NET3 TECHNOLOGY

Data Protection

DR Runbook Template:

Contents

Roles

Net3 Technology, Inc.

- Net3 is to provide a team of two dedicated engineers to assist in recovery effort.
- Net3 is to assist in recovery efforts until VM's are bootable and networking is tested and available.
- Net3 will act as Liaison between the customer and Zerto for advanced technical support.
- Provider will provide application recovery assistance upon request at a rate of $150.00/hr.

Provide full and accurate documentation of the environment. This includes documentation of the following items:

  - Network
  - Application Support Contact Information
  - Application Dependencies
  - Current Infrastructure
  - Usernames and Passwords for network and infrastructure.
  - Licensing
  - Domain Registrar and DNS Record Management information

- Declare the Disaster
- Define scope and level of disaster recovery.
- Provide networking and infrastructure resources for recovery
- Provide application and customer infrastructure support.
- Fully test and approve each phase of restoration according to the approved test plan.

Declaration of Disaster Recovery

Although disaster recovery is self-service through the customer portal at https://portal.palmettovdc.com situations may arise when you need to declare a disaster recovery event and involve Net3 personnel.

To declare a disaster, please contact Net3 Technology at 1-888-499-0862 and specify the following information:

Name:
Customer:
Contact Information:
Disaster Level According to the predefined levels:

Customer Defined Disaster Levels

Disaster Level	Description

The online dispatcher will immediately notify the Net3 Technology disaster recovery team and a technician will establish the conference bridge and notify the contact list below of that information.

Contact List

MSP Disaster Recovery Escalation List

Role	Name	Phone	E-Mail

Customer Disaster Recovery Team

Role	Name	Phone	E-Mail

Escalation Process

MSP will only open Disaster Declarations for the following approved contacts:

Role	Name	Phone	E-Mail

Network Preparedness

The following infrastructure is currently in place and tested in readiness for a disaster event:

Cloud Networking (Provided By MSP Engineering)

PvDC Organizations		
Org Name	Type	Location

PvDC Networking		
Type	Network Name	Subnet
Additional Networks		

Edge Gateway Information					
Edge Name	External IP(s)	DHCP	Load Balancer	SSL VPN	IPSEC VPN

VPN Information				
Site Name	Endpoint IP	Subnet(s)	Peer Endpoint IP	Peer Subnet(s)

On-Premise Networking (Provided By Customer)

We suggest that the following information be documented in the runbook so that personnel providing assistance have it readily available:

- Application Support Contact Information (Vendor, Phone Numbers, Contract Numbers)
- Usernames and Passwords for network and infrastructure.
- Licensing Information (Keys, Logins)
- Domain Registrar and DNS Record Management information. (Logins, Important DNS records that need to be changed.)

On-Premise Networking		
Type	Network Name	Subnet
Additional Networks		
Type	Network Name	Subnet

ACL List

Dynamics Server – Port 443 – dynamics.thestrutinc.com

Web Server – Port 8443 – web.thestrutinc.com

Pre-Failover

Network connectivity must be established and tested before attempting to failover servers.
Once all networking is established, devices in each subnet should be able to establish communication with each other.

Disaster Recovery Scale

Business Name:

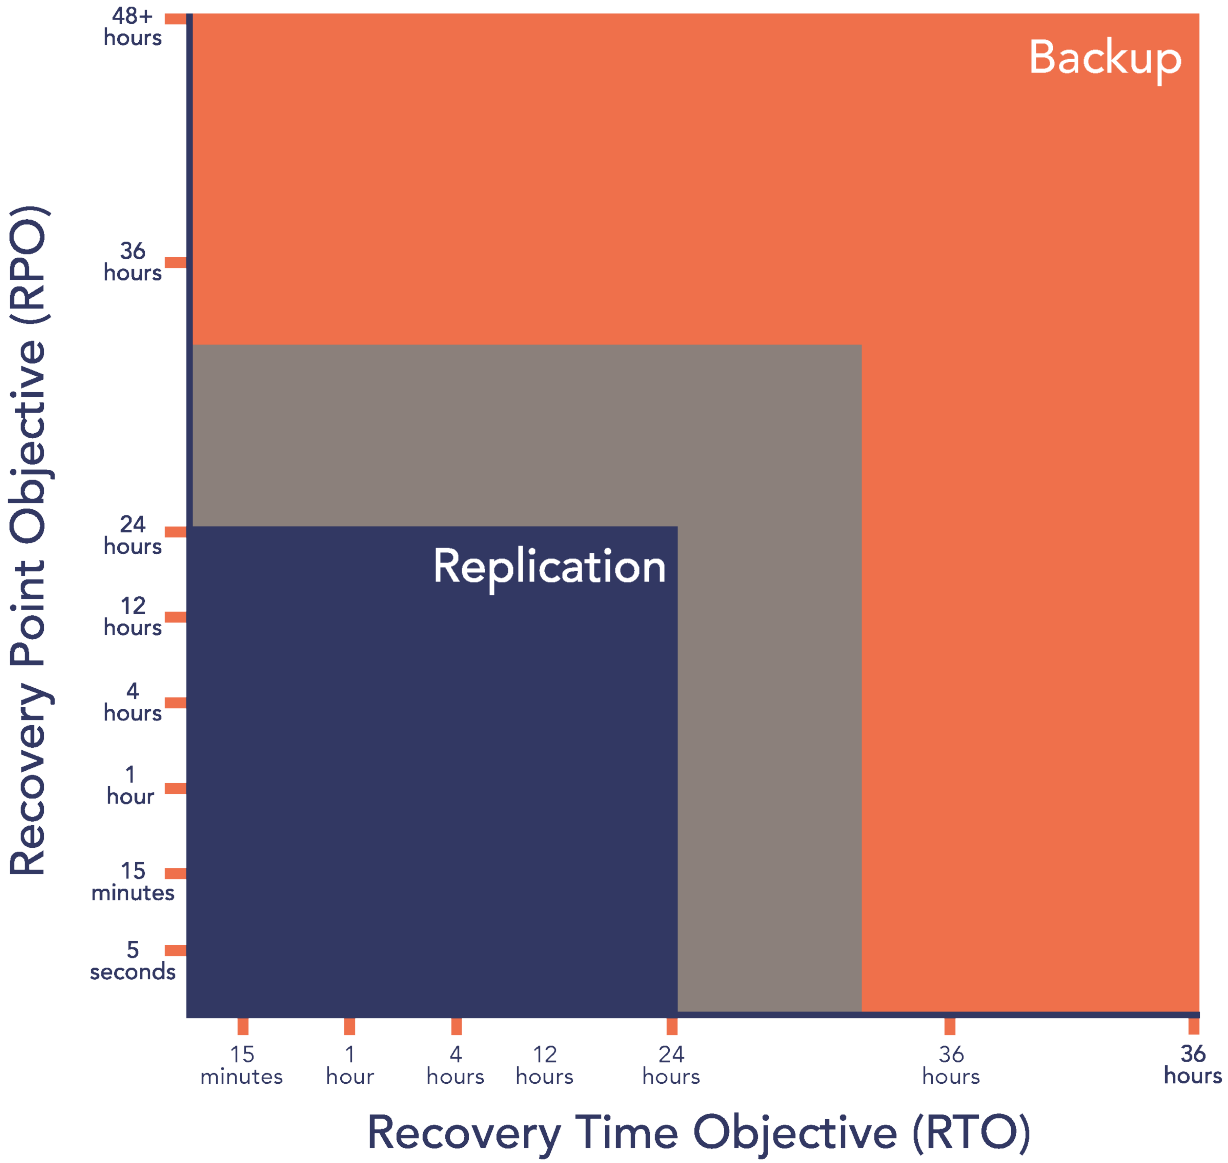

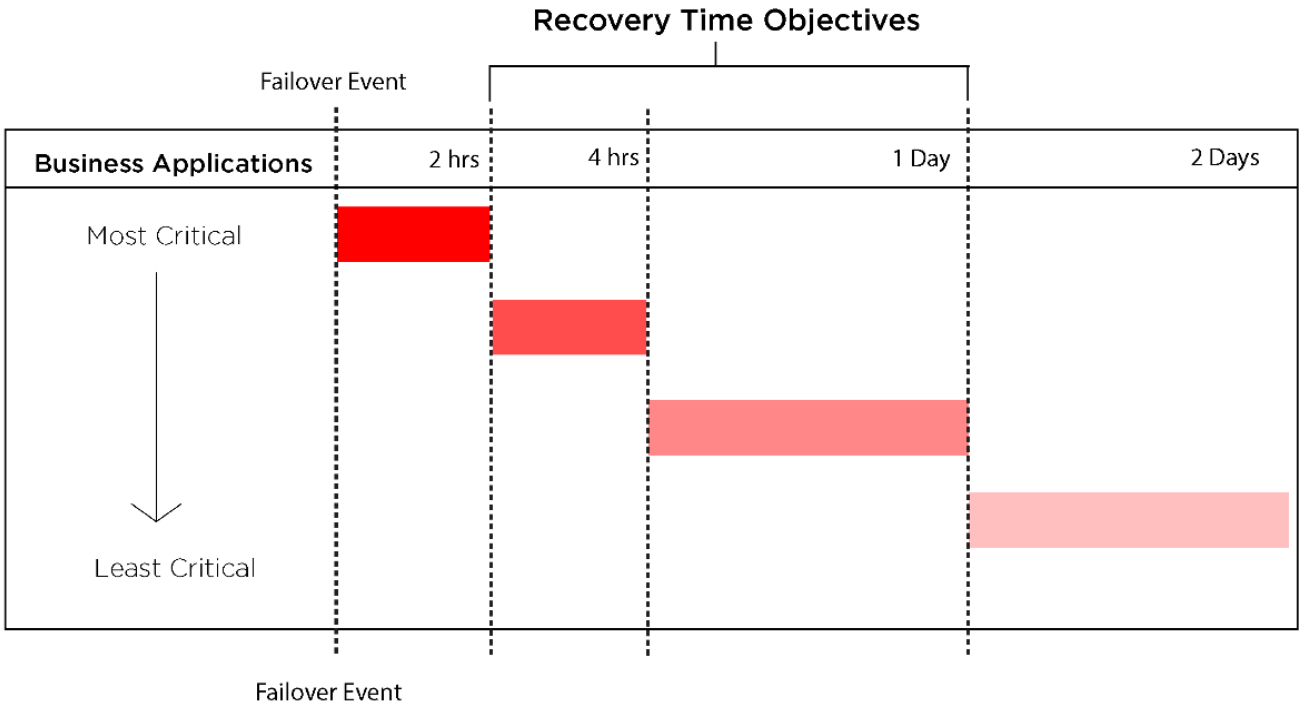
Recovery Time Objectives
Failover Event
Business Applications
2 hrs
4 hrs
1 Day
2 Days
Most Critical
Least Critical
Failover Event

Failover Environment

VPG Configuration

Status	VPG Name	#VMs	Protected Site	Recovery Site	ZORG

VM Configuration

Status	VM Name	VPG Name	Protected Site	Recovery site	Protection Status	Actual RPO

Failover / DR Order of Completion

Recovery Phases

Customer Servers are grouped into recovery phases.

Recovery Phase	Description	Recovery Method	VPG
1			
2			
3			
4			
5			

Each recovery phase includes the dependencies for successful recovery of the next phase.

Each phase also has a documented test plan to be finished before moving to the next phase.

Recovery Phase 1 Test Plan - Dynamics

Test Step				Test Step

Recovery Phase 2 Test Plan - Webserver

Test Step				Test Step

Initiating a Failover

You can initiate a failover, whereby the virtual machines in the virtual protection group are replicated to a set checkpoint in the recovery site. As part of the process you can also set up reverse replication, whereby you create a virtual protection group on the recovery machine for the virtual machines being replicated, pointing back to the protected site.

You can initiate a failover to the last checkpoint recorded in the journal, even if the protected site is no longer up.

If you have time to initiate the failover from the protected site you can. However, if the protected site is down, you initiate the failover from the recovery site.

Note: Any VPGs that are in the process of being synchronized, cannot be recovered, unless the synchronization is a bitmap synchronization.

To initiate a failover:

1. In the Zerto User Interface set the operation to *LIVE* and click *FAILOVER*.

   The *Failover* wizard is displayed.

