

This document describes the following:

- [Basic Recovery Operations Provided by Zerto Virtual Replication](#)
- [Live Disaster Recovery Best Practices](#)
- [Frequently Asked Questions](#)

Basic Recovery Operations Provided by Zerto Virtual Replication

Zerto Virtual Replication provides a number of operations to recover virtual machines at the peer site, as follows:

- [The Move Operation](#)
- [The Failover Operation](#)
- [The Failover Test Operation](#)
- [The Clone Operation](#)

The Move Operation

Use the Move operation to migrate protected virtual machines from the protected (source) site to the recovery (target) site in a planned migration.

When you perform a planned migration of the virtual machines to the recovery site, Zerto Virtual Replication assumes that both sites are healthy and that you planned to relocate the virtual machines in an orderly fashion without loss of data.

The Move operation has the following basic steps:

1. Gracefully shutting down the protected virtual machines. This ensures data integrity.
If the machines cannot be gracefully shut down, for example, when VMware Tools is not available, you can manually shut down the machines before starting the Move operation or you specify as part of the operation to forcibly power off the virtual machines. If the machines cannot be gracefully shut down automatically and are not manually shut down and the Move operation is not set to forcibly power them off, the Move operation stops and Zerto Virtual Replication rolls back the virtual machines to their original status.
2. Inserting a clean checkpoint. This avoids potential data loss since the virtual machines are not on and the new checkpoint is created after all I/Os have been written to disk.
3. Transferring all the latest changes that are still in the queue to the recovery site, including the new checkpoint.
4. Creating the virtual machines in the recovery site and attaching each virtual machine to its relevant virtual disks, based on the checkpoint inserted in step 2.
Note: The virtual machines are created without CD-ROM drives, even if the protected virtual machines had CD-ROM drives.
5. Setting VMware HA to prevent DRS. This prevents automatic vMotioning of the affected virtual machines during the move operation.
6. Powering on the virtual machines making them available to the user. If applicable, the boot order defined in the VPG settings is used to power on the machines.
Note: If the virtual machines do not power on, the process continues and the virtual machines must be powered on manually.
7. Committing the Move operation. The default is to automatically commit the Move operation without testing. However, you can also run basic tests on the machines to ensure their validity to the clean checkpoint. Depending on the commit/rollback policy that you specified for the operation, the operation is committed, finalizing the move, or rolled back, aborting the operation.
8. Removing the protected virtual machines from the inventory.

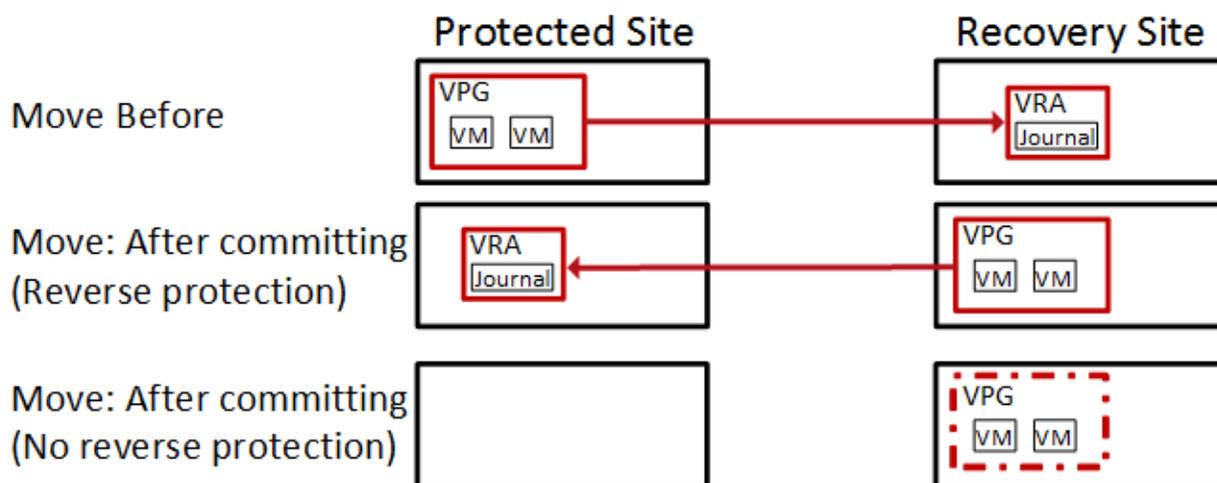
- Promoting the data from the journal to the machines. The machines can be used during the promotion and Zerto Virtual Replication ensures that the user sees the latest image, even if this image, in part, includes data from the journal. That is, when accessing the migrated virtual machine, Zerto Virtual Replication can present both data from the disks and from the journal, to ensure that information is current.

