DRaaS with Zerto Reference Guide

Revision 1.0

Table of Contents

Purpose

The purpose of this document is to catalog a group of processes having to do with GreenCloud's DRaaS with Zerto service. The Zerto software replicates server images from a virtual host to GreenCloud's infrastructure, and enables Disaster Recovery of the replicated servers. This document will include processes that will be helpful for partners managing this service.

What is DRaaS with Zerto?

DRaaS with Zerto uses Zerto's replication software to replicate server images across a VPN. The Zerto software tracks block-level changes on virtual servers in real time, and constantly streams these incremental changes to GreenCloud's Cisco Powered IaaS platform. The partner can, at any time, use the images on GreenCloud's infrastructure to create a replica VM, identical to the source VM to within 15 or fewer seconds. Zerto will automatically shut down the source VM to prevent version conflicts. Once the replica VM is established, the partner can use the IaaS platform, along with any DNS record changes, to immediately replace the source VM's functionality in a very short amount of time – as little as 5 minutes after a disaster.

Zerto Infrastructure

Zerto is managed through the Zerto Virtual Manager (ZVM), a web application and service with access to the partner's vCenter server. The ZVM enables the partner to deploy Virtual Replication Appliances (VRAs) to their ESX host. The VRAs monitor the changes to the Virtual Machines on the host. A connection is brokered across a VPN using a Zerto Cloud Connector (ZCC), and disk changes are transferred to the replication site. This process is monitored and controlled through the ZVM. The Zerto Self-Service Portal (ZSSP) can also be used to monitor replication.

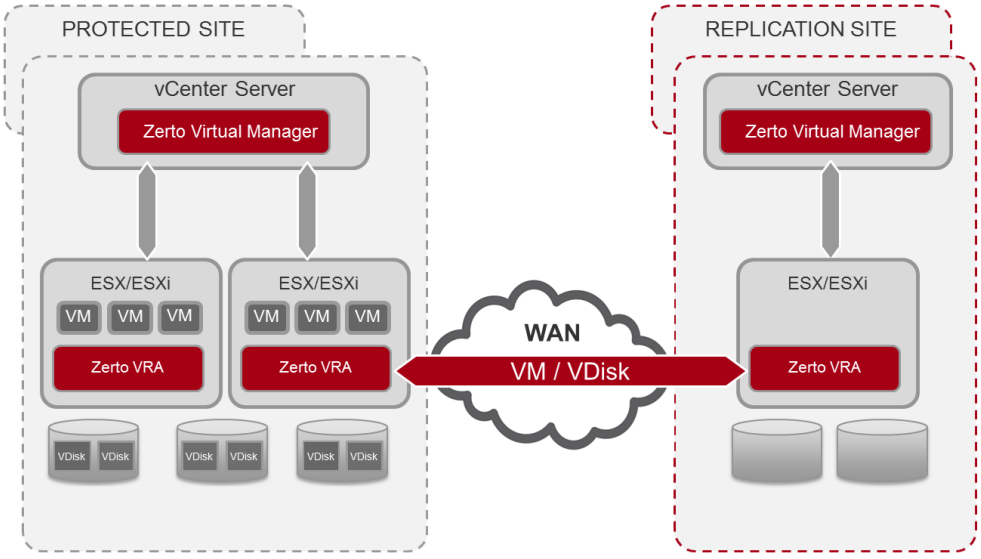

Figure 1 - Zerto Infrastructure Diagram

Common Zerto Procedures

Common procedures for Zerto include preparation/setup, failover tests, and performing full VM migrations or disaster recovery.

Zerto Setup Requirements

Requirements List

The partner should verify that the following are all present and/or true for the target environment:

- vCenter Environment 4.x or greater, or Hyper-V SCVMM, on Windows 2012 R2
- Minimum bandwidth of 10 Mbps upload for replication
- A VPN capable device
- Capacity on each host for a VRA
    - VRAs require 1 CPU, 3GB RAM, 12.5GB Storage, and 1 IP Address
- Windows Server 2008 or higher for the ZVM
    - Also requires 2 CPUs, 4GB RAM, 2GB Disk space, and Microsoft .NET Framework 4 or 4.5
- A service account with full administrator rights within the vCenter environment

Open Port Requirements

The following ports should be open between the ZVM and the GreenCloud side of the VPN.

Port	Description
22	During Virtual Replication Appliance (VRA) installation on ESXi 4.x and 5.x hosts for communication between the Zerto Virtual Manager (ZVM) and the ESXi hosts IPs and for ongoing communication between the ZVM and a Zerto Cloud Connector.
443	When installing a VRA on an ESX/ESXi host, for communication between the ZVM and the ESX/ESXi host IP and communication between the ZVM and vCenter Server and vCloud Director.
4005	Log collection between the ZVM and VRAs on the same site.
4006	TCP communication between the ZVM and VRAs on the same site.
4007	TCP communication between VRAs on different sites.
4008	TCP communication between VRAs on different sites.
4009	TCP communication between the ZVM and VRAs on the same site.
9669	HTTPS communication between the machine running the vSphere client console and a ZVM.
9080	HTTP communication between the machine running the vSphere Client console and a ZVM.
9081	TCP communication between ZVMs and between ZVMs and Zerto Cloud Connectors. This port must not be changed.
9082 and up	Two ports for each VRA handled by a cloud connector, starting at 9082. If an organization requires a total of twenty three VRAs and the cloud provider uses seven VRAs, ports 9082-9127 must be open in the firewall for the cloud connector: Ports 9082-9127 for the organization VRAs and 9082-9095 for the cloud provider VRAs. It is recommended to open a range of ports from port 9082 for VRAs.

Zerto Setup Preparation

GreenCloud Required Information

Please send the following to GreenCloud prior to setup:

- VPN device external IP
- ESX version
- /29 non-routable subnet (See Networking Setup)

Design Considerations

Please ensure the following before commencing setup:

- vCenter is set up and managing the target VMs.
- The target VMs are Hardware version 11 or under.
- The client's upload bandwidth is suitable to stay up-to-date on backups.
- Applications across multiple servers are noted, and will be grouped into Virtual Protected Groups (VPGs).
- Application importance is evaluated, and priority levels can be applied to VPGs.
- Internal network structure is mirrored in the vCloud Org before replication (see below).
- DNS records can be updated quickly in the event of a disaster.
- Site-to-site VPNs at the partner site can be cut over to the DR site.

