

Zerto Virtual Replication provides a business continuity (BC) and disaster recovery (DR) solution in a virtual environment, enabling the replication of mission-critical applications and data as quickly as possible, with minimal data loss. When devising a recovery plan, these two objectives, minimum time to recover and maximum data to recover, are assigned target values: the recovery time objective (RTO) and the recovery point objective (RPO). Zerto Virtual Replication enables a virtual-aware recovery with low values for both the RTO and RPO. In addition, Zerto Virtual Replication enables protecting virtual machines for extended, longer term recovery from an offsite backup.

This document provides a quick guide to setting up Zerto Virtual Replication in a VMware vCenter Server environment to protect virtual machines.

Table of Contents

Introduction	1
Recommended Installation Best Practices	3
Installation	3
Registering the Zerto Virtual Replication License	4
Installing Virtual Replication Appliances	5
Pairing Sites to Enable Replicating From One Site to Another Site	7
Setting Up the Second Site	7
Enabling Replication to the Same Site	8
Protecting Virtual Machines	9
Testing Disaster Recovery	13

Introduction

Zerto Virtual Replication is installed in every site with virtual machines to be protected and recovered. The installation includes:

Zerto Virtual Manager (ZVM) - A Windows service that manages replication at the site level. The ZVM monitors the vCenter Server to get the inventory of VMs, disks, networks, hosts, etc. For example, a VMware vMotion operation of a protected VM from one host to another is monitored by the ZVM and protection and replication is updated accordingly.

Virtual Replication Appliance (VRA) - A virtual machine installed on each hypervisor hosting virtual machines to be protected or recovered, that manages the replication of data between the protected and recovery sites.

Virtual Backup Appliance (VBA) - A Windows service that manages back-ups within Zerto Virtual Replication. The VBA service runs on the same machine as the Zerto Virtual Manager service and manages the repositories where offsite backups are stored. These repositories can be local or on a shared network.

Zerto User Interface - Recovery using Zerto Virtual Replication is managed in a browser. Each site is managed via the Zerto User Interface, accessed from a browser or from within the vSphere Web Client or Client console.

Zerto Virtual Replication also supports both the protected and recovery sites being managed by a single vCenter Server, to handle small branch offices. When the protected and recovery sites are the same site, only one installation of Zerto Virtual Replication is required.

Requirements for Each Site

- VMware vCenter Server version 4.0U1 and higher with at least one ESX/ESXi host. The Zerto Virtual Manager must have access to the vCenter Server via a user with administrator level privileges to the vCenter Server. When the vCenter Server is installed on a Linux machine via the vCenter Server Linux Virtual Appliance (vCSA), Zerto Virtual Replication must still be installed on a Windows machine.
- On the machines where Zerto Virtual Replication is installed:
 - Windows Server 2008 R2 SP1 with KB3033929 and KB2864202, Windows Server 2012 base, or Windows Server 2012 R2.
Reserve at least 2 CPUs and 4GB RAM for the machine. The following CPU and RAM are recommended by Zerto for the machine running Zerto Virtual Replication, dependent on the size of the site:
Sites protecting up to 750 virtual machines and up to 5 peer sites: 2 CPUs and 4GB RAM
Sites protecting 751-2000 virtual machines and up to 15 peer sites: 4 CPUs or 2 Dual Core CPUs and 4GB RAM
Sites protecting over 2000 virtual machines and over 15 peer sites: 8 CPUs or 4 Dual Core CPUs and 8GB RAM
The clocks on the machines where Zerto Virtual Replication is installed must be synchronized with UTC and with each other (the timezones can be different). Zerto recommends synchronizing the clocks using NTP.
 - Microsoft .NET Framework 4.5.2. (included with the Zerto Virtual Replication installation kit) or higher.
 - At least 4GB of free disk space (plus 1.8GB if you need to install Microsoft .NET Framework).

Routable Networks

The Zerto Virtual Replication architecture supports the following network configurations:

- Flat LAN networks
- VLAN networks, including private VLANs and stretched VLANs
- WAN emulation
- VPN IPsec

The Zerto Virtual Replication architecture does not support NAT (Network Address Translation) firewalls.

Minimum Bandwidth

The connectivity between sites must have the bandwidth capacity to handle the data to be replicated between the sites. The minimum bandwidth must be at least 5 Mb/sec.