Note: Virtual machines cannot be moved to another host during promotion. If the host is rebooted during promotion, make sure that the VRA on the host is running and communicating with the Zerto Virtual Manager before starting up the recovered virtual machines.

- If reverse replication is specified, the virtual disks used by the virtual machines in the protected site are used for the reverse protection. A `Delta Sync` is performed to make sure that the two copies, the new recovery site disks and the original protected site disks, are consistent.

If reverse replication is not specified, the VPG definition is saved but the state is *Needs configuration* and the virtual disks used by the virtual machines in the protected site are deleted. Thus, in the future if reverse protection is required, the original virtual disks are not available and a full synchronization is required.

The following diagram shows the positioning of the virtual machines before and after the completion of a Move operation.



The Failover Operation

Use the Failover operation following a disaster to recover protected virtual machines to the recovery site. A failover assumes that connectivity between the sites might be down, and thus the protected virtual machines and disks are not removed, as they are in a planned Move operation.

When you set up a failover you always specify a checkpoint to which you want to recover the virtual machines. When you select a checkpoint - the latest automatically generated checkpoint, an earlier automatically generated checkpoint, a user-defined checkpoint, or a VSS checkpoint - Zerto Virtual Replication makes sure that virtual machines at the remote site are recovered to this specified point-in-time.

Note: To decide which checkpoint to use, you can perform a number of consecutive test failovers, each to a different checkpoint until the desired checkpoint for recovery is determined.

The Failover operation has the following basic steps:

- Creating the virtual machines at the remote site in the production network and attaching each virtual machine to its relevant virtual disks, configured to the checkpoint specified for the recovery. The virtual machines are created without CD-ROM drives, even if the protected virtual machines had CD-ROM drives.

Note: The original protected virtual machines are not touched since the assumption is that the original protected site is down.

- Preventing automatically moving virtual machines to other hosts: Setting HA to prevent DRS. This prevents automatic vMotioning of the affected virtual machines during the failover operation.

3. Powering on the virtual machines making them available to the user. If applicable, the boot order defined in the VPG settings to power on the machines in a specified order is used.

Note: If the virtual machines do not power on, the process continues and the virtual machines must be manually powered on.
4. The default is to automatically commit the failover operation without testing. However, you can also run basic tests on the machines to ensure their validity to the specified checkpoint. Depending on the commit/rollback policy that you specified for the operation, after testing either the operation is committed, finalizing the failover or rolled back, aborting the operation.
5. If the protected site is still available, for example, after a partial disaster, and reverse protection is possible and specified for the failover operation, the protected virtual machines are powered off and removed from the inventory. The virtual disks used by the virtual machines in the protected site are used for reverse protection. A `Delta Sync` is performed to make sure that the two copies, the new target site disks and the original site disks, are consistent.