Networking Setup

A Non-routable subnet of the size /29 which is not in use at the client site (e.g. 172.16.27.0/29) will be required for replication. Please provide a suitable subnet to GreenCloud when available. This subnet will be provisioned in the vCloud Org.

GreenCloud will set up a vCloud Org with an Edge Gateway or ASAv, along with replication networks and resources for the server images. The partner will receive:

- The client's Org ID, as well as vCloud login information,
- The External IP assigned to the vCloud Org,
- An IP Address in the non-routable subnet above for a ZCC with which the ZVM can pair,
- A link to download the ZVM installer, and
- The vCloud access URL.

A VPN will also need to be set up between the partner host and the vCloud org. The VPN should connect the vCloud Org's /29 network with the client side network containing the ZVM and VRAs. Ensure that the appropriate Firewall ports are open between the two networks, and that the Firewall and NAT rules are configured in the vCloud Org. Please see the Reference Guide for the network device in the vCloud Org (Edge Gateway or ASAv) for more information on VPN setup.

Zerto Installation

1. Install ZVM

The ZVM will need to be installed on the client site on a network with connectivity to the vCenter environment, the network to be used by the VRAs, and the VPN tunnel. Log in to the destination server with a service account with administrator rights to the vCenter environment, then download and run the ZVM Installer.

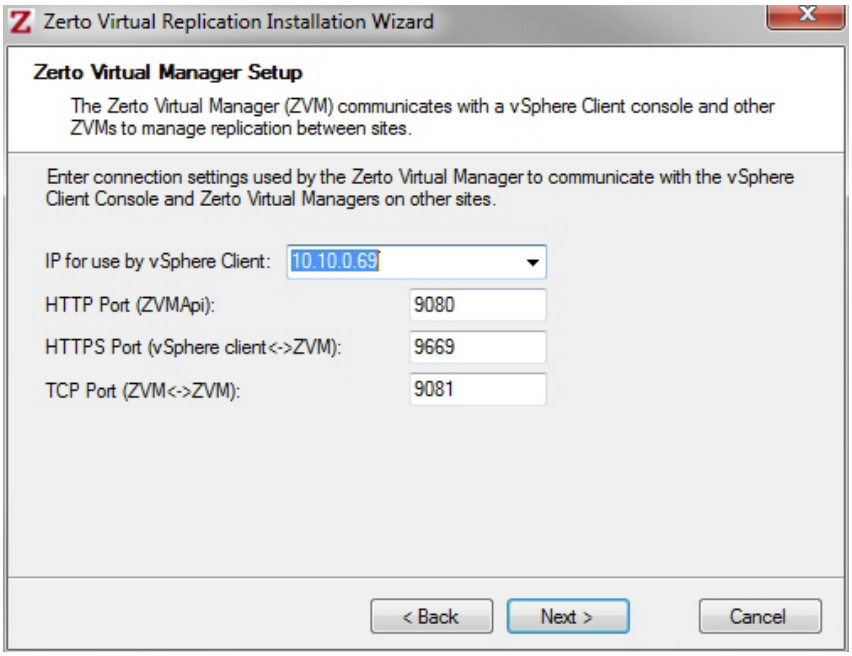

Figure 2 - ZVM Setup 1

Leave the port settings as default, use a Local System account, and choose the appropriate database option. Configure the Site Name as follows: <Org ID>_<Site Name> (e.g. 10001_JohnDoeLawyer).

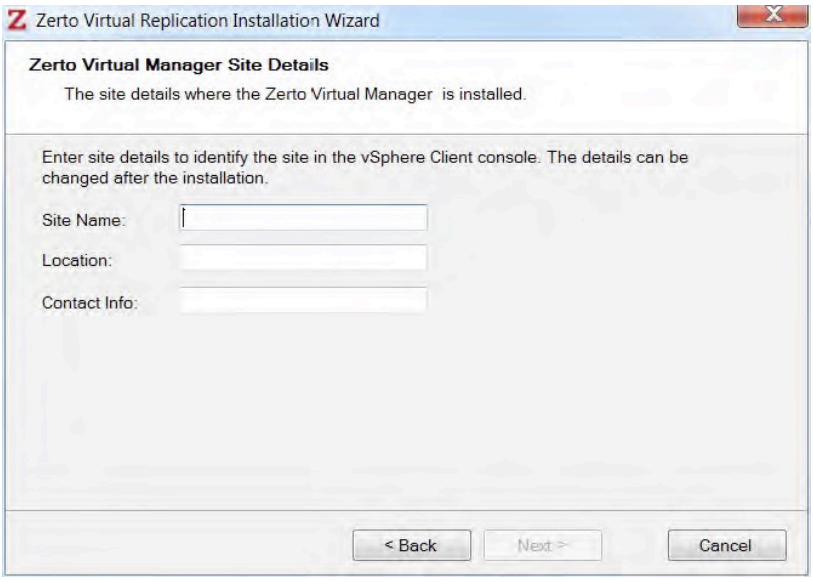

Figure 3 - ZVM Setup 2

2. Pair to GreenCloud's Site

Once the ZVM is installed, navigate to the ZVM address from a remote machine. The ZVM cannot be reached from the host machine. The address should be "https://<host.ip>:9669/zvm/". If the ZVM interface does not appear as shown below, ensure the Zerto service is running and there is network connectivity from your internet browser to the ZVM host.

