage efficiency). NetApp Disaster Recovery needs to be licensed only for 33TiB, not 100TiB. This can be a very large benefit to the total cost of ownership for your DR solution when compared to other DR solutions.

Steps

1. To determine how much data is being consumed on each volume hosting a VMware datastore to be protected, determine the on-disk capacity consumption by running the ONTAP CLI command for each volume: `volume show-space -volume < volume name > -vserver < SVM name > .`

For example:

```
cluster1::> volume show-space
Vserver : vm-nfs-ds1
Volume  : vol0
Feature                                Used      Used%
-----
User Data                             163.4MB    3%
Filesystem Metadata                     172KB     0%
Inodes                                 2.93MB    0%
Snapshot Reserve                       292.9MB    5%
Total Metadata                          185KB     0%
Total Used                             459.4MB    8%
Total Physical Used                     166.4MB    3%
```

2. Note the **Total Physical Used** value for each volume. This is the amount of data that NetApp Disaster Recovery needs to protect, and it is the value that you will use to determine how much capacity you need to license.

Add sites in NetApp Disaster Recovery for Amazon EVS

Before you can protect your VM infrastructure, identify which VMware vCenter clusters are hosting the VMs to be protected and where those vCenters are located. The first step is to create a site to represent the source and destination datacenters. A site is a failure domain or a recovery domain.

You need to create the following:


- A site to represent each production datacenter where your production vCenter clusters reside
- A site for your Amazon EVS/Amazon FSx for NetApp ONTAP cloud datacenter

Create on-premises sites


Create a production vCenter site.

Steps


1. From the NetApp Console left navigation bar, select **Protection > Disaster Recovery**.
2. From any page in NetApp Disaster Recovery, select the **Sites** option.



3. From the Sites option, select **Add**.



4. In the Add site dialog box, provide a site name.
5. Select “On-prem” as the Location.
6. Select **Add**.




If you have other production vCenter sites, you can add them using the same steps.


Create Amazon cloud sites

Create a DR site for Amazon EVS using Amazon FSx for NetApp ONTAP storage.


1. From any page in NetApp Disaster Recovery, select the **Sites** option.



2. From the Sites option, select **Add**.



3. In the Add site dialog box, provide a site name.
4. Select "AWS-EVS" as the Location.
5. Select **Add**.



Result

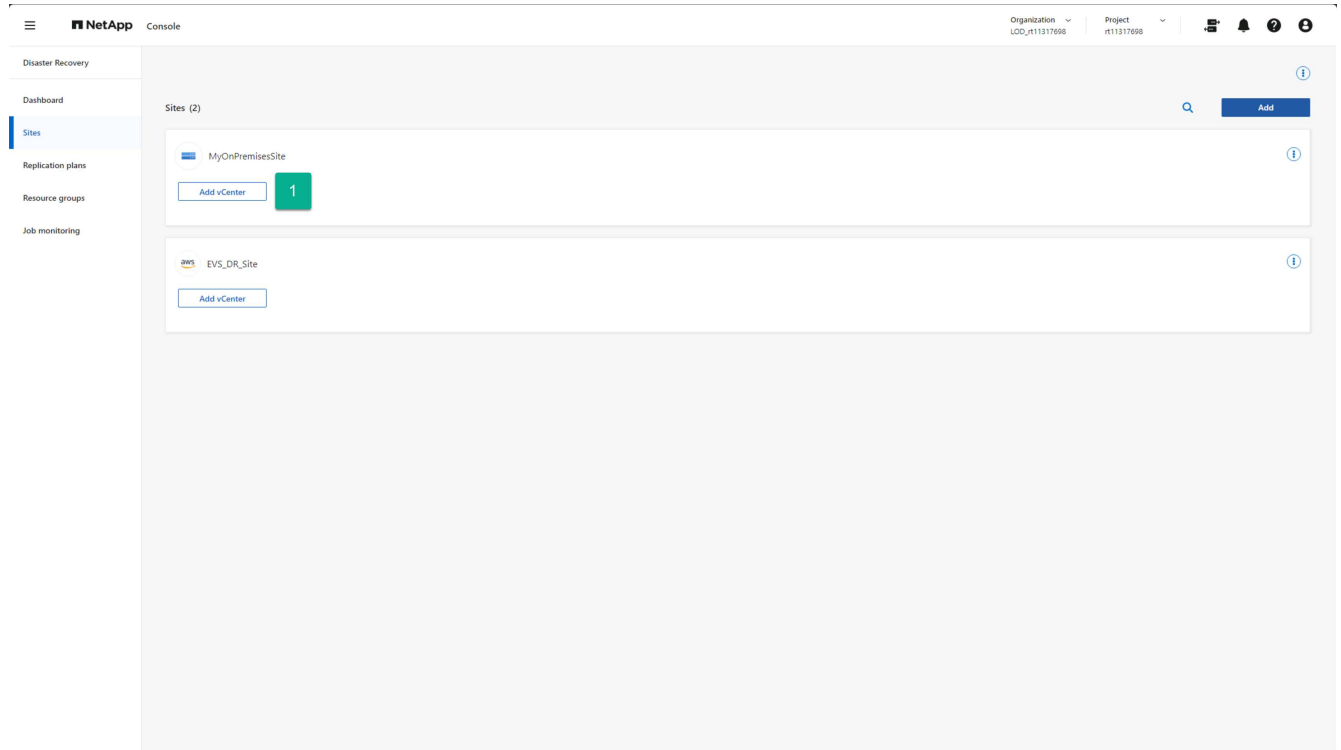
You now have a production (source) site and a DR (destination) site created.

Add on-premises and Amazon EVS vCenter clusters in NetApp Disaster Recovery


With sites created, you now add your vCenter clusters to each site in NetApp Disaster Recovery. When we created each site, we indicated each type of site. This tells NetApp Disaster Recovery what type of access is required for the vCenters hosted in each site type. One of the advantages of Amazon EVS is that there is no real differentiation between an Amazon EVS vCenter and an on-premises vCenter. Both require the same connection and authentication information.

Steps to add a vCenter to each site

1. From the **Sites** option, select **Add vCenter** for the site you want.



2. In the Add vCenter server dialog box, select or provide the following information:
 - a. The NetApp Console agent hosted within your AWS VPC.
 - b. The IP address or FQDN for the vCenter to be added.
 - c. If different, change the port value to the TCP port used by your vCenter cluster manager.
 - d. The vCenter username for the account created earlier that will be used by NetApp Disaster Recovery to manage the vCenter.
 - e. The vCenter password for the provided username.
 - f. If your company uses an external Certificate Authority (CA) or the vCenter Endpoint Certificate Store to gain access to your vCenters, uncheck the **Use self-signed certificates** checkbox. Otherwise, leave the box checked.
3. Select **Add**.