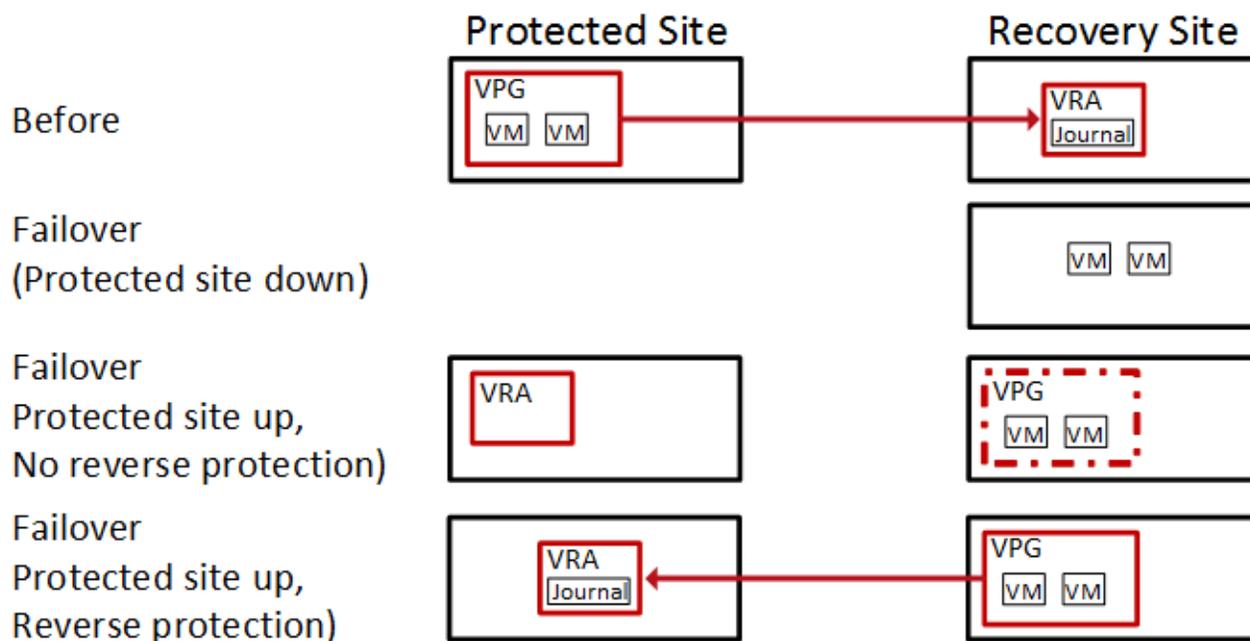
Note: If reverse protection is not possible, or reverse protection is configured not to use the original disks, the source site virtual machines are not powered off and removed. In the latter case, if possible, the virtual machines should be shut down manually before starting the failover.
6. The data from the journal is promoted to the machines. The machines can be used during the promotion and Zerto Virtual Replication ensures that the user sees the latest image, even if this includes data from the journal.

Note: The virtual machines cannot be moved to another host during promotion.

Failback after the Original Site is Operational

To perform a failback to the source site, the VPG that is now protecting the virtual machines on the target site has to be configured and then a `Delta Sync` is performed with the disks in the source site. Once the VPG is in a protecting state the virtual machines can be moved back to the source site.

The following diagram shows the positioning of the virtual machines before and after the completion of a Failover operation.



Note: The Failover operation without reverse protection does not remove the VPG definition but leaves it in a `Needs Configuration` state.

The Failover Test Operation

Use the Failover Test operation to test that during recovery the virtual machines are correctly replicated at the recovery site.

The Failover Test operation creates test virtual machines in a sandbox, using the test network specified in the VPG definition as opposed to a production network, to a specified point-in-time, using the virtual disks managed by the VRA. All testing is written to scratch volumes. The longer the test period the more scratch volumes are used, until the maximum size is reached, at which point no more testing can be done. The maximum size of all the scratch volumes is determined by the journal size hard limit and cannot be changed. The scratch volumes reside on the same datastore defined for the journal.

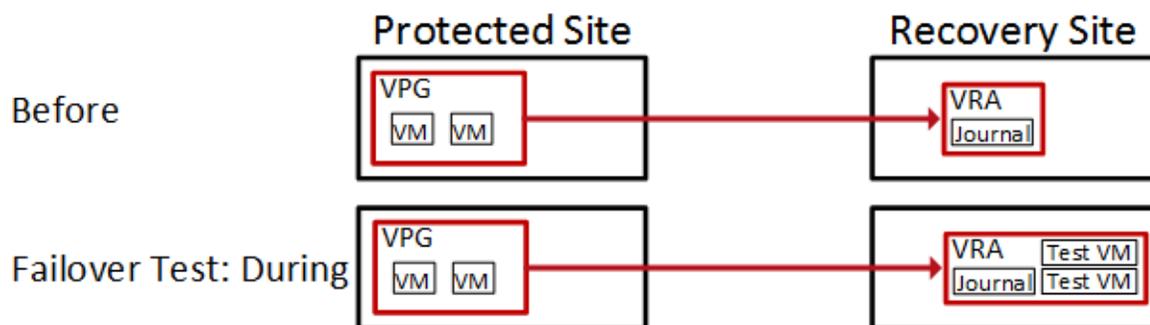
Note: During the test, any changes to the protected virtual machines at the protected site are sent to the recovery site and new checkpoints continue to be generated, since replication of the protected machines continues throughout the test. You can also add your own checkpoints during the test period.