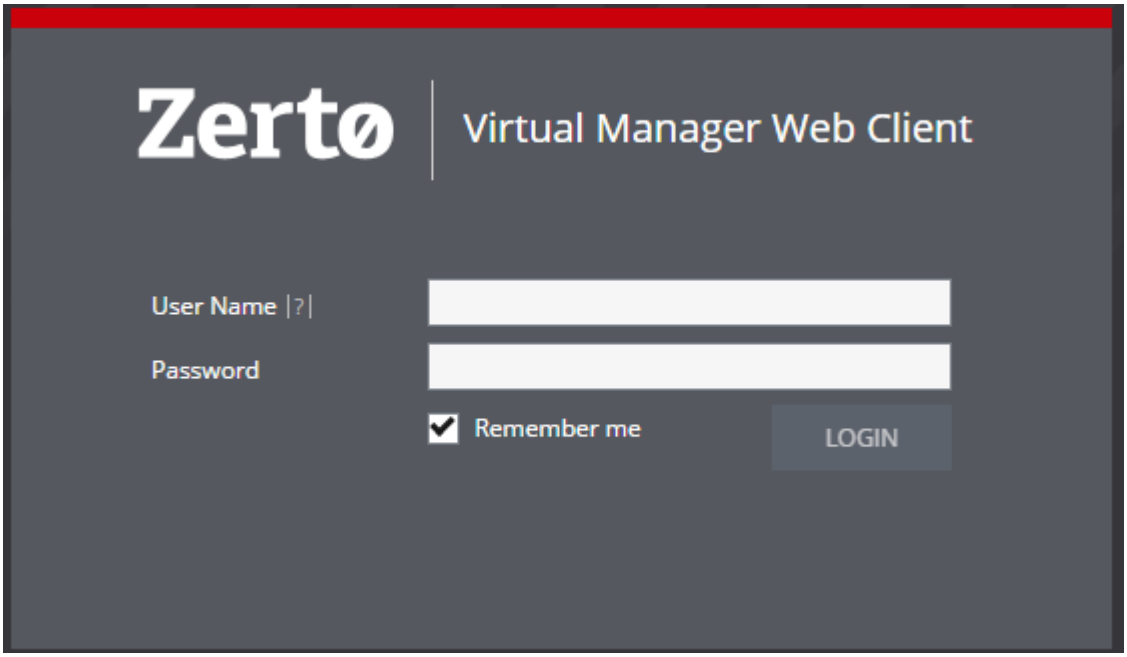

Figure 4 - ZVM Login

Log in with the service account used to install the ZVM. At this point the ZVM will prompt for a site with which to pair. Use the ZCC IP provided by GreenCloud, and select port 9081.

3. Zerto Site Settings

In the main Zerto interface, select the menu icon in the upper right-hand corner and select Site Settings.

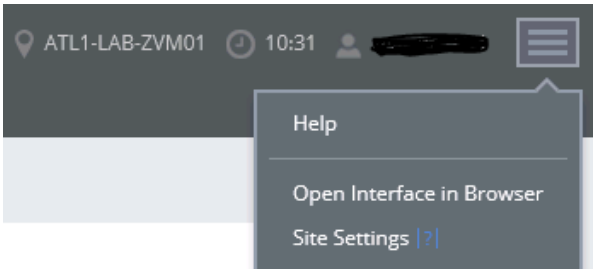

Figure 5 - ZVM Site Settings

Set Maximum bandwidth for replication if necessary during business hours. Otherwise Zerto will use whatever bandwidth is available to transfer data across the VPN.

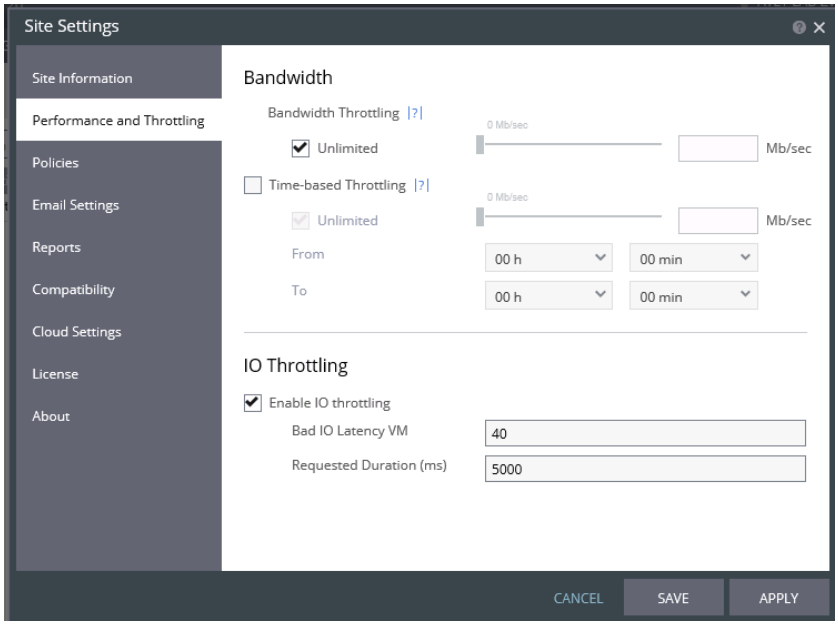

Figure 6 - Bandwidth Throttling

Select Email Settings to set up alert notifications. Check the other Site Settings fields to ensure that all policies appear correct, then select Apply.

4. Deploy VRAs

In the main Zerto interface, select the Setup tab to begin deploying VRAs to all hosts.

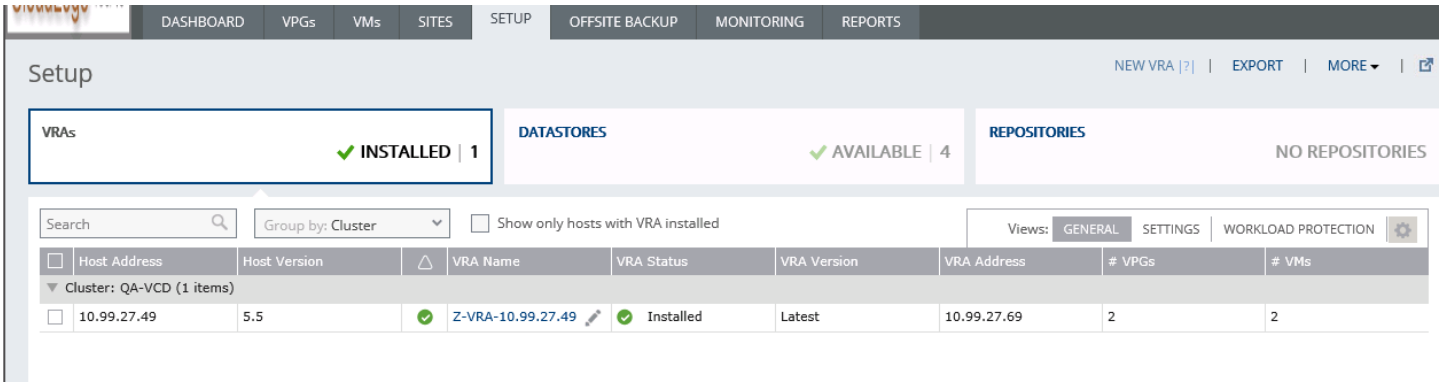

Figure 7 - VRA Setup 1

Select New VRA in the upper right to show the dialog.