Create replication plans for Amazon EVS

Create replication plans in NetApp Disaster Recovery overview

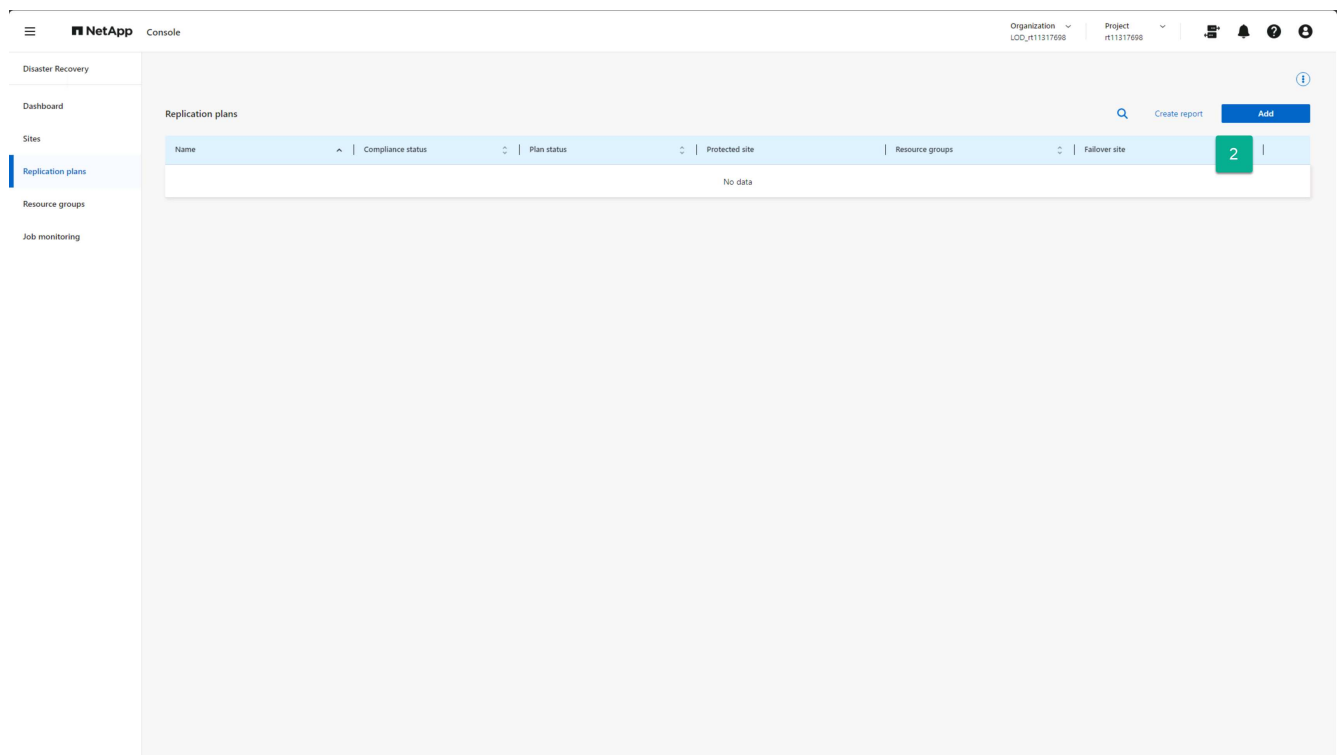
After you have vCenters to protect on the on-premises site and you have an Amazon EVS site configured to use Amazon FSx for NetApp ONTAP that you can use as a DR destination, you can create a replication plan (RP) to protect any set of VMs hosted on the vCenter cluster within your on-premises site.

To start the replication plan creation process:

1. From any NetApp Disaster Recovery screen, select the **Replication plans** option.



2. From the Replication plans page, select **Add**.



This opens the Create replication plan wizard.

Continue with [Create replication plan wizard Step 1](#).

Create a replication plan: Step 1 - Select vCenters in NetApp Disaster Recovery

First, using NetApp Disaster Recovery, provide a replication plan name and select the source and destination vCenters for the replication.

1. Enter a unique name for the replication plan.

Only alpha-numeric characters and underscores (_) are allowed for replication plan names.

2. Select a source vCenter cluster.
3. Select a destination vCenter cluster.
4. Select **Next**.

Continue with [Create replication plan wizard Step 2](#).

Create a replication plan: Step 2 - Select VM resources in NetApp Disaster Recovery

Select the virtual machines to be protected using NetApp Disaster Recovery.

There are several ways to select VMs for protection:

- **Select individual VMs:** Clicking on the **Virtual machines** button enables you to select individual VMs to protect. As you select each VM, the service adds it to a default resource group located on the right-hand side of the screen.
- **Select previously created resource groups:** You can create custom resource groups beforehand using the Resource group option from the NetApp Disaster Recovery menu. This is not a requirement as you can use the other two methods to create a resource group as part of the replication plan process. For details, see [Create a replication plan](#).
- **Select entire vCenter datastores:** If you have a lot of VMs to protect with this replication plan, it may not be as efficient to select individual VMs. Because NetApp Disaster Recovery uses volume-based SnapMirror replication to protect the VMs, all VMs residing on a datastore will be replicated as part of the volume. In most cases, you should have NetApp Disaster Recovery protect and restart any VMs located on the datastore. Use this option to tell the service to add any VMs hosted on a selected datastore to the list of protected VMs.

For this guided instruction, we select the entire vCenter datastore.

Steps to access this page

1. From the **Replication plan** page, continue to the **Applications** section.
2. Review the information in the **Applications** page that opens.

Steps to select the datastore or datastores:

1. Select **Datastores**.
2. Check the checkboxes beside each datastore you want to protect.
3. (Optionally) Rename the resource group to a suitable name by selecting the pencil icon next to the resource group name.

4. Select **Next**.

Continue with [Create replication plan wizard Step 3](#).



Create a replication plan: Step 3 - Map resources in NetApp Disaster Recovery

After you have a list of VMs that you want to protect using NetApp Disaster Recovery, provide failover mapping and VM configuration information to use during a failover.

You need to map four primary types of information:

- Compute resources
- Virtual networks
- VM reconfiguration
- Datastore mapping

Each VM requires the first three types of information. Datastore mapping is required for each datastore that hosts VMs to be protected.

- The sections with the caution icon () require that you provide mapping information.
- The section marked with the check icon () have been mapped or have default mappings. Review them to make sure that the current configuration meets your requirements.

Steps to access this page

1. From the **Replication plan** page, continue to the **Resource mapping** section.
2. Review the information on the **Resource mapping** page that opens.
3. To open each category of mappings required, select the down arrow (▼) beside the section.

Compute resource mapping

Because a site could host multiple virtual datacenters and multiple vCenter clusters, you need to identify which vCenter cluster to recover VMs on in the event of a failover.

Steps to map compute resources

1. Select the virtual datacenter from the list of datacenters located at the DR site.
2. Select the cluster to host the datastores and VMs from the list of clusters within the selected virtual datacenter.
3. (Optional) Select a target host in the target cluster.

This step is not required because NetApp Disaster Recovery selects the first host added to the cluster in vCenter. At that point, the VMs either continue to run on that ESXi host or VMware DRS moves the VM to a different ESXi host as needed based on DRS rules configured.

4. (Optional) Provide the name of a top-level vCenter folder to place the VM registrations into.

This is for your organizational needs and is not required.

Map virtual network resources

Each VM can have one or more virtual NICs connected to virtual networks within the vCenter network infrastructure. To ensure that each VM is properly connected to the desired networks upon restarting in the DR site, identify which DR site virtual networks to connect these VMs. Do this by mapping each virtual network in the on-premises site to an associated network on the DR site.

Select which destination virtual network to map each source virtual network

1. Select the Target segment from the drop-down list.
2. Repeat the previous step for each source virtual network listed.

Define options for VM reconfiguration during failover

Each VM might require modifications to work correctly in the DR vCenter site. The Virtual machines section enables you to provide the necessary changes.

By default, NetApp Disaster Recovery uses the same settings for each VM as used on the source on-premises site. This assumes that VMs will use the same IP address, virtual CPU, and virtual DRAM configuration.

Network reconfiguration

Supported IP address types are static and DHCP. For static IP addresses, you have the following Target IP settings:

- **Same as source:** As the name suggests, the service uses the same IP address on the destination VM that was used on the VM at the source site. This requires that you configure the virtual networks that were mapped in the previous step for the same subnet settings.
- **Different from source:** The service provides a set of IP address fields for each VM that must be configured for the appropriate subnet used on the destination virtual network, which you mapped in the previous section. For each VM you must provide an IP address, subnet mask, DNS, and default gateway values. Optionally, use the same subnet mask, DNS, and gateway settings for all VMs to simplify the process when all VMs attach to the same subnet.
- **Subnet mapping:** This option reconfigures each VM's IP address based on the destination virtual network's CIDR configuration. To use this feature, ensure that each vCenter's virtual networks have a defined CIDR setting within the service, as changed in the vCenter information in the Sites page.

After you configure subnets, Subnet mapping uses the same unit component of the IP address for both source and destination VM configuration, but replaces the subnet component of the IP address based on the provided CIDR information. This feature also requires that both the source and destination virtual networks have the same IP address class (the /xx component of the CIDR). This ensures that there are enough IP addresses available at the destination site to host all of the protected VMs.

For this EVS setup, we assume that the source and destination IP configurations are the same and do not require any additional reconfiguration.

Make changes to network settings reconfiguration

1. Select the type of IP addressing to use for failed over VMs.
2. (Optional) Provide a VM renaming scheme for restarted VMs by providing an optional prefix and suffix

value.