The Failover Test operation has the following basic steps:

1. Start the test.
 - a) Choose a checkpoint to use for the test. The checkpoint can be an existing checkpoint or you can create a new one before starting the test.
 - b) Create the test virtual machines at the remote site using the network specified for testing in the VPG settings and configured to the checkpoint specified for the recovery.
 - c) Power on the virtual machines making them available to the user. If applicable, use the boot order defined in the VPG to power on the machines.
2. Stop the test.
 - a) Power off the test virtual machines and remove them from the inventory.
 - b) Add the following tag to the checkpoint specified for the test: `Tested at startDateAndTimeOfTest`.

Note: The updated checkpoint can be used to identify the point-in-time to restore the virtual machines in the VPG during a failover.

The following diagram shows the positioning of the virtual machines before and during a Failover test operation.



The Clone Operation

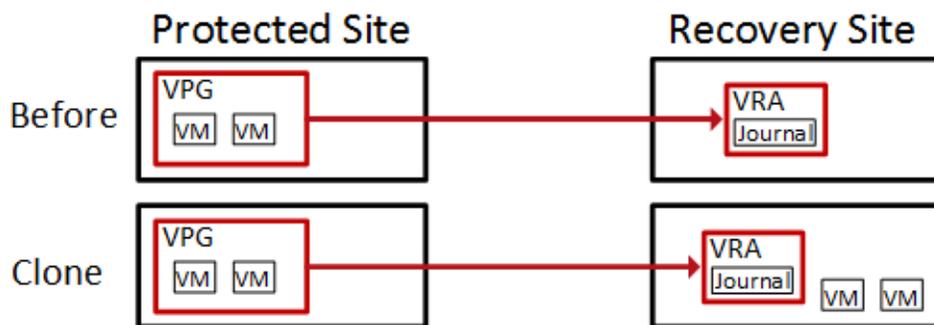
Use the Clone operation to create a copy of the VPG virtual machines on the recovery site. The virtual machines on the protected site remain protected and live.

The Clone operation has the following basic steps:

1. Create the cloned disks with the data from the journal to the specified checkpoint.
2. Create the virtual machines at the remote site in the production network and attach each virtual machine to its relevant cloned disks, configured to the checkpoint specified for the clone.

The cloned machines are named with the names of the protected machines, with the timestamp of the checkpoint used to create the clone. The cloned virtual machines are not powered on and are not protected by Zerto Virtual Replication.

The following diagram shows the positioning of the virtual machines before and after the completion of a Clone operation.