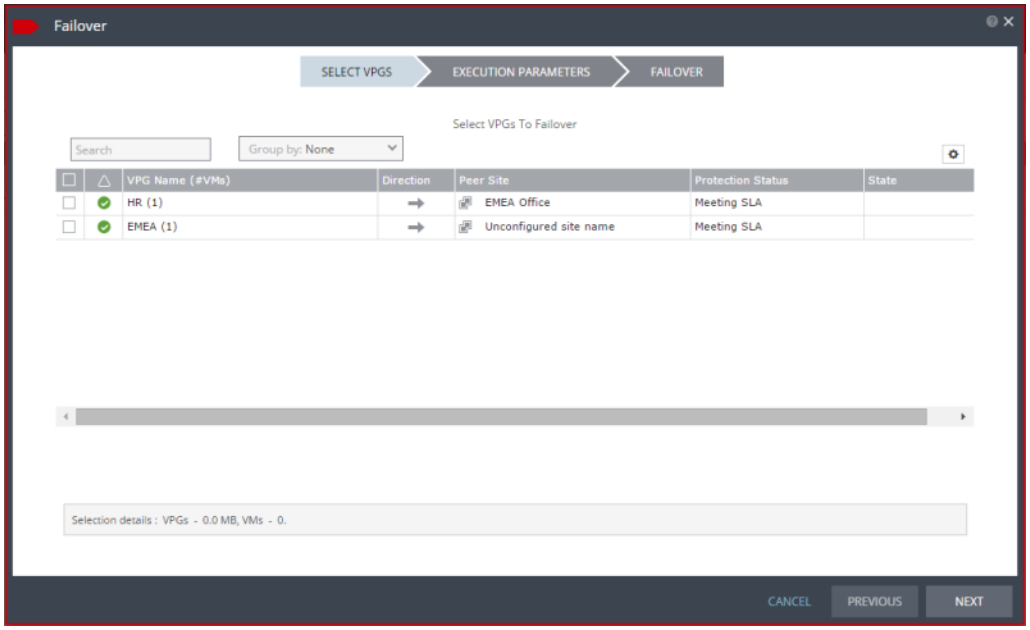

2. Select the VPGs to failover. By default, all VPGs are listed.

   At the bottom, the selection details show the amount of data and the total number of virtual machines selected.

   The Direction arrow shows the direction of the process: from the protected site to the peer, recovery, site.

3. Click *NEXT*.

The *EXECUTION PARAMETERS* step is displayed.

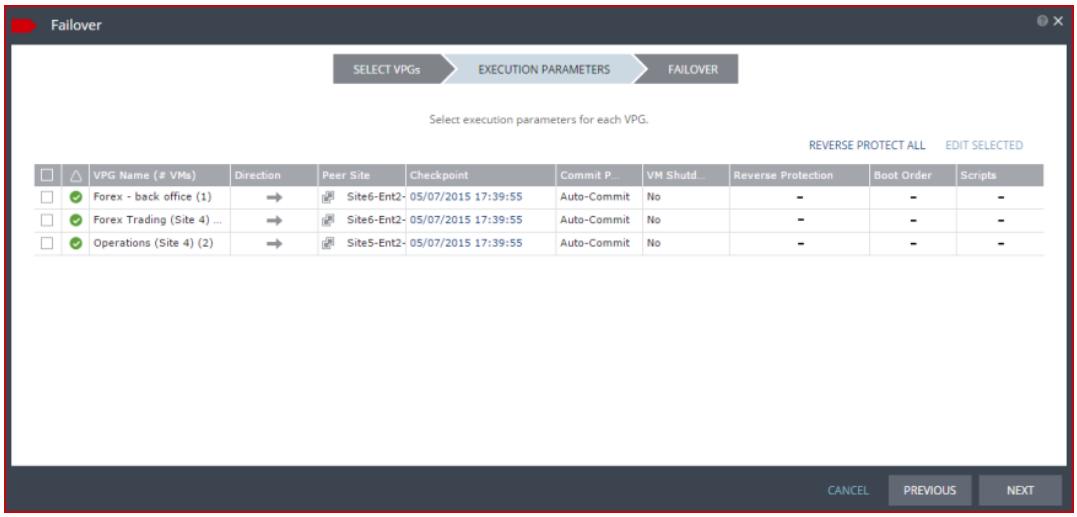

4. By default, the last checkpoint added to the journal is displayed. If you want to use this checkpoint, go to step

4. . If you want to change the checkpoint, click on the checkpoint that is displayed.

   The {*VPG-Name}: Checkpoints* dialog is displayed.

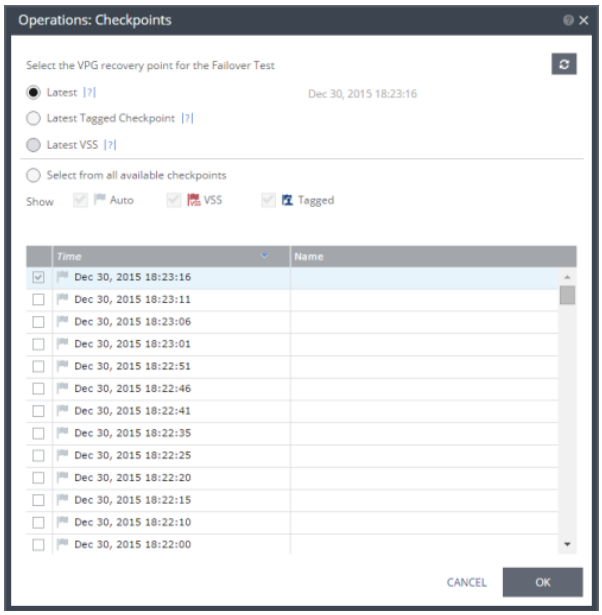

5. Select the checkpoint to use. Click the refresh button to refresh the list. You can choose from one of the following checkpoints:

   Latest – The recovery or clone is to the latest checkpoint. This ensures that the data is crash-consistent for the recovery or clone. When selecting the latest checkpoint, the checkpoint used is the latest at this point. If a checkpoint is added between this point and starting the failover or clone, the later checkpoint is not used.

   Latest Tagged Checkpoint – The recovery operation is to the latest checkpoint created manually. Checkpoints added to the virtual machine journals in the VPG by the Zerto Virtual Manager ensure that the data is crash-consistent to this point. If a checkpoint is added between this point and starting the operation, this later checkpoint is not used.

Latest VSS – When VSS is used, recovery or clone is to the latest VSS snapshot, ensuring that the data is both crash-consistent and application consistent to this point. The frequency of VSS snapshots determines how much data can be recovered.

If you do not want to use the latest checkpoint, latest tagged checkpoint, or latest VSS checkpoint, choose Select from all available checkpoints. By default, this option displays all checkpoints in the system. You can choose to display only automatic, VSS, or tagged checkpoints, or any combination of these types.

6. Click *SAVE*.
7. To change the commit policy, double-click it.
   a) To commit or roll back the Failover operation without user interaction, select Auto-Commit or Auto-Rollback. The commit or rollback will happen automatically after the specified time, if there is no user interaction.
   If you do not want an automatic commit or rollback, select None.
   a) To commit or roll back the recovery operation automatically, without any checking, select Auto-Commit or Auto-Rollback and 0 minutes.
   b) If you do not want an automatic commit or rollback, select None. You must manually commit or roll back.

To allow checking before committing or rolling back, specify an amount of time to check the recovered machines, in minutes, before the automatic commit or rollback action is performed. During this time period, check that the new virtual machines are OK and then commit the operation or roll it back. The maximum amount of time you can delay the commit or rollback operation is 1440 minutes, which is 24 hours.

Checking that involves I/O is done on scratch volumes. The longer this period the more scratch volumes are used, until the maximum size is reached, at which point no more checking 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.

When deciding to commit the operation, you can decide to configure reverse protection, regardless of the reverse protection setting when the operation started.

To specify the shutdown policy, double-click the VM Shutdown field and select the shutdown policy:

No (default) – The protected virtual machines are not touched before starting the failover. This assumes that you do not know the state of the protected machines, or you know that they are not serviceable.

Yes – If the protected virtual machines have VMware Tools available, the virtual machines are gracefully shut down, otherwise the Failover operation fails. This is similar to performing a Move operation to a specified checkpoint.

Force – The protected virtual machines are forcibly shut down before starting the failover. This is similar to performing a Move operation to a specified checkpoint. If the protected virtual machines have VMware Tools available, the procedure waits five minutes for the virtual machines to be gracefully shut down before forcibly powering them off.

8. To specify reverse protection, whereby the virtual machines in the VPG are moved to the recovery site and then protected in the recovery site, back to the original site, double-click the Reverse Protection field and configure the VPG for the reverse protection by clicking the *REVERSE* link.

   The *Edit Reverse VPG* wizard is displayed.

   You can edit the reverse protection configuration with the following differences:
   - You cannot add or remove virtual machines to the reverse protection VPG.
   - By default, reverse replication is to the original protected disks. You can specify a different storage to be used for the reverse replication.
   - If VMware Tools is available, for each virtual machine in the VPG, the IP address of the originally protected virtual machine is used. Thus, during failback the original IP address of the virtual machine on the site where the machine was originally protected is reused. However, if the machine does not contain the utility, DHCP is used.
     The vSphere version must be 4.1 or higher for re-IP to be enabled.

   Note: When committing the failover, you can reconfigure reverse protection, regardless of the reverse protection settings specified here.

9. If you want the machines in the recovery site to be booted in the order you defined when you created the VPG, click the Boot Order field and check the field.
10. If you want the procedure to run the scripts you defined when you created the VPG, click the Scripts field and check the field.
11. Click *NEXT*.
12. Click *START FAILOVER* to start the failover.

    If a commit policy was set with a timeout greater than zero, you can check the failed over virtual machines on the recovery site before committing the failover operation.

