



Use BlueXP disaster recovery

BlueXP disaster recovery

NetApp
August 29, 2024

Table of Contents

- Use BlueXP disaster recovery 1
 - Use BlueXP disaster recovery overview 1
 - View the health of your disaster recovery plans on the Dashboard 1
 - Add vCenter sites 2
 - Create a replication plan 4
 - Replicate applications to another site 12
 - Migrate applications to another site 13
 - Fail over applications to a remote site 13
 - Fail back applications to the original source 16
 - Manage sites, resource groups, replication plans, datastores and virtual machines information 16
 - Monitor disaster recovery jobs 19

Use BlueXP disaster recovery

Use BlueXP disaster recovery overview

Using BlueXP disaster recovery, you can accomplish the following goals:


- [View the health of your disaster recovery plans](#)
- [Add vCenter sites](#).
- [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 disaster recovery plans on the Dashboard

Using the BlueXP 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.

Steps

1. From the BlueXP left nav, select **Protection > Disaster recovery**.
2. From the BlueXP disaster recovery top menu, select **Dashboard**.



3. 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."

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

Add vCenter sites

Before you can create a disaster recovery plan, you need to add a primary vCenter site and a target vCenter disaster recovery site in BlueXP.

After they are added, BlueXP 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.

Steps

1. Log in to BlueXP and select **Protection > Disaster recovery** from the left nav.

You'll land on BlueXP 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.

2. **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 BlueXP Connector.

Site: BlueXP Connector:

vCenter IP address:

vCenter user name: vCenter password:

Use self-signed certificates

- a. Add a site, select the BlueXP Connector, and provide vCenter credentials.
- b. 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.

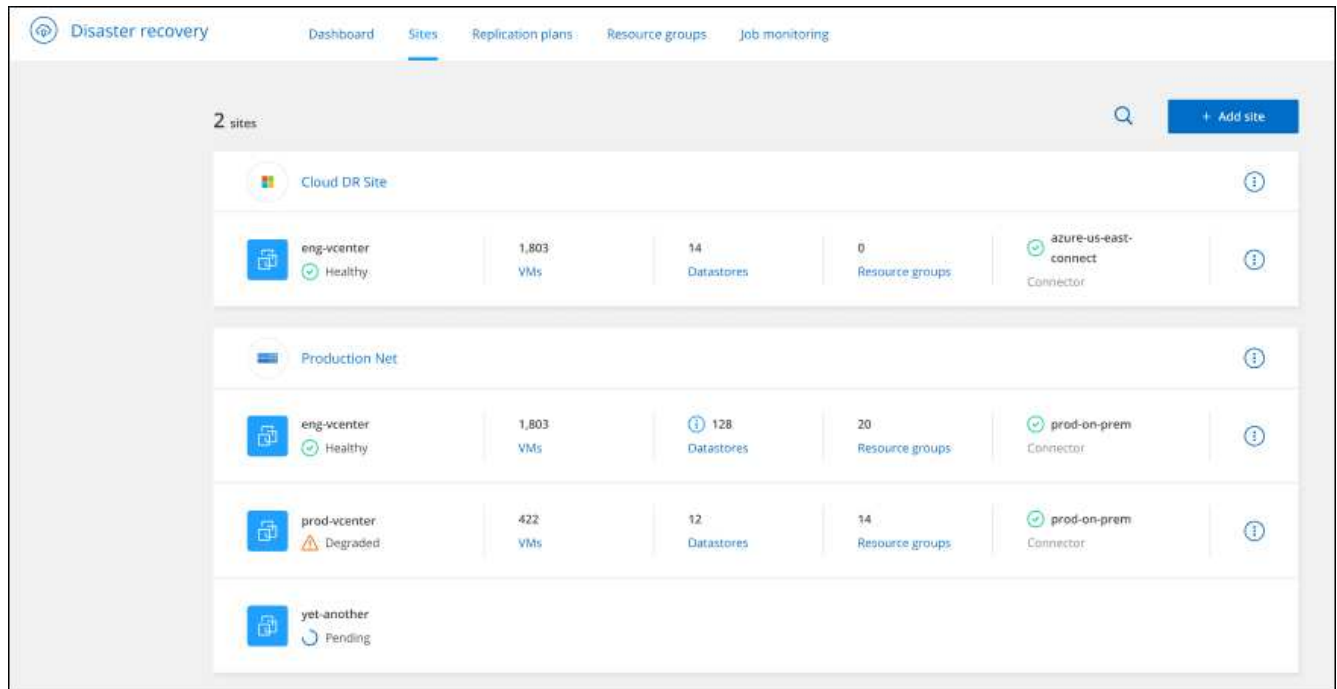
- c. Select **Add**.