VM compute resource reconfiguration

There are several options for reconfiguring VM compute resources. NetApp Disaster Recovery supports changing the number of virtual CPUs, the amount of virtual DRAM, and the VM name.

Specify any VM configuration changes

1. (Optional) Modify the number of virtual CPUs each VM should use. This might be needed if your DR vCenter cluster hosts do not have as many CPU cores as the source vCenter cluster.
2. (Optional) Modify the amount of virtual DRAM each VM should use. This might be needed if your DR vCenter cluster hosts do not have as much physical DRAM as the source vCenter cluster hosts.

Boot order

NetApp Disaster Recovery supports an ordered restart of VMs based on a boot order field. The Boot order field indicates how the VMs in each resource group start. Those VMs with the same value in the Boot order field boot in parallel.

Modify the boot order settings

1. (Optionally) Modify the order you would like your VMs to be restarted. This field takes any numeric value. NetApp Disaster Recovery tries to restart VMs that have the same numeric value in parallel.
2. (Optionally) Provide a delay to be used between each VM restart. The time is injected after this VM's restart has completed and before the VM(s) with the next higher boot order number. This number is in minutes.

Custom guest OS operations

NetApp Disaster Recovery supports performing some guest OS operations for each VM:

- NetApp Disaster Recovery can take application-consistent backups of VMs for VMs running Oracle databases and Microsoft SQL Server databases.
- NetApp Disaster Recovery can execute custom defined scripts suitable for the guest OS for each VM. Executing such scripts requires user credentials acceptable to the guest OS with ample privileges to execute the operations listed in the script.

Modify each VM's custom guest OS operations

1. (Optional) Check the **Create application consistent replicas** checkbox if the VM is hosting an Oracle or SQL Server database.
2. (Optional) To take custom actions within the guest OS as part of the startup process, upload a script for any VMs. To run a single script in all VMs, use the checkbox highlighted and complete the fields.
3. Certain configuration changes require user credentials with adequate permissions to perform the operations. Provide credentials in the following cases:
 - A script will be executed within the VM by the guest OS.

- An application-consistent snapshot needs to be performed.

Map datastores

The final step in creating a replication plan is identifying how ONTAP should protect the datastores. These settings define the replication plans recovery point objective (RPO), how many backups should be maintained, and where to replicate each vCenter datastore's hosting ONTAP volumes.

By default, NetApp Disaster Recovery manages its own snapshot replication schedule; however, optionally, you can specify that you would like to use the existing SnapMirror replication policy schedule for datastore protection.

In addition, you can optionally customize which data LIFs (logical interfaces) and export policy to use. If you don't provide these settings, NetApp Disaster Recovery uses all data LIFs associated with the appropriate protocol (NFS, iSCSI, or FC) and uses the default export policy for NFS volumes.

To configure datastore (volume) mapping

1. (Optional) Decide whether you want to use an existing ONTAP SnapMirror replication schedule or have NetApp Disaster Recovery manage protection of your VMs (default).
2. Provide a starting point for when the service should start taking backups.
3. Specify how often the service should take a backup and replicate it to the DR destination Amazon FSx for NetApp ONTAP cluster.
4. Specify how many historical backups should be retained. The service maintains the same number of backups on the source and destination storage cluster.
5. (Optional) Select a default logical interface (data LIFs) for each volume. If none is selected, all the data LIFs in the destination SVM that support the volume access protocol are configured.
6. (Optional) Select an export policy for any NFS volumes. If not selected, the default export policy is used

Continue with [Create replication plan wizard Step 4](#).

Create a replication plan: Step 4 - Verify settings in NetApp Disaster Recovery


After you add the replication plan information in NetApp Disaster Recovery, verify that the information you entered is correct.

Steps

1. Select **Save** to review your settings before activating the replication plan.

You can select each tab to review the settings and make changes on any tab by selecting the pencil icon.

Replication plan settings review




2. When you are satisfied that all settings are correct, select **Add plan** at the bottom of the screen.

Continue with [Verify the replication plan](#).

Verify that everything is working in NetApp Disaster Recovery

After you add the replication plan in NetApp Disaster Recovery, you return to the Replication plans page where you can view your replication plans and their status. You should verify that the replication plan is in the **Healthy** state. If it is not, you should check the status of the replication plan and correct any issues before proceeding.

Figure: Replication plans page



NetApp Disaster Recovery performs a series of tests to verify that all the components (ONTAP cluster, vCenter clusters, and VMs) are accessible and in the proper state for the service to protect the VMs. This is called a compliance check, and it is run on a regular basis.

From the Replication plans page, you can see the following information:

- Status of the last compliance check
- The replication plan's replication state
- The name of the protected (source) site
- The list of resource groups protected by the replication plan
- The name of the failover (destination) site

Perform replication plan operations with NetApp Disaster Recovery

Use NetApp Disaster Recovery with Amazon EVS and Amazon FSx for NetApp ONTAP to perform the following operations: failover, test failover, refresh resources, migrate, take a snapshot now, disable/enable replication plan, clean up old snapshots, reconcile snapshots, delete replication plan, and edit schedules.

Fail over


The primary operation that you might need to perform is the one you hope never happens: failing over to the DR (destination) datacenter in the event of a catastrophic failure at the production on-premises site.

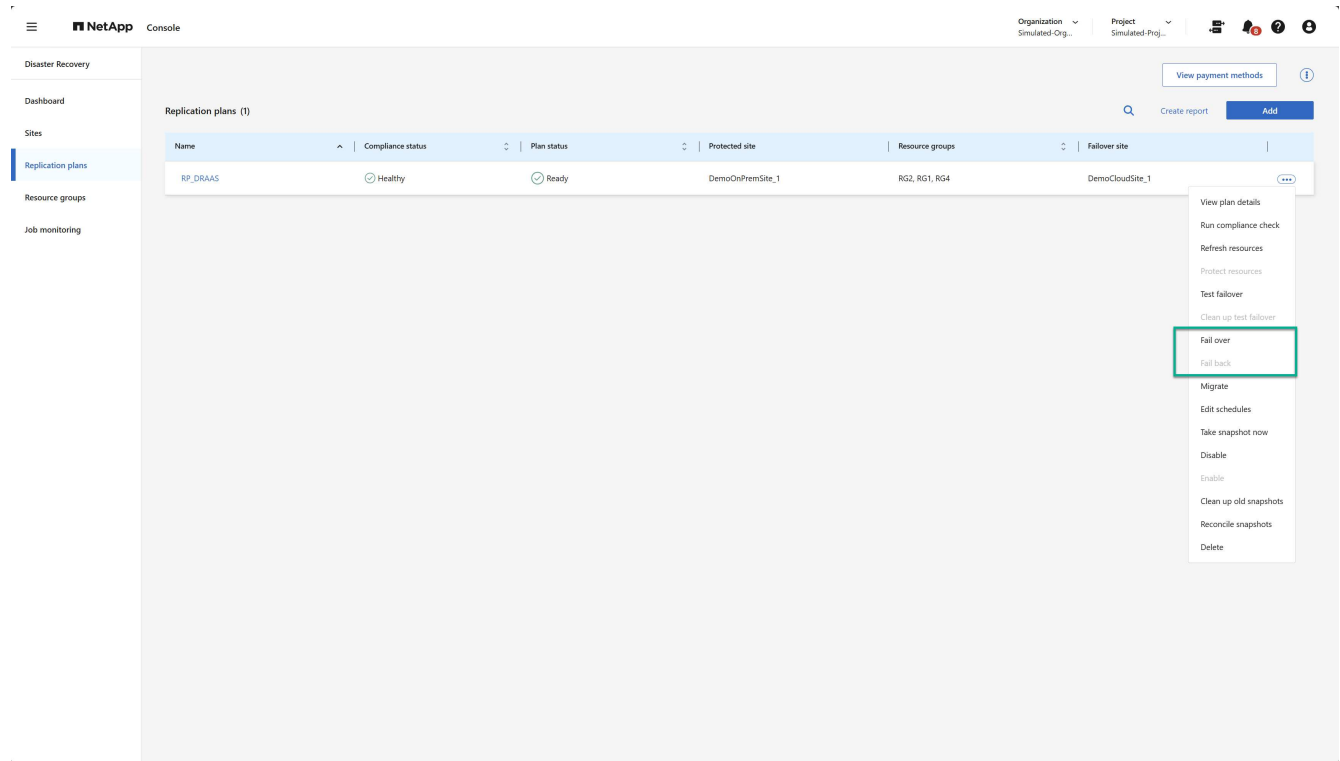
Failover is a manually initiated process.