    The failover starts, by creating the virtual machines in the recovery site to the point-in-time specified: either the last data transferred from the protected site or to one of the checkpoints written in the journal.

    Note: If a virtual machine exists on the recovery site with the same name as a virtual machine being failed over, the machine is created and named in the peer site with a number added as a suffix to the name, starting with the number 1.

    If the original protected site is still up and reverse replication configured to use the protected virtual machines virtual disks, these virtual machines are powered off.

    The status icon changes to orange and an alert is issued, to warn you that the procedure is waiting for either a commit or rollback.

All testing done during this period, before committing or rolling back the failover operation, is written to thin-provisioned scratch virtual disks. These virtual disks are automatically defined when the machines are created on the recovery site for testing. 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. Using these scratch volumes makes committing or rolling back the failover operation more efficient.

Note: You cannot take a snapshot of a virtual machine before the failover operation is committed and the data from the journal promoted to the moved virtual machine disks, since the virtual machine volumes are still managed by the VRA and not directly by the virtual machine. Using a snapshot of a recovered machine before the failover operation has completed will result in a corrupted virtual machine being created.

13. After checking the virtual machines on the recovery site, choose one of the following:
    - Wait for the specified Commit Policy time to elapse, and the specified operation, either Commit or Rollback, is performed automatically.
    - Click the *Commit* or *Rollback* icon ( ) in the specific VPG tab.
      Click *Commit*. The *Commit* dialog is displayed to confirm the commit and, if necessary set, or reset, the reverse protection configuration. If the protected site is still up and you can set up reverse protection, you can reconfigure reverse protection by checking the Reverse Protection checkbox and then click the *Reverse* link. Configuring reverse protection here overwrites any of settings defined when initially configuring the move. Click *Rollback* to roll back the operation, removing the virtual machines that were created on the recovery site and rebooting the machines on the protected site. The *Rollback* dialog is displayed to confirm the rollback.

If the original protected site is still up and reverse replication is configured to use the virtual disks of the protected virtual machines, these virtual machines are removed from this site, unless the original protected site does not have enough storage available to fail back the failed over virtual machines. Finally, data is promoted from the journal to the recovered virtual machines.

During promotion of data, you cannot move a host on the recovered virtual machines. 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.

By default the virtual machines are started with the same IPs as the protected machines in the protected site. If you do not specify reverse protection, the original machines still exist in the protected site and this can create clashes. In this case, Zerto recommends ensuring that a different IP is assigned to the virtual machines when they start, when configuring each virtual machine NIC properties in the VPG, during the definition of the VPG. If you have defined the new virtual machines so that they will be assigned different IPs, the re-IP cannot be performed until the new machine is started. Zerto Virtual Replication changes the machine IPs and then reboots these machines with their new IPs.

Note: If the virtual machines do not power on, the process continues and the virtual machines must be manually powered on. The virtual machines cannot be powered on automatically in a

number of situations, such as when there is not enough resources in the resource pool or the required MAC address is part of a reserved range or there is a MAC address conflict or IP conflict, for example, if a clone was previously created with the MAC or IP address.

Reverse Protection For a Failed Over VPG

When you specify reverse protection, the virtual machines are recovered on the recovery site and then protected using the values specified during the failover. The original virtual machines are removed from the original protected site and then on the target site the data is promoted from the journal to the recovered virtual machines and then synchronization with the original site is performed so that the VPG is fully protected. The synchronization used is either a Delta Sync or if there is only one volume to synchronize, a Volume Delta Sync is performed.

For the Failover operation to complete successfully, when reverse protection is specified, the original protected site must have enough storage available to fail back the failed over virtual machines.

Note: When recovering the VPG to a vCloud Director site, reverse replication is configured to a vCD vApp.

If you do not specify reverse protection, the VPG definition is kept with the status *Needs Configuration* and the reverse settings in the VPG definition are not set.

Initiating a Failover During a Test

Replication continues during a test. If you need to initiate a failover during a test, you initiate the failover. The test stops to enable the failover and then a normal failover is performed. Any changes made to test the failover are not replicated, as only changes to the protected machines in the VPG are replicated.

Note: You cannot initiate a failover while a test is being initialized or closed.

Failback

Once customer infrastructure is back online, failback follows the same procedures as fail over.

If using replication, it will be reversed and can be failed back at any time. If using a backup product, backups must be completed and then restored on-premise from cloud.

If wanted, Zerto can be reconfigured to failback servers that were restored using a backup product.