Next, you will add a target vCenter.

3. Target:

- a. Choose the target site and the location. If the target is cloud, select **AWS**.
- b. Select **Add**.

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



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

Create a replication plan

After you've added vCenter sites, you're ready to create a disaster recovery or *replication plan*. 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 associated with one application or you might group applications that have similar tiers.

Such plans are sometimes called *blueprints*.

You can create a replication plan and also edit schedules for compliance and testing.

Create the plan

A wizard takes you through these steps:

- Select vCenter servers
- Select the VMs you want to replicate and assign groups
- Map how resources from the source environment map to the destination.
- Identify recurrence
- Review the plan

While you are creating the replication plan, 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 in a Consistency Group
- Many to many

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 BlueXP disaster recovery.

Select vCenter servers

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

Steps

1. From the BlueXP left nav, select **Protection > Disaster recovery**.
2. From the BlueXP disaster recovery top menu, select **Replication plans** and select **Add**. Or, if you are just beginning to use the service, from the Dashboard, select **Add replication plan**.

3. Create a name for the replication plan.
4. Select the source and target vCenters from the Source and Target vCenter lists.
5. Select **Next**.

Select applications to replicate and assign resource groups

The next step is to group the required virtual machines into functional resource groups. Resource groups enable you to group a set of dependent virtual machines into logical groups that meet your requirements. For example, groups could contain delayed boot orders that can be run upon recovery.