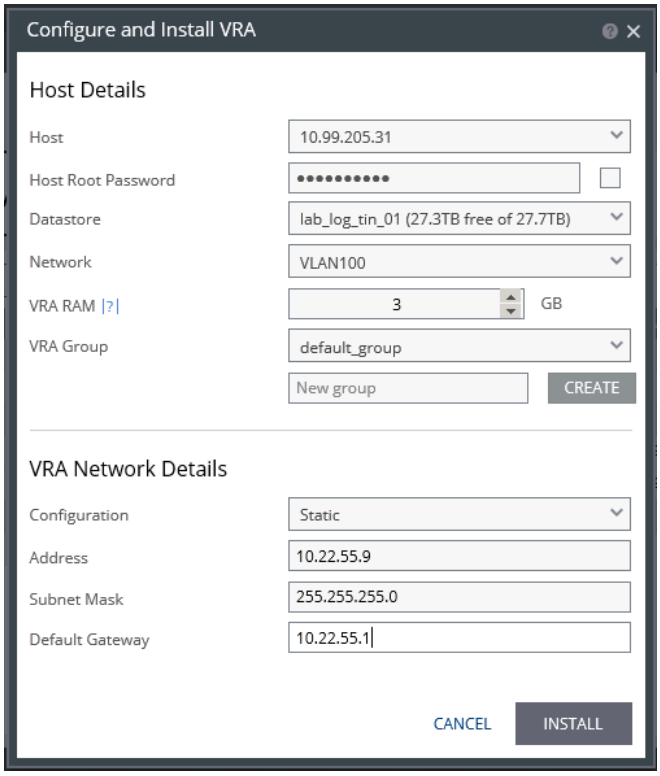

Figure 8 - VRA Setup 2

Select a datastore, select the correct network with connectivity to the GreenCloud side of the VPN, and enter available IP information for the VRA. Leave the VRA group set to “default_group”. Then select “Install” and repeat for every ESX or Hyper-V host. If the process fails, try entering the Host Root password at the top.

Please note that VRAs should never be transferred by vMotion or Storage vMotion. If a storage move is necessary, move all protected VMs to a new host, uninstall the VRA, and re-install it on the correct datastore.

At this point Zerto has been installed, and VPGs can be configured.

VPG Configuration

A Virtual Protected Group, or VPG, is a group of servers protected by Zerto through replication. Once they are replicated, Disaster Recovery onto GreenCloud's IaaS platform is available.

1. Create a new VPG

Select the VPGs tab, then "New VPG" in the upper right-hand corner:

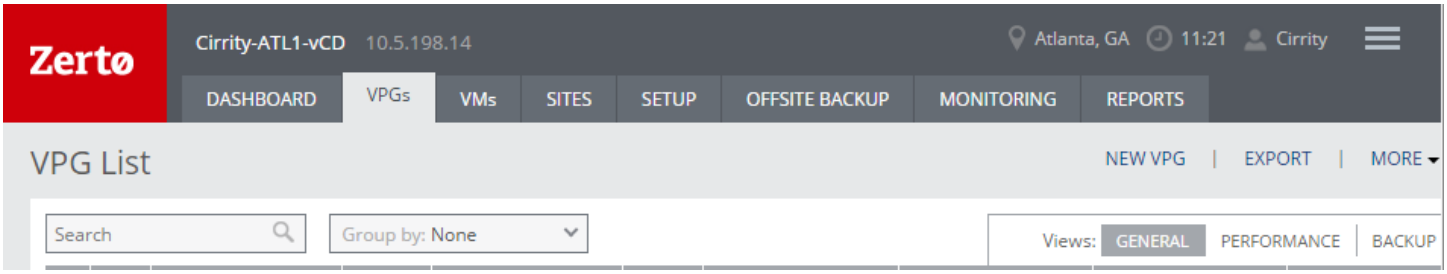

Figure 9 - Add New VPG

2. Select VMs

Enter a VPG name and select Next, then select a VM or group of VMs from the list. This list should include all VMs which can be contacted by the VRA. If a VM does not appear in this list, verify that the VRA deployed to that VM's host is on the same network. Click the arrow to add the selected VMs to the VPG. Define the boot order at this step if multiple dependent VMs are contained in this VPG.