The Zerto Virtual Manager Web Client

Microsoft Windows Explorer 9 is not supported and version 10 does not work well with the user interface. Zerto recommends using Chrome, Firefox, or later versions of Internet Explorer. The minimum recommended screen resolution is 1024*768.

Open Firewall Ports

Zerto Virtual Manager requires the following ports to be open in the protected and recovery site firewalls:

PORT	DESCRIPTION
22 ^a	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.
443	During VRA installation on ESX/ESXi hosts for communication between the ZVM and the ESX/ESXi hosts IPs. Also, for ongoing communication between the ZVM and VMware vCenter Server.
4005	Log collection between the ZVM and VRAs on the same site.
4006	TCP communication between the ZVM and VRAs and the VBA on the same site.
4007	TCP control communication between protecting and recovering VRAs.
4008	TCP communication between VRAs to pass data from protected virtual machines to a VRA on a recovery site.
4009	TCP communication between the ZVM and site VRAs to handle checkpoints.
9080	HTTP communication between the ZVM and Zerto internal APIs, and a VSS Agent.
9081 ^b	TCP communication between ZVMs ^c .

PORT	DESCRIPTION
9180	Communication between the VBA and VRA.
9669	HTTPS communication between the Zerto User Interface and a ZVM, and for invoking Zerto RESTful APIs.

- If the ESX/ESXi hosts are given names, make sure that the Zerto Virtual Manager can resolve these names.
- The default port set during the Zerto Virtual Replication installation. When pairing the ZVM to a Zerto Cloud Connector, this value must not be changed.
- When the same vCenter Server is used for protection and recovery, Zerto Virtual Replication is installed on one site only and this port is ignored.

If a proxy server is used at the site, specify the IP address of the Zerto Virtual Manager in the exception list in the Proxy Server settings.

Recommended Installation Best Practices

Zerto recommends the following best practices:

- Install Zerto Virtual Replication on a dedicated virtual machine with a dedicated administrator account and with VMware high availability (HA) enabled, and no other applications installed on this machine, and especially not on the machine running the vCenter Server service. If other applications are installed, the Zerto Virtual Manager service must receive enough resources and HA must remain enabled.
- Install a VRA on every host in a cluster so that if protected virtual machines are moved from one host to another, there is always a VRA to protect the moved virtual machines.
- Install VRAs using static IP addresses and not DHCP.
- Prepare an administrator account for the machine where Zerto Virtual Replication is installed.
- It is required to exclude the Zerto Virtual Replication folder from antivirus scanning. Failure to do so may lead to the ZVR folder being incorrectly identified as a threat and in some circumstances corrupt the ZVR folder.
- Synchronize the clocks on the machines where Zerto Virtual Replication is installed using NTP.

Installation

The Zerto Virtual Replication installation deploys the Zerto Virtual Manager and copies the installation software for the Virtual Replication Appliance.

A complete installation includes installing Zerto Virtual Replication on the protected and peer, recovery, sites.

Note: When both these sites are managed by a single vCenter Server, Zerto Virtual Replication is installed on only one site. In this case, Zerto recommends installing Zerto Virtual Replication in the main site where protected machines will be recovered.

You can install Zerto Virtual Replication using the defaults provided by Zerto or perform a custom install, in which you can define the ports that will be used by Zerto Virtual Replication.

Performing an Express Installation

You can install Zerto Virtual Replication using the defaults provided by Zerto. Site information and information to connect to vCloud Director can be provided, if required, after the installation in the Zerto User Interface.

Note: You cannot install Zerto Virtual Replication on the same machine where another version of Zerto Virtual Replication has been installed, for example, if the *Zerto Virtual Replication for Microsoft Hyper-V* version has been installed on the machine.

To perform an express install of Zerto Virtual Replication:

1. Run `Zerto Virtual Replication Installer.exe`.

Note: If the required version of Microsoft .NET Framework is not installed, you are prompted to install the required version of .NET Framework, which is included as part of the Zerto Virtual Replication installation package. After .NET is installed the machine automatically restarts and the Zerto Virtual Replication installation begins.

- Follow the wizard through the installation until the *Choose Installation Type* dialog and select the *Express* installation option.
 - Click *NEXT*.
- The *vCenter Server Connectivity* dialog is displayed.