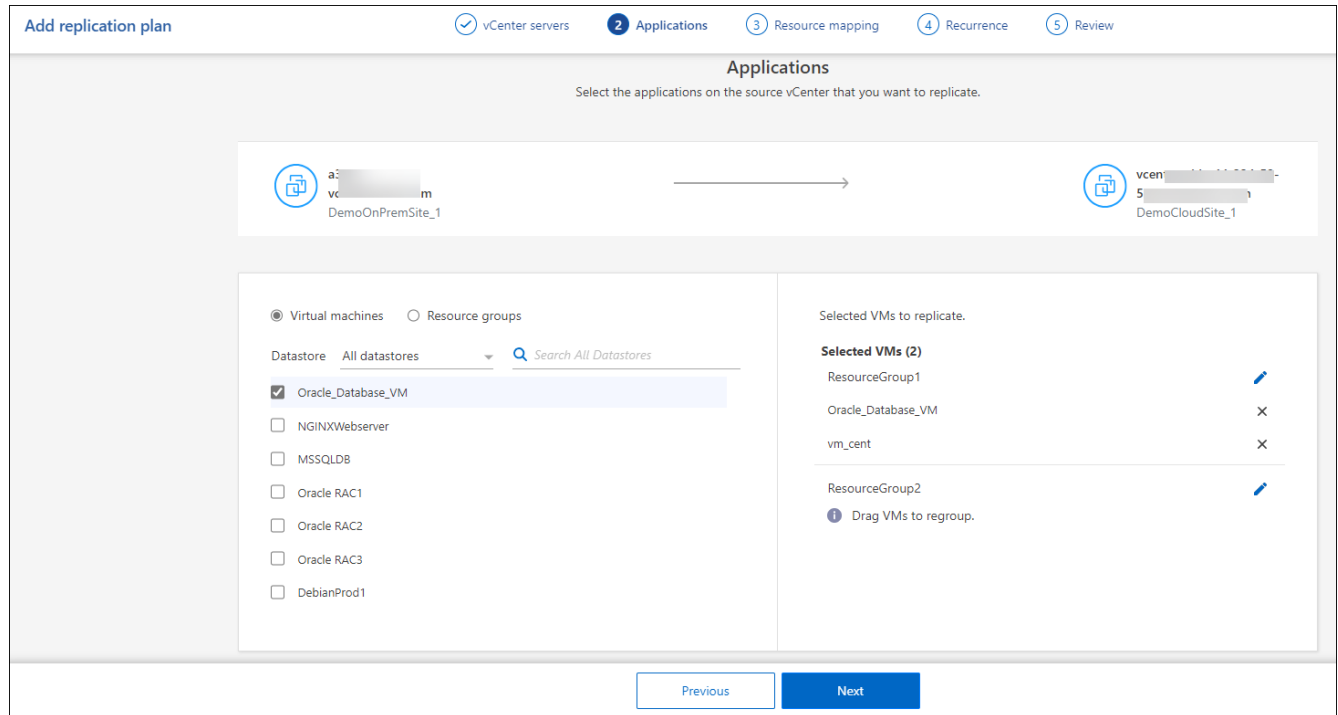



Each resource group can include one or more virtual machines. The virtual machines will power on based on the sequence in which you include them here.

Steps

1. On the left side of the Applications page, optionally search for datastores.
2. On the left side of the Applications page, select the virtual machines that you want to replicate and assign to the selected group.

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



3. Optionally, do any of the following:
 - To change groups, click the group **Edit**  icon.
 - To remove a virtual machine from a group, select **X**.
 - To move a virtual machine to a different group, drag and drop it in the new group.
4. When you have multiple resource groups, ensure that the sequence of the groups matches the operational sequence that should occur.

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

5. Optionally, rename the group by clicking the **Edit** icon.
6. Select **Next**.

Map source resources to the target

In the Resource mapping step, specify how the resources from the source environment should map to the target.

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 BlueXP 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.
3. **Compute resources:** Select the down arrow next to **Compute resources**.
 - Source cluster
 - Target cluster (optional)
4. **Virtual networks:** In the Failover mappings tab, select the down arrow next to **Virtual networks**. Enter or select the source virtual LAN and target segment.
5. **Virtual machines:** In the Failover mappings tab, select the down arrow next to **Virtual machines**.

Select the network mapping to the appropriate segment. The segments should already be provisioned, so select the appropriate segment to map the virtual machine.

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

SnapMirror is at the volume level. So, all virtual machines are replicated to the replication target. Make sure to select all virtual machines that are part of the datastore. If they are not selected, only the virtual machines that are part of the replication plan are processed.

- **IP address type:** When you are mapping networking between source and target locations in the virtual machines section of the replication plan, BlueXP disaster recovery offers two options: DHCP or static IP. For static IPs, configure the subnet, gateway, and DNS servers. Additionally, enter credentials for virtual machines.
 - **DHCP:** If you choose this option, you provide just the 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 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.

- **Source VM CPU and RAM:** Under the Virtual machines details, you can optionally resize the VM CPU and RAM parameters.
 - **Boot delay:** You can modify the boot order for all the selected virtual machines across the resource groups. By default, the boot order selected during resource-group selection is used; however, you can make changes at this stage. With this field, you can adjust the delay in minutes of the boot action.
 - **Create application-consistent replicas:** Indicate whether 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.
6. **Datastores:** In the Failover mappings tab, select the down arrow next to **Datastores**. Based on the selection of virtual machines, datastore mappings are automatically selected.

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

- **RPO:** Enter the Recovery Point Objective (RPO) to indicate the amount of data to recover (measured in time). For example, if you enter an RPO of 60 minutes, the recovery must have data that is not older than 60 minutes at all times. If there is a disaster, you are allowing the loss of up to 60 minutes of data. Also enter the number of Snapshot copies to retain for all datastores.
- **Retention count:** Enter the number of snapshots you want to retain.
- **Source and Target datastores:** If a volume has a SnapMirror relationship already established, you can select the corresponding source and target datastores. If you select a volume that does not have a SnapMirror relationship, you can create one now by selecting the working environment and its peer SVM.