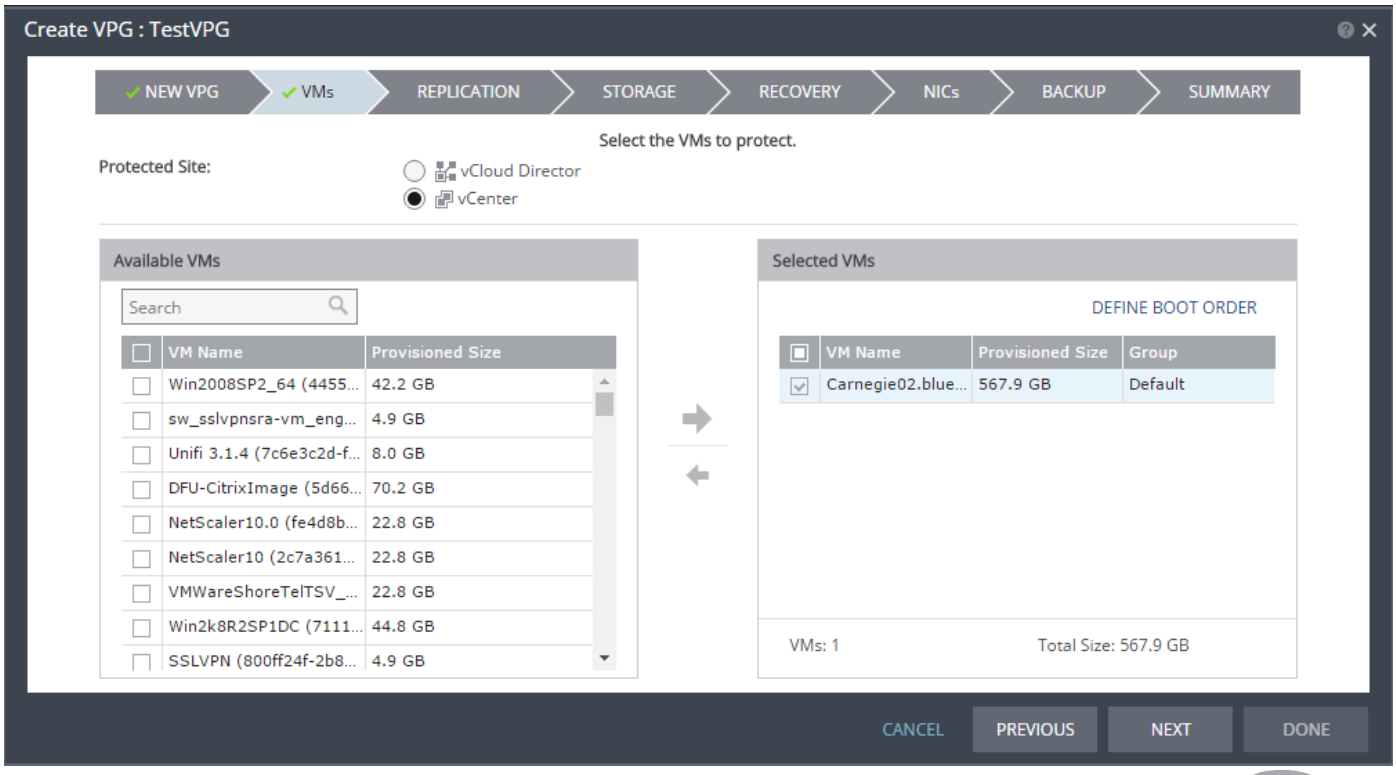

Figure 10 - VPG Configuration: Add VMs

3. Configure Replication and Recovery

At the next screen select the Recovery site. This should be the site defined in the Pair to GreenCloud's site step during installation. Then select the ZORG, which should be the Org ID provided by GreenCloud. The Recovery Org VDC field will appear, and the dropdown should contain only one option. Leave Service Profile as the default. Advanced VM settings can also be configured from this menu, such as Journal Size and Journal Warning Threshold. For more information on these options please contact GreenCloud Support.

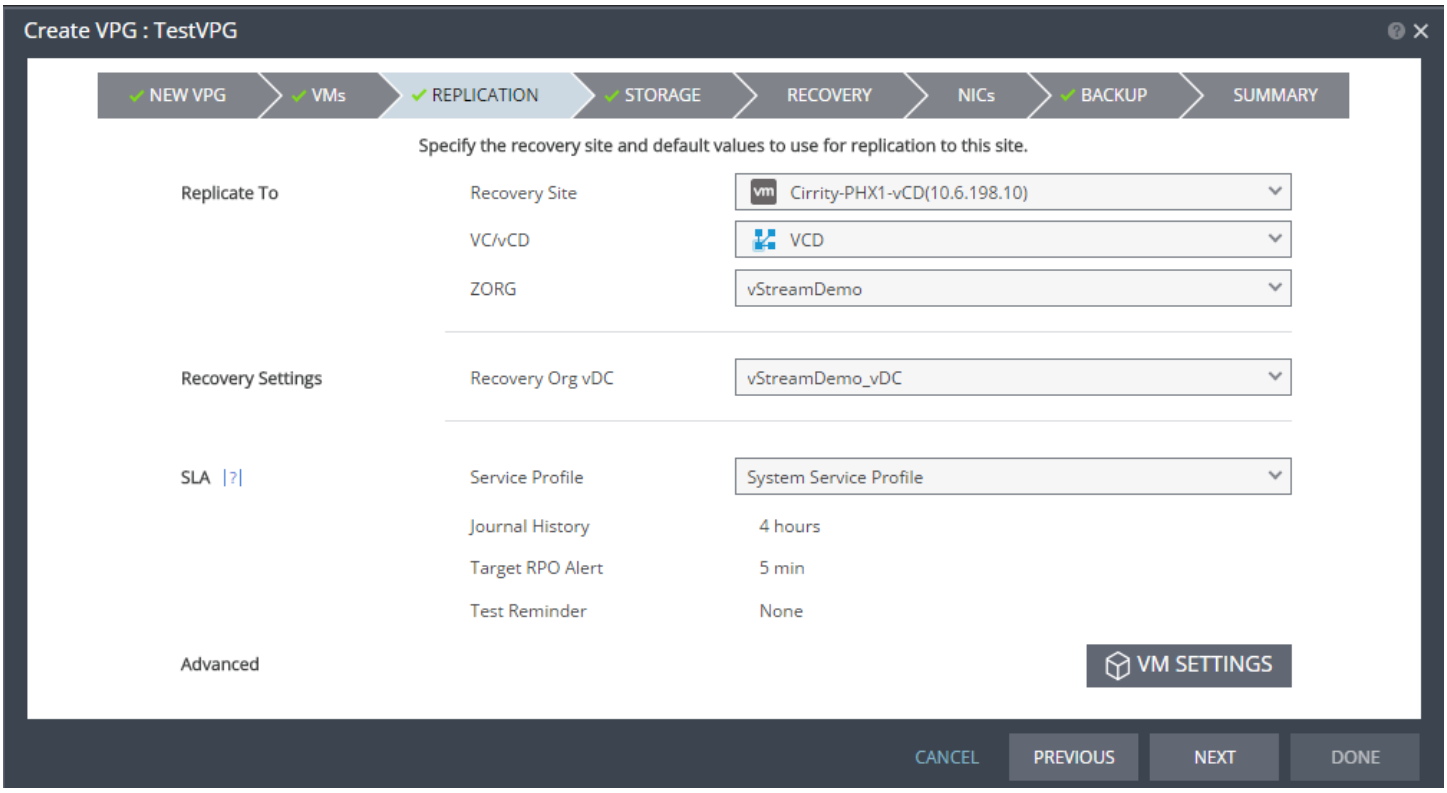

Figure 11 - VPG Configuration: Replication

At the next screen, ensure that the storage options match the source VM's provisioned storage options.