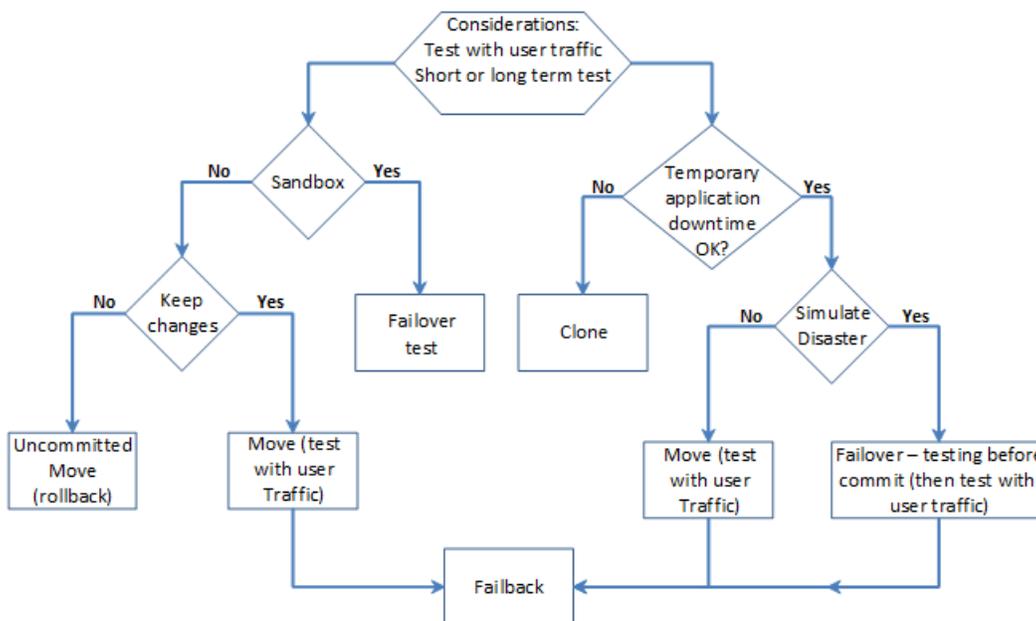
Live Disaster Recovery Best Practices

This section describes how to use the basic Zerto Virtual Replication recovery operations to perform live disaster recovery tests, in different situations.

When performing a live DR test you need to consider the following:

- The purpose of the live DR test: Do you only want to verify that the VMs can recover properly or do you want to conduct a full DR test that will include running user traffic against the recovered VMs?
- The length of time you want to test the recovery, a few hours or several days.
- Whether the changes to the recovered machine need to be retained after the test or can they be discarded?
- Whether you are willing to accept temporary downtime of the application.
- Whether you want to simulate an actual disaster at the protected site, for example by simulating a network outage or bringing down the protected site.

The following flowchart shows the testing decision flow:



During any live test, Zerto recommends that you only maintain one working version of the same virtual machine. Thus, the first step in any test, except for a Failover Test or Clone, is to make sure that the protected virtual machines are shut down before starting to test recovered machines. During a Zerto Virtual Replication Move operation the first step Zerto Virtual Replication performs is to shut down the protected machines, to ensure data integrity. However, a Zerto Virtual Replication Failover operation assumes that the protected virtual machines are no longer accessible (the total site disaster scenario) and does not

attempt to shut them down at the beginning of the operation. In a live test using a Failover operation you have to manually shut down the virtual machines to be tested at the beginning of the test in order to prevent potential *split-brain* situations where two instances of the same applications are live at the same time.

If you want to perform a live DR test that includes a simulated disaster you can simulate the disaster, for example, by disconnecting the network between the two sites. In this type of test, once the disaster is simulated a Move operation cannot be used, since it requires both sites to be healthy, while a Failover operation can be used.

Basic Verification – User Traffic Is Not Run against the Recovered VMs

Basic testing that the virtual machines can recover is done using either a Failover Test operation or an uncommitted Move operation, using the `Rollback` setting.

Using a Failover Test Operation

You use a Failover Test operation if recovering the virtual machines in a sandbox, using the test network specified in the VPG definition for network isolation, is sufficient for the test. The Failover Test operation is described in the *Zerto Virtual Manager Administration Guide for the VMware vSphere Environment*.

Recommended Procedure for a Live DR Test

1. Change the VPG `Failover Test Network` to the production network used at the recovery site.
2. Manually shut down the virtual machines in the VPG.
3. Insert a new checkpoint. This avoids potential data loss since the virtual machines are shut down and the new checkpoint is added after all I/Os have been written to disk.
4. Optionally simulate a disaster, for example by disconnecting the two sites.
5. Perform a test failover on the VPG, choosing the checkpoint you added in step 3.
6. Verify that the test machines are recovered as expected.
7. Run user traffic against the virtual machines.
8. Stop the failover test.
9. Reconnect the sites.

Failover Test Considerations

- You do not have to shut down the protected virtual machines, and changes from the test phase are not kept or applied to the protected applications.
- You can recover to a specific point-in-time.
- You can use an isolated network to enable testing in a sandbox environment and not a live DR environment. This is the recommended practice.
- During the testing period, every change is recorded in a scratch volume. Thus, since both the scratch volume and virtual machines being tested are on the same site, performance can be impacted by the increased IOs during the failover test. Also, the longer the test period the more scratch volumes are used, until the maximum size is reached, at which point no more testing can be done. The maximum size of all the scratch volumes is determined by the journal size hard limit and cannot be changed. The scratch volumes reside on the storage defined for the journal.
- At the end of the test, if you powered off the virtual machines in the protected site, you can power them back on and continue to work without the need to save or replicate back any data changed during the test.