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 BlueXP disaster recovery.

- **Consistency Groups:** When you create a replication plan, you can include VMs that are from different volumes and different SVMs. BlueXP disaster recovery creates a Consistency Group

Snapshot.

- If you specify the Recovery Point Objective (RPO), the service schedules a primary backup based on the RPO and updates the secondary destinations.
 - 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 on-demand backup, updates the destination, and uses that Snapshot for the failover.
7. To set different mappings for the test environment, uncheck the box and select the **Test mappings** tab. 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.


Identify the recurrence

Select whether you want to migrate data (a one-time move) to another target or replicate it at the SnapMirror frequency.

If you want to replicate it, identify how often data should be mirrored.

Steps

1. In the Recurrence page, select **Migrate** or **Replicate**.
 - **Migrate**: Select to move the application to the target location.
 - **Replicate**: Keep the target copy up to date with changes from the source copy in a recurring replication.



2. Select **Next**.

Confirm the replication plan

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



You can later disable or delete the replication plan.

Steps

1. Review information in each tab: Plan Details, Failover Mapping, 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.

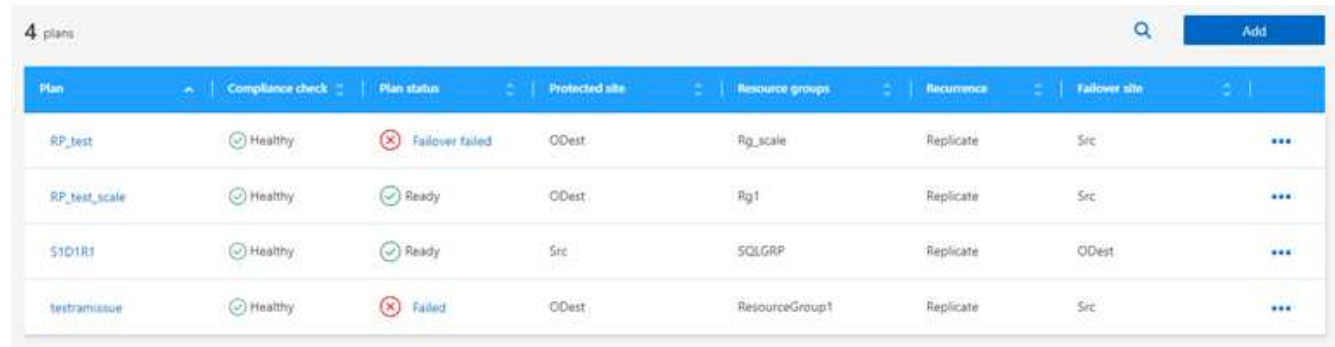
A test failover creates a FlexClone volume, mounts the datastore, and moves the workload on that datastore. A test failover operation does *not* impact production workloads, the SnapMirror relationship used

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 BlueXP disaster recovery top menu, select **Replication plans**.



Plan	Compliance check	Plan status	Protected site	Resource groups	Recurrence	Failover site	
RP_test	Healthy	Failover failed	ODest	Rg_scale	Replicate	Src	...
RP_test_scale	Healthy	Ready	ODest	Rg1	Replicate	Src	...
S1D1R1	Healthy	Ready	Src	SQLGRP	Replicate	ODest	...
tetramissue	Healthy	Failed	ODest	ResourceGroup1	Replicate	Src	...

2. Select the **Actions** **...** icon and select **Edit schedules**.
3. Enter how frequently in minutes that you want BlueXP 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_test_scale

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

Compliance check

Frequency (min) i

Test failover

Run test failovers on a monthly schedule

Day of the month i Time Start date i

▼ ▼ 📅

Required **Required** **Required**

Automatically clean up after test failover i

- To clean up the test environment after the failover test finishes, check **Automatically clean up after test failover**.