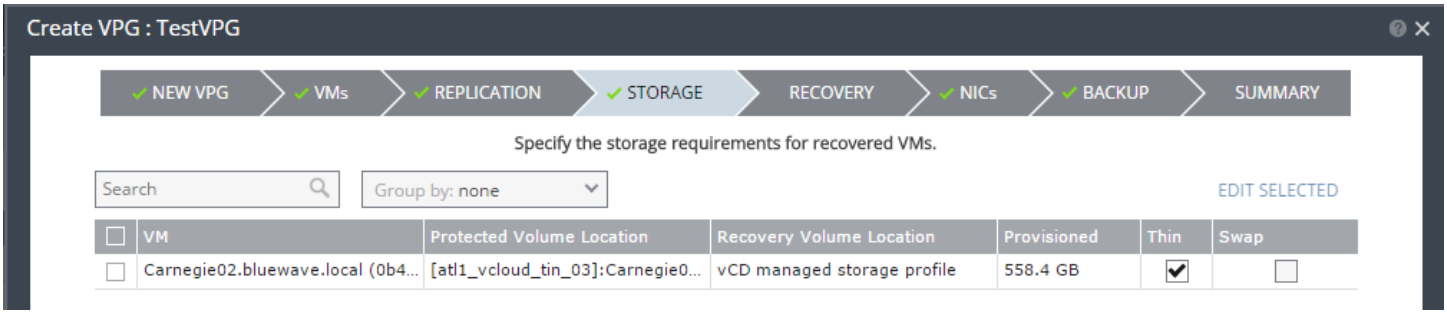

Figure 12 - VPG Configuration: Storage

The Recovery screen configures Guest Customization, networking, and configuration scripts for the replicated image upon restoration. Generally Guest Customization should be off – please see the vCloud Reference Guide for more information regarding Guest Customization. The Failover/Move network and Failover Test network should both be the internal network configured during setup. Configuration scripts can also be configured at this screen.

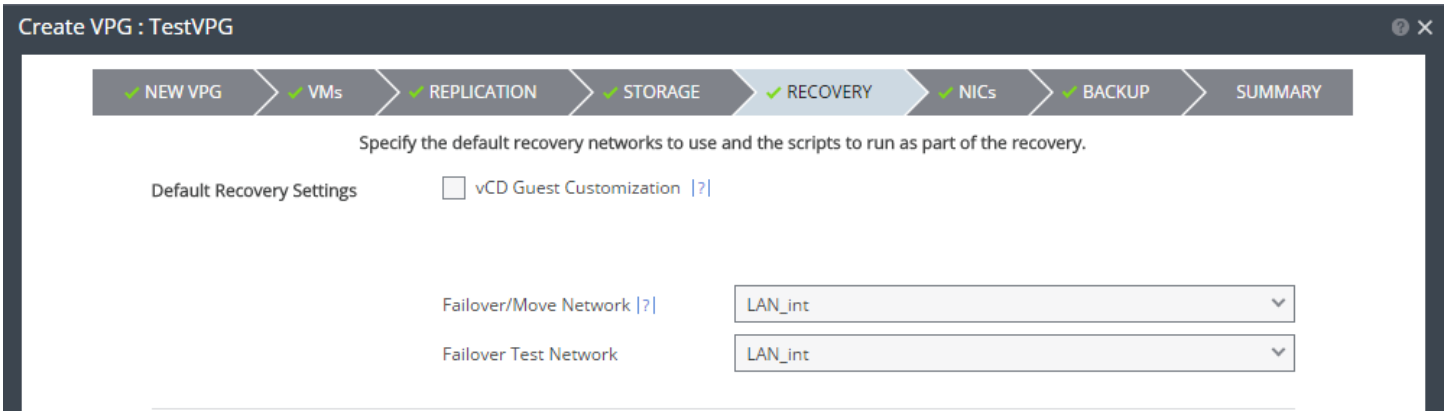

Figure 13 - VPG Configuration: Recovery

The next screen allows configuration of the restored server's NIC. This includes vCloud MAC and IP address. Please note that VMware Tools must be running on the target server in order for these settings to function properly, otherwise the restored VM will automatically receive an IP from the Failover Network.

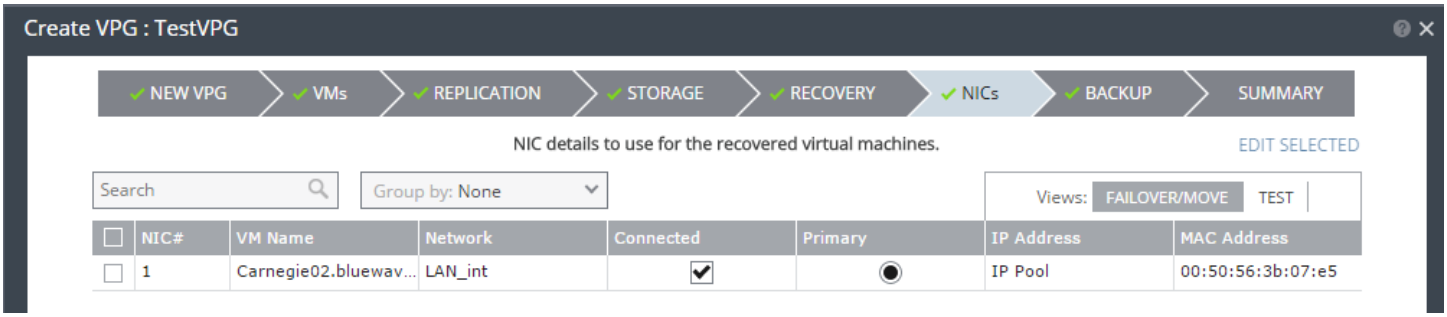

Figure 14 - VPG Configuration: NICs

4. VPG Complete

After this point select "Done". The VPG will be created and replication will begin immediately. Once synchronization is complete, the VM is ready for failover. At this point the VPG will report its status as "Protected". Replication status for each VM is available in the ZVM page.

VPG Statuses

Needs Configuration

One or more configuration settings are missing, for example, when reverse protection is not specified or a virtual machine is added to a vApp

Recovery is Possible

Communication with the Zerto Virtual Manager at the protected site is down so continuing protection is halted, but recovery on the remote site is available (compare with Site disconnection).

VM not protected

A virtual machine in the VPG is no longer being protected. For example, when the virtual machine was moved to another host without a VRA.

Bitmap Syncing

A change tracking mechanism of the protected machines during a disconnected state or when a VRA buffer is full. In these situations, Zerto Virtual Replication starts to maintain a smart bitmap in memory, in which it tracks and records the storage areas that changed. Since the bitmap is kept in memory, Zerto Virtual Replication does not require any LUN or volume per VPG at the protected side

Performing a Failover Test

Failover Test versus Live Failover

There are two types of Failovers available through Zerto. Test Failovers are intended to test configuration and verify image content in preparation for a Live Failover. Live Failovers are intended for a full Disaster Recovery situation, and should only be used when the original (local) VM is offline or unavailable, or for permanent migrations into GreenCloud IaaS.