Steps to access the failover operation


1. From the NetApp Console left navigation bar, select **Protection > Disaster Recovery**.
2. From the NetApp Disaster Recovery menu, select **Replication plans**.

Steps to perform a failover

1. From the Replication plans page, select the replication plan's Actions option .
2. Select **Fail over**.



3. If the production (protected) site is not accessible, select a previously created snapshot as your recovery image. To do this, select **Select**.
4. Select the backup to be used for the recovery.
5. (Optional) Select whether you want NetApp Disaster Recovery to force the failover process regardless of the state of the replication plan. This should only be done as a last resort.
6. (Optional) Select whether you want NetApp Disaster Recovery to automatically create a reverse protection relationship after the production site has been recovered.
7. Type the word "Failover" to verify that you would like to proceed.
8. Select **Failover**.



Test failover


A test failover is similar to a failover except for two differences.


- The production site is still active and all VMs are still operating as expected.
- NetApp Disaster Recovery protection of the production VMs continues.

This is accomplished by using native ONTAP FlexClone volumes at the destination site. To learn more about test failover, see [Fail over applications to a remote site | NetApp Documentation](#).


The steps for executing a test failover are identical to those used to execute a real failover except that you use the Test failover operation on the replication plan's context menu.

Steps

1. Select the replication plan's Actions option .
2. Select **Test failover** from the menu.




3. Decide if you want get the latest state of the production environment (Take snapshot now) or use a previously created replication plan backup (Select)
4. If you chose a previously created backup, then select the backup to be used for the recovery.
5. Type the word “Test failover” to verify that you would like to proceed.
6. Select **Test failover**.

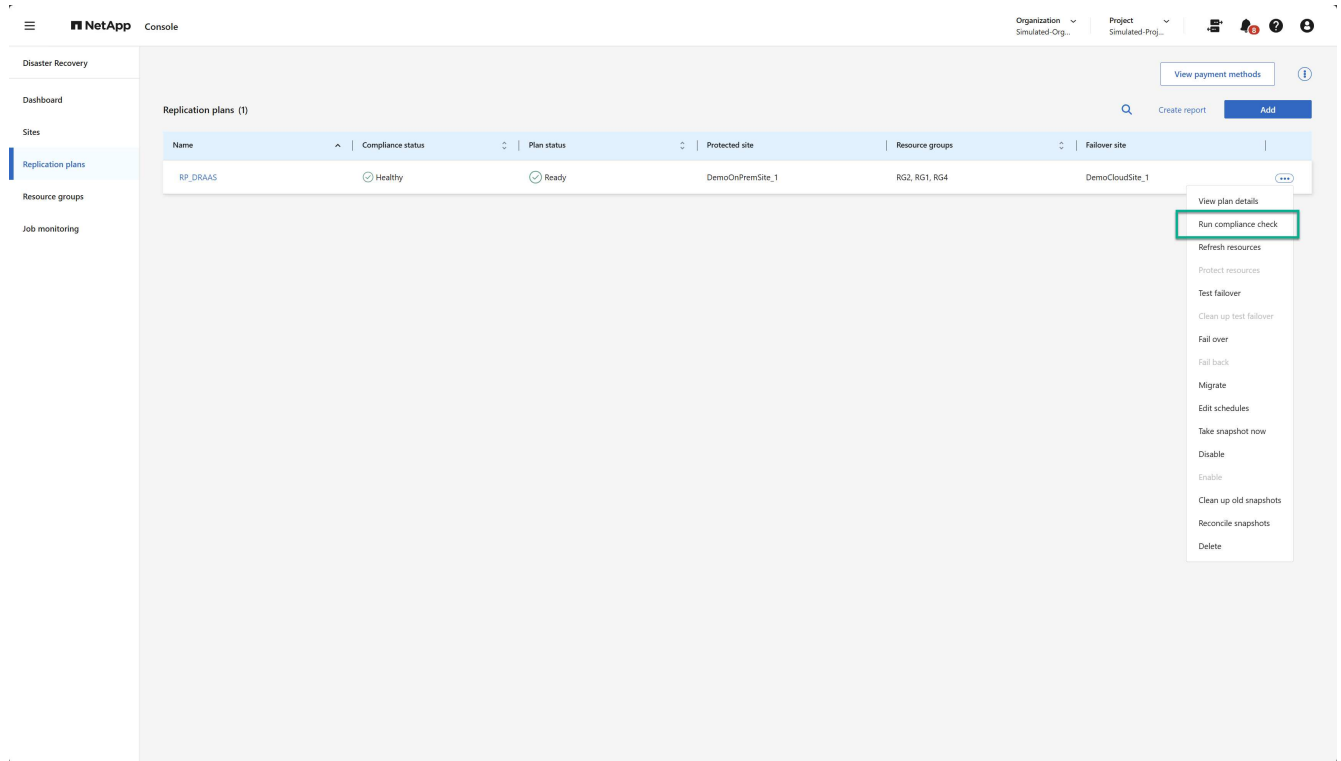


Run a compliance check

Compliance checks are run every three hours, by default. At any time, you might want to manually run a compliance check.

Steps

1. Select the **Actions** option  next to the replication plan.
2. Select the **Run compliance check** option from the replication plan's Actions menu:




3. To change how often NetApp Disaster Recovery automatically runs compliance checks, select **Edit schedules** option from the replication plan's Actions menu.

Refresh resources

Any time you make changes to your virtual infrastructure — such as adding or deleting VMs, adding or deleting datastores, or moving VMs between datastores — you need to perform a refresh of the impacted vCenter clusters in NetApp Disaster Recovery service. The service does this automatically once every 24 hours by default, but a manual refresh ensures that the latest virtual infrastructure information is available and taken into account for DR protection.


There are two instances where a refresh is necessary:


- vCenter refresh: Perform a vCenter refresh anytime VMs are added or deleted from or moved out of a vCenter cluster:
- Replication plan refresh: Perform a replication plan refresh anytime a VM is moved between datastores in the same source vCenter cluster.



Migrate

While NetApp Disaster Recovery is primarily used for disaster recovery use cases, it can also enable one-time moves of a set of VMs from the source site to the destination site. This could be for a concerted migration to cloud project or it could be used for disaster avoidance — such as bad weather, political strife, or other potential temporary catastrophic events.


1. Select the **Actions** option  next to the replication plan.
2. To move the VMs in a replication plan to the destination Amazon EVS cluster, select **Migrate** from the replication plan's Actions menu:

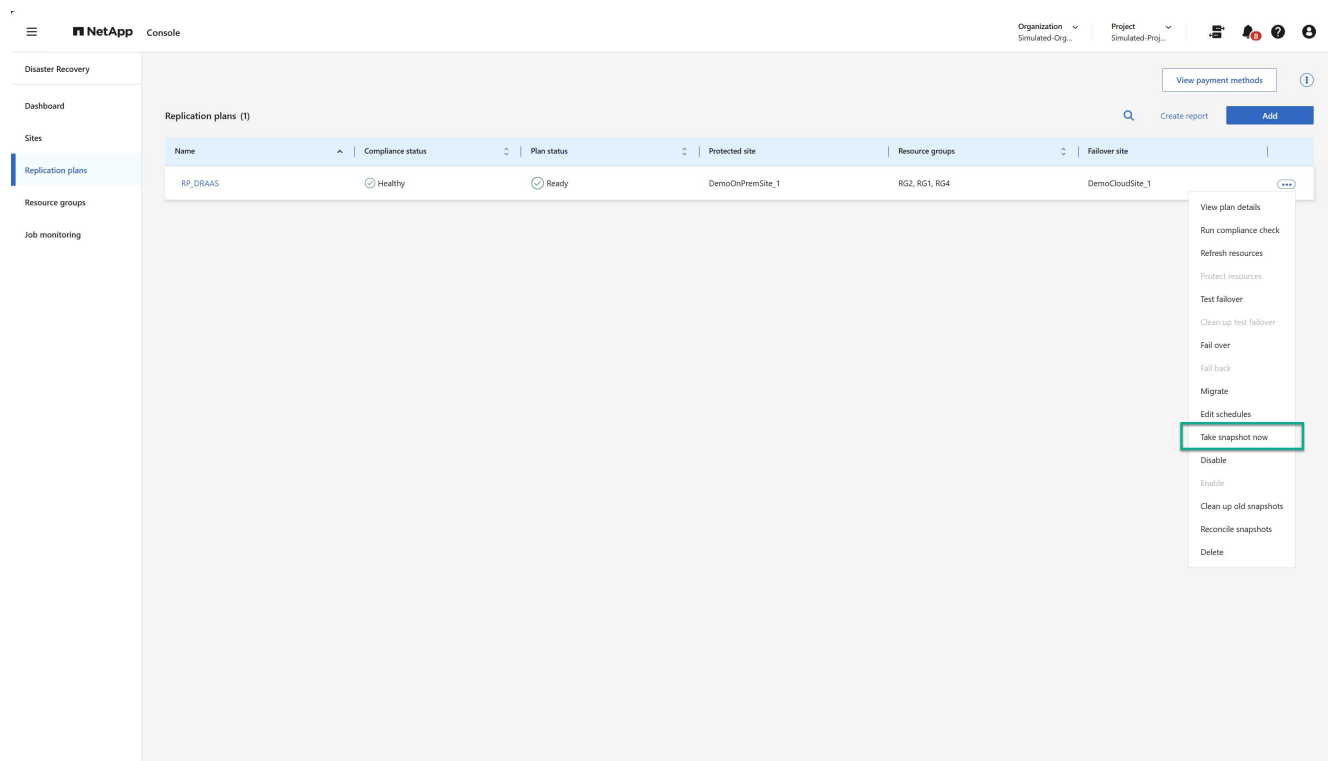


3. Enter information in the Migrate dialog box.

Take a snapshot now

At any time, you can take an immediate snapshot of the replication plan. This snapshot is included in the NetApp Disaster Recovery considerations set by the replication plan's snapshot retention count.

1. Select the **Actions** option  next to the replication plan.
2. To take an immediate snapshot of the replication plan's resources, select **Take snapshot now** on the replication plan's Actions menu:




Disable or enable replication plan

You might need to temporarily stop the replication plan to perform some operation or maintenance that could impact the replication process. The service provides a method to stop and start replication.


1. To temporarily stop replication, select **Disable** on the replication plan's Actions menu.
2. To restart replication, select **Enable** on the replication plan's Actions menu.

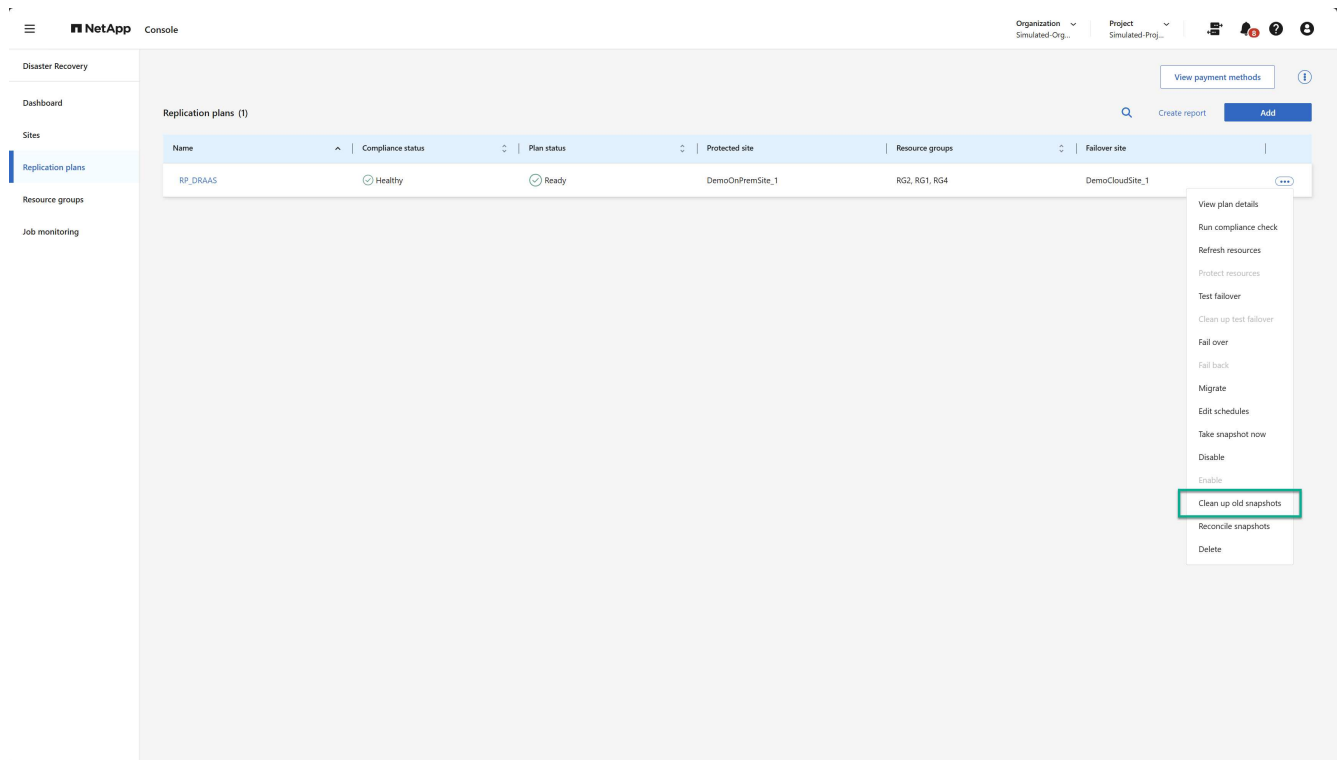
When the replication plan is active, the **Enable** command is grayed out. When the replication plan is disabled, the **Disable** command is grayed out.



Clean up old snapshots


You might want to clean up older snapshots that have been retained on the source and destination sites. This can happen if the replication plan's snapshot retention count is altered.


1. Select the **Actions** option  next to the replication plan.
2. To remove these older snapshots manually, select **Clean up old snapshots** from the replication plan's Actions menu.



Reconcile snapshots


Because the service orchestrates ONTAP volume snapshots, it is possible for an ONTAP storage administrator to directly delete snapshots using either ONTAP System Manager, the ONTAP CLI, or the ONTAP REST APIs without the service's knowledge. The service automatically deletes any snapshots on the source that are not on the destination cluster automatically every 24 hours. However, you can perform this on demand. This feature enables you to ensure that the snapshots are consistent across all sites.


1. Select the **Actions** option  next to the replication plan.
2. To delete snapshots from the source cluster that do not exist on the destination cluster, select **Reconcile snapshots** from the replication plan's Actions menu.



Delete replication plan


If the replication plan is no longer needed, you can delete it.