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

- Select **Save**.

Replicate applications to another site

Using BlueXP 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 BlueXP disaster recovery verifies that the replication is actually occurring according to the plan. You can monitor the progress in the Job Monitor page.

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. From the BlueXP left nav, select **Protection > Disaster recovery**.
2. From the top menu, select **Replication plans**.
3. Select the replication plan.
4. On the right, select the **Actions** option  and select **Replicate**.

Migrate applications to another site

Using BlueXP 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 BlueXP 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. From the BlueXP left nav, select **Protection > Disaster recovery**.
2. From the top menu, select **Replication plans**.
3. Select the replication plan.
4. On the right, select the **Actions** option  and select **Migrate**.

Fail over applications to a remote site

In case of a disaster, fail over your primary on-premises VMware site to another on-premises VMware site or VMware Cloud on AWS.

During the failover, the most recent SnapMirror Snapshot copy is used. Or, you can select a specific Snapshot copy from a point-in-time Snapshot copy (per the retention policy of SnapMirror). The point-in-time option can be helpful if you are facing a corruption event such as ransomware, where the most recent replicas are already compromised or encrypted. BlueXP disaster recovery shows all available points in time.

This procedure breaks the replication relationship, places the vCenter source VMs offline, and enables read/write on the target site.

You can include custom scripts in .sh, .bat, or .ps1 format as post failover processes. With custom scripts, you can have BlueXP 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.

Before you start the failover, you can test the process, ensuring success when you need it. The test does not place the virtual machines offline.

Test the failover process

Before you start an actual failover, you should test the failover process.


During a failover test, virtual machines are temporarily created. BlueXP disaster recovery does not map the target volume. Instead, it makes a new FlexClone volume from the selected Snapshot, and a temporary datastore backing the FlexClone volume is mapped to the ESXi hosts.

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.


Steps

1. From the BlueXP left nav, select **Protection > Disaster recovery**.
2. From the BlueXP disaster recovery top menu, select **Replication plans**.
3. Select the replication plan.
4. On the right, select the **Actions** option  and select **Test failover**.
5. In the Test failover page, enter "Test failover" and select **Test fail over**.
6. 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 BlueXP disaster recovery top 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:

- 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.
- The target virtual machines are registered and powered on based on the order captured in the Resource groups page.
- The SnapMirror relationship is reversed from target to source virtual machine.






After the failover starts, the recovered VMs can be seen 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 BlueXP left nav, select **Protection > Disaster recovery**.
2. From the BlueXP disaster recovery top menu, select **Replication plans**.
3. Select the replication plan.
4. 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.

 **Warning:** PowerOn State found. All virtual machine(s) in replication plan 'Rp_script' are powered-On. 

Snapshot copy for volume recovery Take snapshot now Select

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

Force failover 

Enter **Failover** to confirm

Failover

Failover

Cancel

5. 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.

6. 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.
7. Type "failover" in the box.
8. Select **Fail over**.
9. To check the progress, in the top menu, select **Job monitoring**.


Fail back applications to the original source

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.

In this workflow, BlueXP 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:

- On the target site, the virtual machines are powered off and unregistered, and volumes are unmounted.
- The SnapMirror relationship on the original source is broken to make it read/write.
- The SnapMirror relationship is resynchronized to reverse the replication.
- The source virtual machines are powered on and registered, and volumes are mounted on the source.

Steps

1. From the BlueXP left nav, select **Protection > Disaster recovery**.
2. From the BlueXP disaster recovery top menu, select **Replication plans**.
3. Select the replication plan.
4. On the right, select the **Actions** option  and select **Fail back**.
5. Enter the replication plan name to confirm and start the failback.
6. Choose the Snapshot for the datastore from which to recover. The default is the latest.
7. To check the progress, in the top menu, select **Job monitoring**.