Test Failovers leave the original VM online, and create a VM using the "Test Network" specified during VPG creation. All writes to the new VM are made to temporary volumes, so Test Failovers cannot last indefinitely as the temporary volumes will eventually run out of space. The original VM remains online, any changes are tracked, and new checkpoints will continue to be generated. Once the test is complete, it can be halted using the ZVM or the ZSSP. This will power off the new VM and resume normal Zerto operation. **Live failovers**, on the other hand, **shut down the original VM** in favor of the new VM. All writes made to the new VM are persistent. Reverse protection is Zerto's method of writing any changes made on the new VM to the original, so as to fail back to the original environment at a later time. Please only enable reverse protection after careful consideration of how the new VM will be handled in GreenCloud's environment, and how this will affect the original VM.

1. Verify "Test" Failover is selected

Check the Failover button in the lower right-hand corner. Verify that the toggle switch indicates "Test" as shown below.

Figure 15 - Failover Test Toggle

Then click the "Failover" button to start the dialog.

2. Configure Failover

Select the name of the VPGs to be tested.

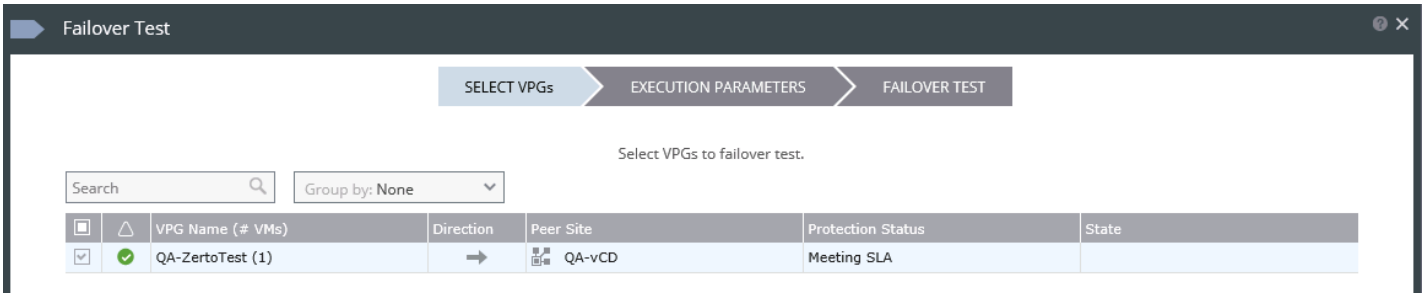

Figure 16 - Failover Test VPG Selection

Change the point in time checkpoint to be the desired restore point, select Next, and choose Start Failover Test.

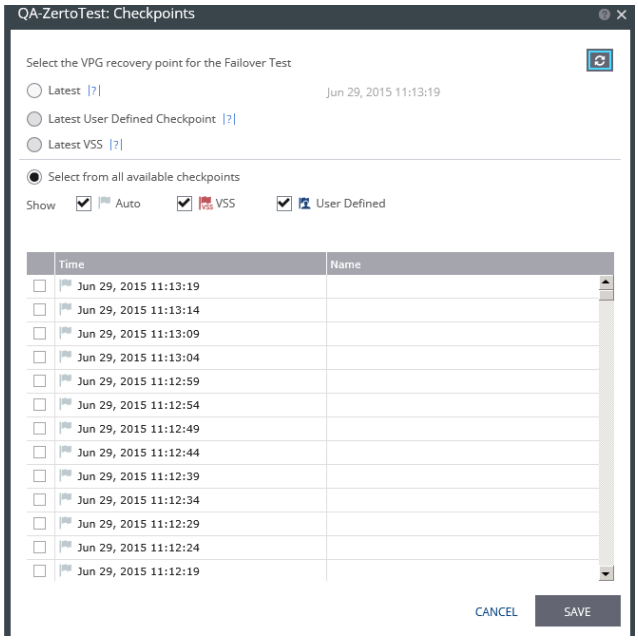

Figure 17 - Failover Test Checkpoint Selection

3. Verify Failover Test

The VM will be created and deployed in the vCloud Org. Once the Failover Test is “waiting for input” in Zerto, the VM is ready for testing. Ensure that the network rules and operating system on the VM are in adequate condition.

4. End Failover Test

Select the Task Waiting for Input and click the Stop button in the test dialog box as shown below.

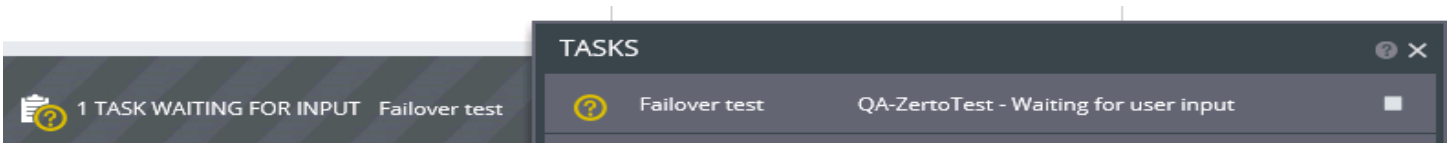

Select a Test status and enter any notes, then select Stop Selected to finish the test and clean up the vCloud org.

Performing a Failover

1. Verify "Live" Failover is Selected

Check the Failover button in the lower right-hand corner. Verify that the toggle switch indicates "Live" as shown below.

Figure 18 - Live Failover Toggle

2. Configure Live Failover

Select the VPG(s) to be failed over or moved as shown below.

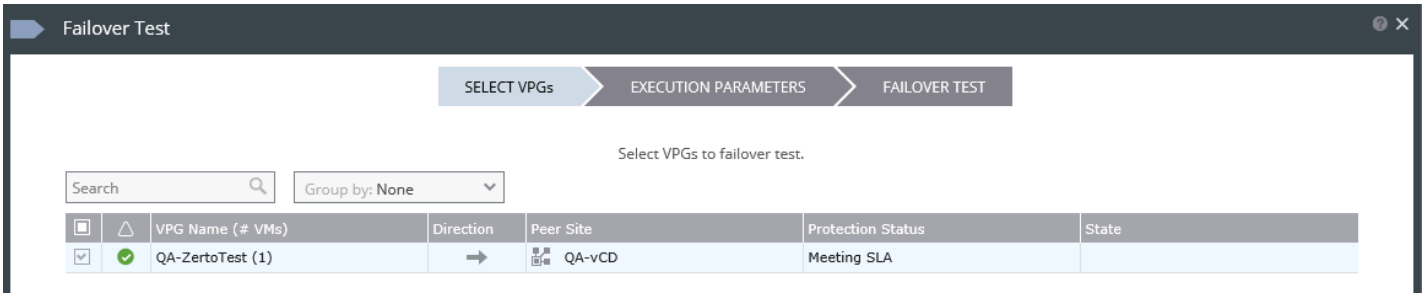

Figure 19 - Failover VPG Selection

For Failovers, select the point in time from which to restore.