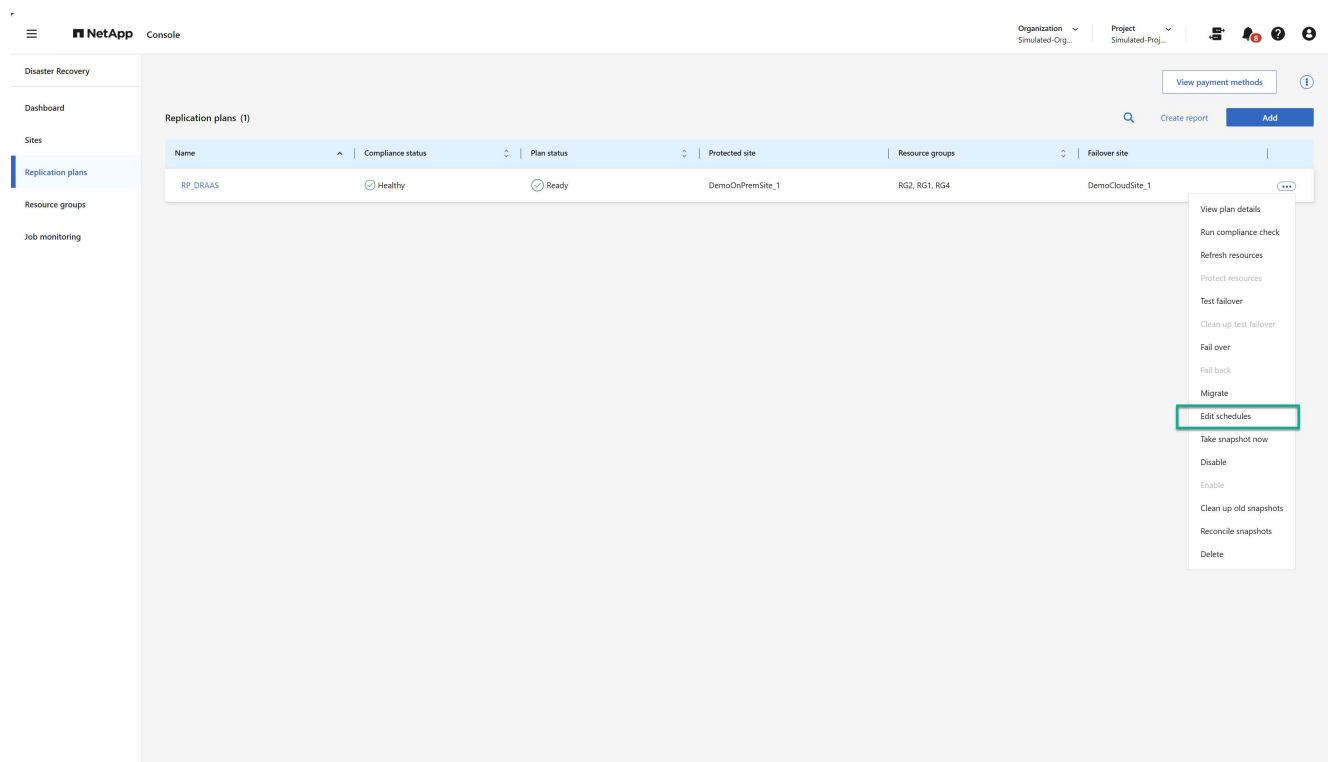
1. Select the **Actions** option  next to the replication plan.
2. To delete the replication plan, select **Delete** from the replication plan's context menu.



Edit schedules

Two operations are performed automatically on a regular schedule: test failovers and compliance checks.

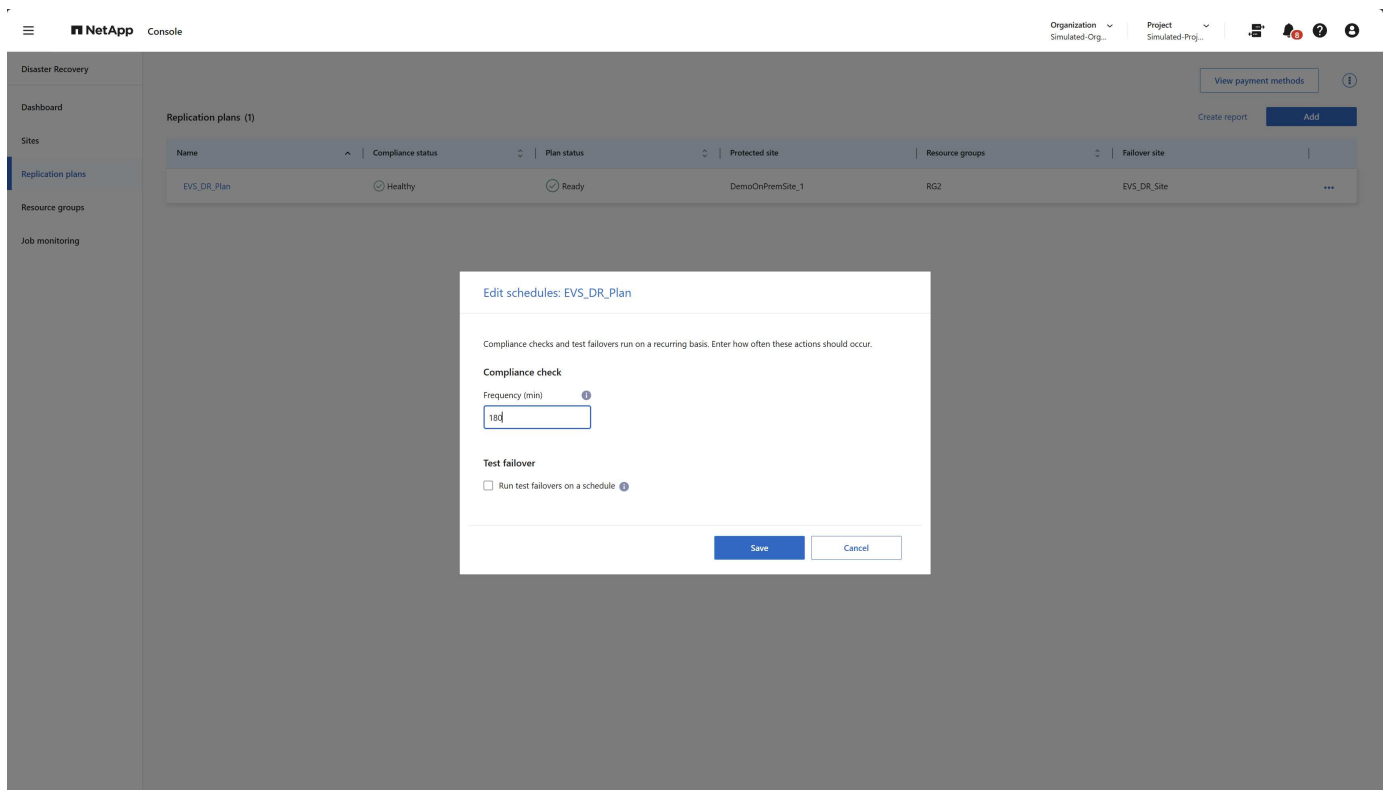
1. Select the **Actions** option  next to the replication plan.
2. To change these schedules for either of these two operations, select **Edit schedules** for the replication plan.



Change compliance check interval

By default, compliance checks are performed every three hours. You can change this to any interval between 30 minutes and 24 hours.


To change this interval, change the Frequency field in the Edit schedules dialog box:



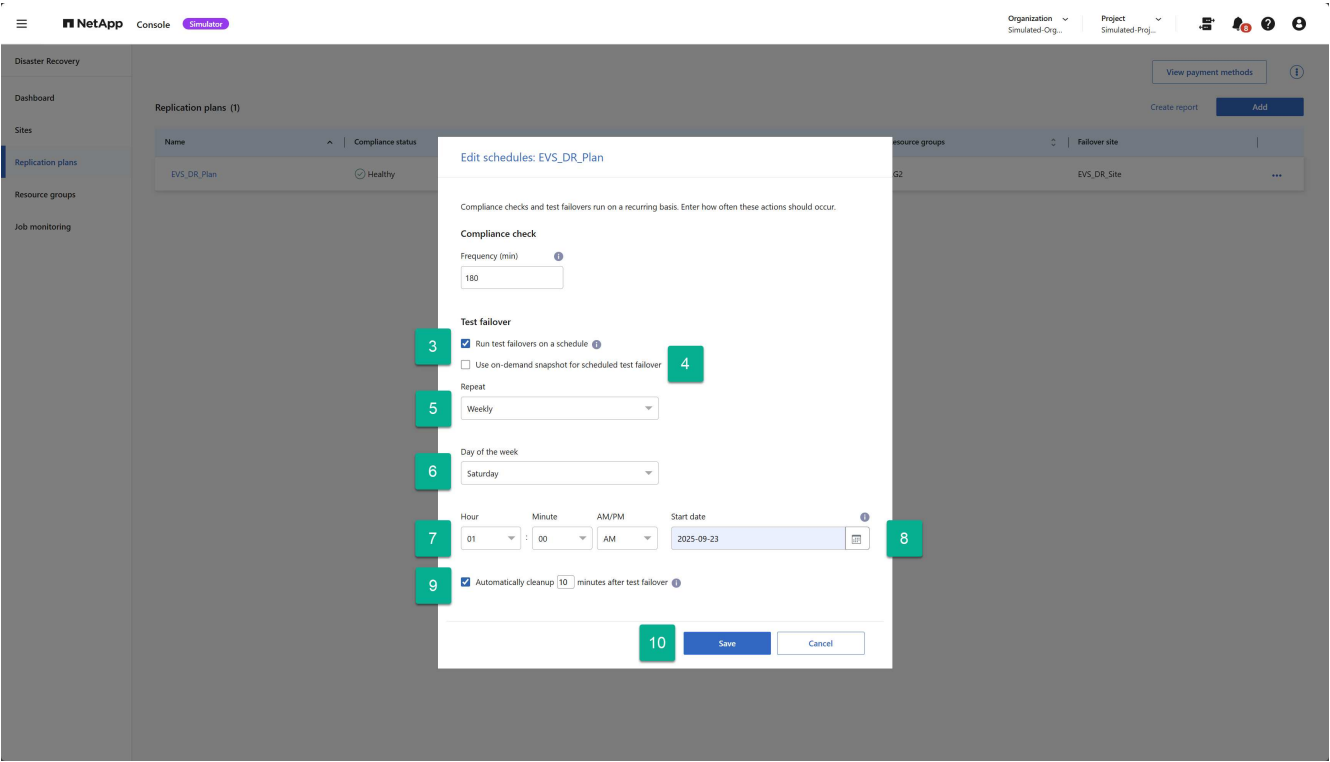
Schedule automated test failovers

Test failovers are manually executed by default. You can schedule automatic test failovers, which helps ensure that your replication plans perform as expected. To learn more about the test failover process, see [Test the failover process](#).

Steps to schedule test failovers

1. Select the **Actions** option  next to the replication plan.
2. Select **Run failover**.
3. Check the **Run test failovers on a schedule** checkbox.
4. (Optional) Check the **Use on-demand-snapshot for scheduled test failover**.
5. Select an interval type in the Repeat drop-down.
6. Select when to perform the test failover
 - a. Weekly: select the Day of the Week
 - b. Monthly: select the Day of the month
7. Choose the time of day to run the test failover
8. Chose the start date.
9. Decide if you want the service to automatically clean up the test environment and how long you would like the test environment to run before the clean up process starts.

10. Select **Save**.



Frequently asked questions for NetApp Disaster Recovery

This FAQ can help if you're just looking for a quick answer to a question.

What's the NetApp Disaster Recovery URL?

For the URL, in a browser, enter: <https://console.netapp.com/> to access the NetApp console.

Do you need a license to use NetApp Disaster Recovery?

A NetApp Disaster Recovery license is required for complete access. However, you can try it out with the free trial.

For details about setting up licensing for NetApp Disaster Recovery, refer to [Set up NetApp Disaster Recovery licensing](#).

How do you access NetApp Disaster Recovery?

NetApp Disaster Recovery does not require any enablement. The disaster recovery option automatically appears on the NetApp Console left navigation.

Knowledge and support

Register for support

Support registration is required to receive technical support specific to BlueXP and its storage solutions and services. Support registration is also required to enable key workflows for Cloud Volumes ONTAP systems.

Registering for support does not enable NetApp support for a cloud provider file service. For technical support related to a cloud provider file service, its infrastructure, or any solution using the service, refer to "Getting help" in the BlueXP documentation for that product.

- [Amazon FSx for ONTAP](#)
- [Azure NetApp Files](#)
- [Google Cloud NetApp Volumes](#)

Support registration overview

There are two forms of registration to activate support entitlement:

- Registering your BlueXP account serial number (your 20 digit 960xxxxxxx serial number located on the Support Resources page in BlueXP).

This serves as your single support subscription ID for any service within BlueXP. Each BlueXP account-level support subscription must be registered.

- Registering the Cloud Volumes ONTAP serial numbers associated with a subscription in your cloud provider's marketplace (these are 20 digit 909201xxxxxxx serial numbers).

These serial numbers are commonly referred to as *PAYGO serial numbers* and get generated by BlueXP at the time of Cloud Volumes ONTAP deployment.

Registering both types of serial numbers enables capabilities like opening support tickets and automatic case generation. Registration is completed by adding NetApp Support Site (NSS) accounts to BlueXP as described below.

Register BlueXP for NetApp support

To register for support and activate support entitlement, one user in your BlueXP organization (or account) must associate a NetApp Support Site account with their BlueXP login. How you register for NetApp support depends on whether you already have a NetApp Support Site (NSS) account.

Existing customer with an NSS account

If you're a NetApp customer with an NSS account, you simply need to register for support through BlueXP.

Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.
2. Select **User Credentials**.

3. Select **Add NSS credentials** and follow the NetApp Support Site (NSS) Authentication prompt.
4. To confirm that the registration process was successful, select the Help icon, and select **Support**.

The **Resources** page should show that your BlueXP organization is registered for support.

Note that other BlueXP users will not see this same support registration status if they have not associated a NetApp Support Site account with their BlueXP login. However, that doesn't mean that your BlueXP organization is not registered for support. As long as one user in the organization has followed these steps, then your organization has been registered.

Existing customer but no NSS account

If you're an existing NetApp customer with existing licenses and serial numbers but *no* NSS account, you need to create an NSS account and associate it with your BlueXP login.

Steps


1. Create a NetApp Support Site account by completing the [NetApp Support Site User Registration form](#)
 - a. Be sure to select the appropriate User Level, which is typically **NetApp Customer/End User**.
 - b. Be sure to copy the BlueXP account serial number (960xxxx) used above for the serial number field. This will speed up the account processing.
2. Associate your new NSS account with your BlueXP login by completing the steps under [Existing customer with an NSS account](#).

Brand new to NetApp

If you are brand new to NetApp and you don't have an NSS account, follow each step below.

Steps

1. In the upper right of the BlueXP console, select the Help icon, and select **Support**.



2. Locate your account ID serial number from the Support Registration page.



96015585434285107893
Account serial number

⚠ Not Registered

Add your NetApp Support Site (NSS) [credentials](#) to BlueXP
Follow these [instructions](#) to register for support in case you don't have an NSS account yet.

3. Navigate to [NetApp's support registration site](#) and select **I am not a registered NetApp Customer**.
4. Fill out the mandatory fields (those with red asterisks).
5. In the **Product Line** field, select **Cloud Manager** and then select your applicable billing provider.
6. Copy your account serial number from step 2 above, complete the security check, and then confirm that you read NetApp's Global Data Privacy Policy.

An email is immediately sent to the mailbox provided to finalize this secure transaction. Be sure to check your spam folders if the validation email doesn't arrive in few minutes.

7. Confirm the action from within the email.

Confirming submits your request to NetApp and recommends that you create a NetApp Support Site account.

8. Create a NetApp Support Site account by completing the [NetApp Support Site User Registration form](#)
 - a. Be sure to select the appropriate User Level, which is typically **NetApp Customer/End User**.
 - b. Be sure to copy the account serial number (960xxxx) used above for the serial number field. This will speed up processing.

After you finish

NetApp should reach out to you during this process. This is a one-time onboarding exercise for new users.

Once you have your NetApp Support Site account, associate the account with your BlueXP login by completing the steps under [Existing customer with an NSS account](#).

Associate NSS credentials for Cloud Volumes ONTAP support

Associating NetApp Support Site credentials with your BlueXP organization is required to enable the following key workflows for Cloud Volumes ONTAP:

- Registering pay-as-you-go Cloud Volumes ONTAP systems for support

Providing your NSS account is required to activate support for your system and to gain access to NetApp technical support resources.

- Deploying Cloud Volumes ONTAP when you bring your own license (BYOL)

Providing your NSS account is required so that BlueXP can upload your license key and to enable the subscription for the term that you purchased. This includes automatic updates for term renewals.

- Upgrading Cloud Volumes ONTAP software to the latest release


Associating NSS credentials with your BlueXP organization is different than the NSS account that is associated with a BlueXP user login.

These NSS credentials are associated with your specific BlueXP organization ID. Users who belong to the BlueXP organization can access these credentials from **Support > NSS Management**.