Manage sites, resource groups, replication plans, datastores and virtual machines information


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

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

Manage vCenter sites

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

Steps


1. From the top 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 can also edit and delete resource groups.

Steps

1. From the top 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**.

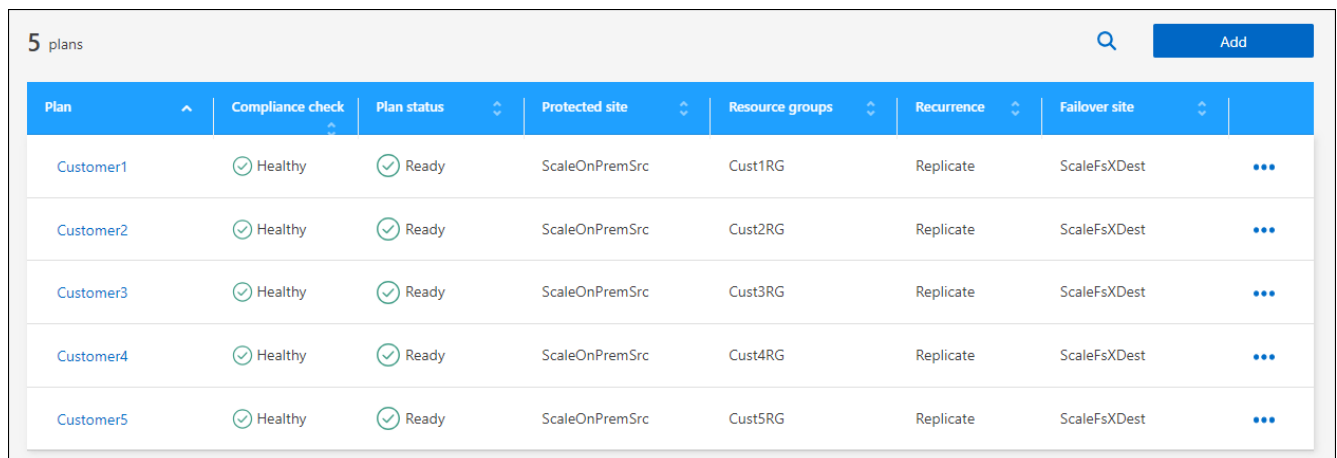
Manage replication plans

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

- 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 top menu, select **Replication plans**.



The screenshot shows a table with 5 plans. The table has columns for Plan, Compliance check, Plan status, Protected site, Resource groups, Recurrence, and Failover site. Each row represents a plan with a name (Customer1 to Customer5), a compliance check status (Healthy), a plan status (Ready), a protected site (ScaleOnPremSrc), a resource group (Cust1RG to Cust5RG), a recurrence (Replicate), and a failover site (ScaleFsXDest). Each row also has a three-dot menu icon on the right.

Plan	Compliance check	Plan status	Protected site	Resource groups	Recurrence	Failover site	
Customer1	Healthy	Ready	ScaleOnPremSrc	Cust1RG	Replicate	ScaleFsXDest	...
Customer2	Healthy	Ready	ScaleOnPremSrc	Cust2RG	Replicate	ScaleFsXDest	...
Customer3	Healthy	Ready	ScaleOnPremSrc	Cust3RG	Replicate	ScaleFsXDest	...
Customer4	Healthy	Ready	ScaleOnPremSrc	Cust4RG	Replicate	ScaleFsXDest	...
Customer5	Healthy	Ready	ScaleOnPremSrc	Cust5RG	Replicate	ScaleFsXDest	...

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 top 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**, **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**.

Change retention count for failover schedules

You can change how many datastores are retained.

1. From the top 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.

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: Retention count for all datastores:

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.

View datastores information

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

1. From the top 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.

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

Monitor disaster recovery jobs

You can monitor all disaster recovery jobs and determine their progress.

Steps

1. From the BlueXP left nav, select **Protection > Disaster recovery**.
2. From the top menu, select **Job monitoring**.
3. Explore all jobs related to operations and review their timestamps and status.
4. To view details of a particular job, select that row.
5. 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.

Steps

1. From the BlueXP left nav, select **Protection > Disaster recovery**.
2. From the top 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**.

Copyright information

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

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

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

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

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

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

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

Trademark information

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