- Specify the following:
 - IP / Host Name** - The IP address or host name of the machine where the vCenter Server runs.
 - Username** - The user name of a user with administrator level privileges in the vCenter Server. The name can be entered using either of the following formats:
 - username
 - domain\username
 - Password** - A valid password for the given user name.
 - Site Name** - A name to identify the site.
 - Click *NEXT*.
- The *Validation* dialog is displayed.
- The installation performs checks to make sure that the installation can proceed successfully.
- After the checks complete successfully, click *RUN* and continue to the end of the installation.
 - If you intend to manage disaster recovery from this machine, open the Zerto User Interface at the end of the installation, logging in with the user name and password for the vCenter Server connected to the Zerto Virtual Manager.
 - Set any antivirus software running on the machine not to scan the folder where Zerto Virtual Replication is installed.
 - Repeat the procedure to install Zerto Virtual Replication on the peer site.

Registering the Zerto Virtual Replication License

Access the Zerto User Interface from a browser as follows:


To use the Zerto Virtual Manager Web Client:

- In a browser, enter the following URL:


```
https://zvm_IP:9669
```

 where *zvm_IP* is the IP address of the Zerto Virtual Manager for the site you want to manage.
- Login using the user name and password for the vCenter Server connected to the Zerto Virtual Manager.

On the very first access to the Zerto User Interface, you must either register your use of Zerto Virtual Replication, by entering the license key supplied by Zerto or pair to a site where a license has already been entered.



After entering a valid license, the *DASHBOARD* tab is displayed with a summary of the site. Before you can start protecting virtual machines in this site, you must configure Zerto Virtual Manager at each site by installing Virtual Replication Appliances on the hosts in the site and pair the protected and recovery sites, as described in the following sections.

Note: Complete the configuration of one site at a time.

Installing Virtual Replication Appliances

The Zerto Virtual Replication installation includes the OVF template for Virtual Replication Appliances (VRAs). A VRA is a Zerto Virtual Replication virtual machine that manages the replication of virtual machines across sites. A VRA must be installed on every host that manages virtual machines that require protecting in the protected site and on every host that manages virtual machines in the recovery site. The VRA compresses the data that is passed across the WAN from the protected site to the recovery site. The VRA automatically adjusts the compression level according to CPU usage, including totally disabling it if needed.

A VRA can manage a maximum of 1500 volumes, whether these volumes are being protected or recovered.

The VRA is a custom, very thin, Linux-based virtual machine with a small footprint, disk – memory and CPU – and increased security since there are a minimum number of services installed.

Zerto recommends installing a VRA on every hypervisor host so that if protected virtual machines are moved from one host in the cluster to another host in the cluster there is always a VRA to protect the moved virtual machines.

VRA Installation Requirements

To install a VRA you require the following:

- 12.5GB datastore space.
- At least 1GB of reserved memory.
- The ESX/ESXi version must be 4.0U1 or higher and Ports 22 and 443 must be enabled on the host during the installation.

You must also know the following information to install a VRA:

- If the ESXi version is 5.5 or higher and the VRA should connect to the host with user credentials, or if the ESXi version is lower than 5.5 (4.x or 5.x), the password to access the host root account.

Note: For ESXi versions 5.5 or higher, by default the VRA connects to the host with a vSphere Installation Bundle, VIB. Therefore, it is not necessary to enter the password used to access the host root account.

- The storage the VRA will use and the local network used by the host.
- The network settings to access the peer site; either the default gateway or the IP address, subnet mask, and gateway.
- If a static IP is used, which is the Zerto recommendation¹, instead of DHCP, the IP address, subnet mask, and default gateway to be used by the VRA.

Note: For the duration of the installation of the VRA, the Zerto Virtual Manager enables SSH in the vCenter Server.