You can also use a Failover Test operation if you want to simulate an actual disaster for around an hour or less and do not want to save any changes on the recovery site.

Using an Uncommitted Move Operation

You use a Move operation with the commit/rollback policy set to rollback after the test period, if recovering the virtual machines needs testing in the recovery site production environment.

Note: Committing the Move operation will necessitate failing the migrated virtual machines back to the production site after a Delta Sync has been performed on the committed machines in the recovery site.

Recommended Procedure for a Live DR Test

1. In the `MOVE` wizard, in the EXECUTION PARAMETERS tab, for commit policy, select None.
2. Either power off the relevant virtual machines or check the `Force Shutdown` checkbox, in the EXECUTION PARAMETERS tab, to make sure that the virtual machines are shut down, if they cannot be powered off using VMware Tools.
3. After testing the machines in the recovery site, roll back the Move operation, which will return the virtual machines to their pre-test state.

Move Considerations

- Changes from the pre-commit phase are not kept or applied to the protected applications.
- The virtual machines are allocated disks and connected to the network for a full test of the environment.
- The protected machines are turned off until the end of the test, ensuring that there are no conflicts between the protected site and recovery site.
- During the testing period, every change is recorded in a scratch volume to enable rolling back. Thus, since both the scratch volume and virtual machines being moved are on the same site, performance can be impacted by the increased IOs during the testing period. Also, the longer the test period the more scratch volumes are used, until the maximum size is reached, at which point no more testing can be done. The maximum size of all the scratch volumes is determined by the journal size hard limit and cannot be changed. The scratch volumes reside on the storage defined for the journal.
- You can only recover to the last checkpoint written to the journal, at the start of the Move operation.

Run User Traffic against the Recovered VMs

Testing actual user traffic against the recovered virtual machines can be done using a Clone, Move, or Failover operation, as follows:

Move operation – When you can shut down the source virtual machines but you don't want or need to simulate an actual disaster.

Failover operation – When you want to simulate an actual disaster.

Clone operation – When the source application has to continue throughout the test.

Using a Move Operation

You use a Move operation when you can shut down the source virtual machines but you don't want to simulate an actual disaster. After the virtual machines have been recovered in the target site they are used as the production machines for as long as the test lasts. The Move operation is described above and in the *Zerto Virtual Manager Administration Guide for the VMware vSphere Environment*.

Recommended Procedure for a Live DR Test

1. To enable using the Move functionality for a DR test, in the `MOVE` wizard, in the EXECUTION PARAMETERS tab, for commit policy, select None.
2. Move the VPG back to the original protected site. A `Delta Sync` is performed to copy the new transactions performed on the virtual machines in the recovery site back to the original protected site.

Move Considerations

- You can test the moved machines before they are committed.
- You can test for as long as you want.

- The virtual machines are allocated disks and connected to the network for a full test of the environment.
- The originally protected disks are maintained for a faster failback when reverse replication is specified.
- The protected machines are turned off until the move is committed and then they are removed from the protected site. This ensures that there are no conflicts between the protected site and recovery site.
- You must test to the **last** checkpoint, taken after the protected virtual machines are shut down.
- An actual disaster is not simulated.
- During the testing period, if reverse replication is not specified, there is no protection for the recovered machines.

Using a Failover Operation

You use a Failover operation when you can shut down the source virtual machines and you want to simulate an actual disaster. After the virtual machines have been recovered in the target site they are used as the production machines for as long as the test lasts.

Using a Failover operation to test DR requires specific steps to ensure that the virtual machines are gracefully migrated to the target site, similar to a Move operation and that, like a Move operation, they can be verified prior to committing the failover. The Failover operation is described above and in the *Zerto Virtual Manager Administration Guide for the VMware vSphere Environment*. The following procedure highlights specific steps for a DR test.