- If you have a customer-level account, you can add one or more NSS accounts.
- If you have a partner or reseller account, you can add one or more NSS accounts, but they can't be added alongside customer-level accounts.

Steps

1. In the upper right of the BlueXP console, select the Help icon, and select **Support**.



2. Select **NSS Management > Add NSS Account**.
3. When you're prompted, select **Continue** to be redirected to a Microsoft login page.

NetApp uses Microsoft Entra ID as the identity provider for authentication services specific to support and licensing.

4. At the login page, provide your NetApp Support Site registered email address and password to perform the authentication process.

These actions enable BlueXP to use your NSS account for things like license downloads, software upgrade verification, and future support registrations.


Note the following:

- The NSS account must be a customer-level account (not a guest or temp account). You can have multiple customer-level NSS accounts.
- There can be only one NSS account if that account is a partner-level account. If you try to add customer-level NSS accounts and a partner-level account exists, you'll get the following error message:

"The NSS customer type is not allowed for this account as there are already NSS Users of different type."

The same is true if you have pre-existing customer-level NSS accounts and try to add a partner-level account.

- Upon successful login, NetApp will store the NSS user name.

This is a system-generated ID that maps to your email. On the **NSS Management** page, you can display your email from the  menu.

- If you ever need to refresh your login credential tokens, there is also an **Update Credentials** option in the **...** menu.

Using this option prompts you to log in again. Note that the token for these accounts expire after 90 days. A notification will be posted to alert you of this.

Get help

NetApp provides support for BlueXP and its cloud services in a variety of ways. Extensive free self-support options are available 24/7, such as knowledgebase (KB) articles and a community forum. Your support registration includes remote technical support via web ticketing.

Get support for a cloud provider file service

For technical support related to a cloud provider file service, its infrastructure, or any solution using the service, refer to "Getting help" in the BlueXP documentation for that product.

- [Amazon FSx for ONTAP](#)
- [Azure NetApp Files](#)
- [Google Cloud NetApp Volumes](#)

To receive technical support specific to BlueXP and its storage solutions and services, use the support options described below.

Use self-support options

These options are available for free, 24 hours a day, 7 days a week:

- **Documentation**

The BlueXP documentation that you're currently viewing.

- [Knowledge base](#)

Search through the BlueXP knowledge base to find helpful articles to troubleshoot issues.

- [Communities](#)

Join the BlueXP community to follow ongoing discussions or create new ones.

Create a case with NetApp support

In addition to the self-support options above, you can work with a NetApp Support specialist to resolve any issues after you activate support.

Before you get started

- To use the **Create a Case** capability, you must first associate your NetApp Support Site credentials with your BlueXP login. [Learn how to manage credentials associated with your BlueXP login.](#)
- If you're opening a case for an ONTAP system that has a serial number, then your NSS account must be

associated with the serial number for that system.

Steps

1. In BlueXP, select **Help > Support**.
2. On the **Resources** page, choose one of the available options under Technical Support:
 - a. Select **Call Us** if you'd like to speak with someone on the phone. You'll be directed to a page on netapp.com that lists the phone numbers that you can call.
 - b. Select **Create a Case** to open a ticket with a NetApp Support specialist:
 - **Service:** Select the service that the issue is associated with. For example, BlueXP when specific to a technical support issue with workflows or functionality within the service.
 - **Working Environment:** If applicable to storage, select **Cloud Volumes ONTAP** or **On-Prem** and then the associated working environment.

The list of working environments are within scope of the BlueXP organization (or account), project (or workspace), and Connector you have selected in the top banner of the service.

- **Case Priority:** Choose the priority for the case, which can be Low, Medium, High, or Critical.

To learn more details about these priorities, hover your mouse over the information icon next to the field name.

- **Issue Description:** Provide a detailed description of your problem, including any applicable error messages or troubleshooting steps that you performed.
- **Additional Email Addresses:** Enter additional email addresses if you'd like to make someone else aware of this issue.
- **Attachment (Optional):** Upload up to five attachments, one at a time.

Attachments are limited to 25 MB per file. The following file extensions are supported: txt, log, pdf, jpg/jpeg, rtf, doc/docx, xls/xlsx, and csv.

ntapitdemo
NetApp Support Site Account

Service

Select

Working Enviroment

Select

Case Priority

Low - General guidance

Issue Description

Provide detailed description of problem, applicable error messages and troubleshooting steps taken.

Additional Email Addresses (Optional)

Type here

Attachment (Optional)

No files selected

Upload

After you finish

A pop-up will appear with your support case number. A NetApp Support specialist will review your case and get back to you soon.

For a history of your support cases, you can select **Settings > Timeline** and look for actions named "create support case." A button to the far right lets you expand the action to see details.

It's possible that you might encounter the following error message when trying to create a case:

"You are not authorized to Create a Case against the selected service"

This error could mean that the NSS account and the company of record it's associated with is not the same company of record for the BlueXP account serial number (ie. 960xxxx) or the working environment serial number. You can seek assistance using one of the following options:

- Use the in-product chat
- Submit a non-technical case at <https://mysupport.netapp.com/site/help>

Manage your support cases (Preview)

You can view and manage active and resolved support cases directly from BlueXP. You can manage the cases associated with your NSS account and with your company.

Case management is available as a Preview. We plan to refine this experience and add enhancements in upcoming releases. Please send us feedback by using the in-product chat.

Note the following:

- The case management dashboard at the top of the page offers two views:
 - The view on the left shows the total cases opened in the past 3 months by the user NSS account you provided.
 - The view on the right shows the total cases opened in the past 3 months at your company level based on your user NSS account.

The results in the table reflect the cases related to the view that you selected.

- You can add or remove columns of interest and you can filter the contents of columns like Priority and Status. Other columns provide just sorting capabilities.

View the steps below for more details.

- At a per-case level, we offer the ability to update case notes or close a case that is not already in Closed or Pending Closed status.

Steps

1. In BlueXP, select **Help > Support**.
2. Select **Case Management** and if you're prompted, add your NSS account to BlueXP.

The **Case management** page shows open cases related to the NSS account that is associated with your BlueXP user account. This is the same NSS account that appears at the top of the **NSS management** page.

3. Optionally modify the information that displays in the table:
 - Under **Organization's cases**, select **View** to view all cases associated with your company.
 - Modify the date range by choosing an exact date range or by choosing a different time frame.

Search: Cases opened on the last 3 months Create a case

Date created	Last updated		Status (5)	
December 22, 2022	December 29, 2022	Last 7 days	Assigned	...
December 21, 2022	December 28, 2022	Last 30 days	Active	...
December 15, 2022	December 27, 2022	Last 3 months	Pending customer	...
December 14, 2022	December 26, 2022	Medium (P3)	Solution proposed	...
		Low (P4)		

Apply Reset

- Filter the contents of the columns.

Search: Cases opened on the last 3 months Create a case

Last updated	Priority	Status (5)	
December 29, 2022	Critical (P1)	Active	...
December 28, 2022	High (P2)	Pending customer	...
December 27, 2022	Medium (P3)	Solution proposed	...
December 26, 2022	Low (P4)	Pending closed	...
		Closed	...

Apply Reset

- Change the columns that appear in the table by selecting + and then choosing the columns that you'd like to display.

Search: Cases opened on the last 3 months Create a case

Last updated	Priority	Status (5)	
December 29, 2022	Critical (P1)	Last updated	...
December 28, 2022	High (P2)	Priority	...
December 27, 2022	Medium (P3)	Cluster name	...
December 26, 2022	Low (P4)	Case owner	...
		Opened by	...


Apply Reset

4. Manage an existing case by selecting ... and selecting one of the available options:

- **View case:** View full details about a specific case.
- **Update case notes:** Provide additional details about your problem or select **Upload files** to attach up to a maximum of five files.

Attachments are limited to 25 MB per file. The following file extensions are supported: txt, log, pdf, jpg/jpeg, rtf, doc/docx, xls/xlsx, and csv.

- **Close case:** Provide details about why you're closing the case and select **Close case**.



Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<https://www.netapp.com/company/legal/copyright/>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

[Notice for NetApp Disaster Recovery](#)

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.