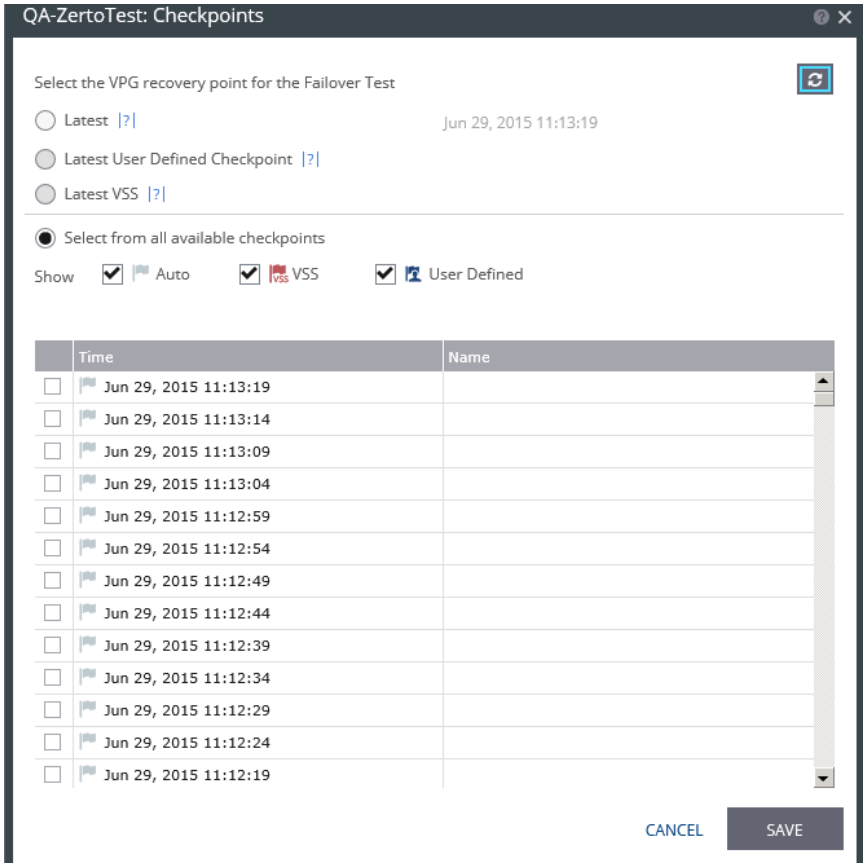

Figure 20 - Failover Checkpoint Selection

Select Reverse Protect All to enable reverse replication after the failover or move so that the VM may be switched back to the primary site when ready. **Note:** Reverse Protection will write all changes made on the restored VM to the source VM, including any failures. **Please only enable reverse protection after careful consideration**. Select a commit policy, choose a shutdown preference, and enable a Boot Order if required. Then select Failover. Once Failover is complete, the VPG status will display as “Needs Configuration” since the source VM will be offline.

3. Configure New VM

Once the failover is complete, a new VM will be in the vCloud org, and will begin the Journal Update process. Once there have been no queued tasks in Zerto for at least 2 minutes the Failover is complete. If reverse replication was configured and disk images are present, a delta synchronization will begin to the primary site. Otherwise, if reverse replication was configured, a full synchronization will begin.

At this time the fully restored VM is accessible and should be up to date. GreenCloud support is always available for assistance should something go wrong during the Failover. Please see the IaaS Reference Guide for more information on configuring the restored server in the vCloud platform.

Appendix A – Glossary of Terms

Zerto Cloud Connector (ZCC)

A virtual machine installed on the cloud side, one for each customer organization replication network. The Zerto Cloud Connector requires both cloud-facing and customer-facing static IP addresses. The ZCC routes traffic between the customer network and the cloud replication network, in a secure manner ensuring complete separation between the customer network and the cloud service provider network. The ZCC has two Ethernet interfaces, one to the customer's network and one to the cloud service provider's network. Within the cloud connector a bidirectional connection is created between the customer and cloud service provider networks. Thus, all network traffic passes through the ZCC, where the incoming traffic on the customer network is automatically configured to IP addresses of the cloud service provider network.

Zerto Cloud Manager (ZCM)

A Windows service, which enables managing all the cloud sites offering disaster recovery using a single interface. The ZCM manages the DR either as a service (DRaaS) or completely within the cloud environment, protecting on one cloud site and recovering to a second site (ICDR).

Virtual Protection Group (VPG)

Virtual machines are protected in virtual protection groups. A virtual protection groups (VPG) is a group of virtual machines that you want to group together for replication purposes. For example, the virtual machines that comprise an application like Microsoft Exchange, where one virtual machine is used for the software, one for the database and a third

Virtual Replication Appliance (VRA)

A virtual machine that manages the replication of protected virtual machine writes across sites. A VRA must be installed on every ESX/ESXi which hosts virtual machines that require protecting in the protected site and on every ESX/ESXi that will host the replicated virtual machines in the recovery site.

Zerto Self-service Portal (ZSSP)

A website for Green Cloud customers to initiate failover or Move operations on their own

Recovery Point Objective (RPO)

The maximum amount of data that may be lost when the activity or service is restored after an interruption. Expressed as a length of time before the interruption.

Appendix B – Revision History

AUTHOR	DATE	COMMENTS	VER.
Stuart Carmichael	2013-04-29	Initial publication	
Stuart Carmichael	2013-09-11	Updated to match Zerto 3.0 look	
Stuart Carmichael	2014-01-07	Updated for Zerto 3.1 changes	
Stuart Carmichael	2015-06-29	Updated for Zerto 4.0 changes	
Alex Reid	2017-07-17	GreenCloud Content Update	
Alex Reid	2017-08-17	GreenCloud formatting update	0.1
Alex Reid	2017-09-01	Added Glossary, minor content updates	
Alex Reid	2017-10-11	Various corrections	1.0