Recommended Procedure for a Live DR Test

1. Manually shut down the virtual machines.
2. Insert a new checkpoint. This avoids potential data loss since the virtual machines are shut down and the new checkpoint is added after all I/Os have been written to disk.
3. Optionally simulate a disaster, for example by disconnecting the two sites.
4. Perform a live failover on the VPG, specifying the commit policy and choosing the checkpoint you added in step 2. Choose a commit policy that will give you the necessary time to check that the failed over virtual machines have been successfully recovered to the correct point-in-time and if they are not, you are able to roll back the failover.
5. Continue to use the recovered virtual machines.
6. The VPG is in a *Needs configuration* state, because there is no access to the protected site.

After testing the recovered virtual machine you can finalize the live DR test and fail the virtual machines back to the original protected site:

1. Reconnect the sites.
2. Enable protection for the virtual machines by editing the VPG and clicking *DONE*.
3. Zerto Virtual Replication uses the original disks to preseed the volumes and expedite the synchronization between the two sites, using a `DELTA_SYNC`. The time it will take for the Delta Sync to complete is based on total size of the disks and storage performance at both sites. After the synchronization completes the VPG enters the *Meeting SLA* state.
4. Perform a Move operation to fail back the virtual machines to the original protected site.
5. In the `MOVE` wizard, in the EXECUTION PARAMETERS tab, for commit policy, set the commit policy to enable basic testing before the move is committed.

The virtual machines are recovered at the original protected site, and the VPG enters a Delta Sync phase before it enters a *Meeting SLA* state.

Failover Considerations

- The originally protected disks are maintained for a faster failback.
- Using the Failover operation for testing is non-intuitive.
- Testing by using the Failover operation requires performing manual procedures, such as shutting down the protected virtual machines.
- During the testing period, there is no protection for the recovered machines.

Using a Clone Operation

You use the Clone operation when the source application has to continue throughout the test. You can create a clone of the virtual machines in a VPG on the peer site to a specific point-in-time. The clone is a copy of the protected virtual machines on the recovery site, while the virtual machines on the protected site remain protected and live. The Clone operation is described above and in the *Zerto Virtual Manager Administration Guide for the VMware vSphere Environment*.

The cloned virtual machines are independent of Zerto Virtual Replication. At the end of the test you can remove these machines or leave them.

Clone Considerations

- You can clone to a specific point-in-time.
- There is no protection for the cloned machines.
- After use of the clone ends, no changes made to the cloned virtual machines are applied to the protected virtual machines.
- The original virtual machines on the source site are live and online throughout the test.

Frequently Asked Questions

Question: Can I use a Failover Test or Move with precommit operation for long test runs?

Answer: Yes. These operations use scratch volumes. The longer the test period the more scratch volumes are used, until the maximum size is reached, at which point no more testing can be done. The maximum size of all the scratch volumes is determined by the journal size hard limit and cannot be changed. The scratch volumes reside on the same datastore defined for the journal.

Question: Can I use a Move operation in conjunction with a simulated disaster?

Answer: No. A Move operation requires a healthy environment and assumes network connectivity between the sites.

Question: Can I use a Clone operation and keep changes to the cloned machines?

Answer: No. The Clone operation creates an extra copy of each virtual machine which is not related to the source virtual machine.

Question: How soon can I failback my virtual machine?

Answer: After reverse protection is set, a Delta Sync will take time, based on the dataset size and storage performance. Note that reverse protection is not available from AWS.

Question: Can I detach the disks from the VRA and attach them to a VM?

Answer: Disks that were not properly promoted to a desired checkpoint might not be in a consistent state and therefore using them might produce unexpected results. Also, the disks will be recovered to a point in time that goes back as far as the history window was defined.

Feedback

Zerto is constantly looking for ways to improve all aspects of its business and your feedback is very important to us. Please send any comments you have about this document, to help us to improve it, to support@zerto.com.

ABOUT ZERTO

Zerto is committed to keeping enterprise and cloud IT running 24/7 by providing scalable business continuity software solutions. Through the Zerto Cloud Continuity Platform, organizations seamlessly move and protect virtualized workloads between public, private and hybrid clouds. The company's flagship product, Zerto Virtual Replication, is the standard for protection of applications in cloud and virtualized datacenters.