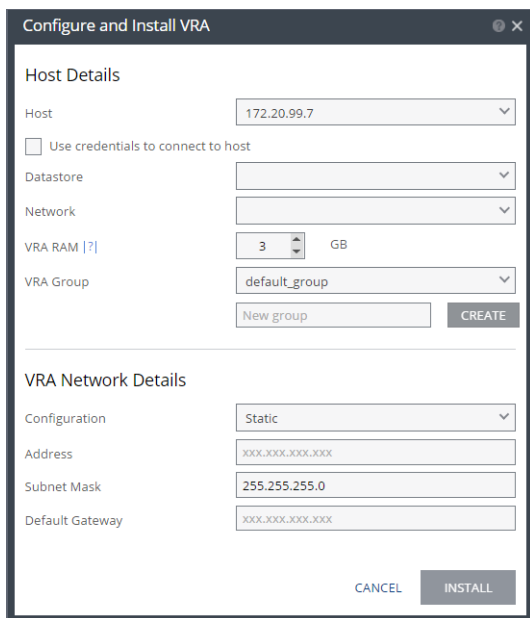
1. In a non-production environment it is often convenient to use DHCP to allocate an IP to the VRA. In a production environment this is not recommended. For example, if the DHCP server changes the IP allocation on a reboot, the VRA does not handle the change.

If the peer site VRAs are not on the same network as the peer site Zerto Virtual Manager, you must set up routing to enable the Zerto Virtual Manager to communicate with the peer site VRAs, as described in the *Zerto Virtual Manager Administration Guide for the VMware vSphere Environment*.

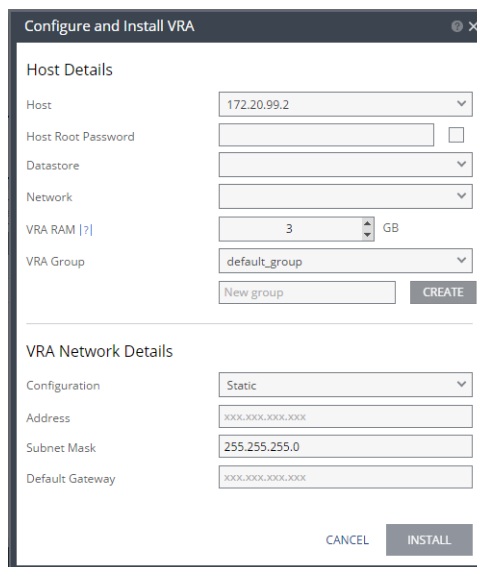
To install Zerto Virtual Replication Appliances (VRAs) on ESX/ESXi hosts:

1. In the Zerto User Interface, click *SETUP > VRAs*.
2. Select a host which requires a VRA and click *NEW VRA*.

The *Configure and Install VRA* dialog is displayed. The dialog displayed depends on the ESX/i version:



ESXi versions from 5.5



ESXi versions before version 5.5

- Note:** If you selected a cluster or multiple hosts, the VRA is installed on the first host in the displayed list.
3. Specify the following in the *Host Details* section:
 - Host** - The host on which the VRA is installed. The drop-down displays the hosts that do not have a VRA installed, with the selected host displayed by default.

From ESXi 5.5, by default, Zerto Virtual Manager uses a vSphere Installation Bundle, VIB, to connect to the host. When using VIB:

 - The user does not enter a password.
 - Once a day, Zerto Virtual Manager checks that the VRA and host can connect. If the connection fails, Zerto Virtual Manager re-initiates the connection automatically and logs it.

For ESX/i versions earlier than 5.5, when using a password, root access is required. Once a day, Zerto Virtual Manager checks that the password is valid. If the password was changed, an alert is issued, requesting the user enter the new password.

 - Use credentials to connect to host** - When unchecked, the Zerto Virtual Manager uses VIB to connect to the host. This field is only relevant for ESXi 5.5 and later.
 - Host Root Password** - When the VRA should connect to the host with a password, check *Use credential to connect to host* and enter the root user password used to access the host. When the box on the right side is checked, the password is displayed in plain text. This field is only relevant for ESXi 4.x and 5.x hosts. This field is disabled for ESX 4.x hosts.
 - Datastore** - The datastore that the VRA will use for protected virtual machine data on the recovery site, including the journals. You can install more than one VRA on the same datastore.
 - Network** - The network used to access the VRA.
 4. Leave the *VRA RAM* and *VRA Group* values with their defaults.
 5. Specify the following in the *VRA Network Details* section:

Configuration – Specify DHCP.

Note: In a production environment the `Static` option is the recommended option.

6. Click *INSTALL*.

The VRA installation starts and the status is displayed in the *TASKS* popup dialog in the status bar and under *MONITORING > TASKS*.

The VRA displayed name and DNS name is `z-vra-hostname`. If a virtual machine with this name exists, for example when a previous VRA was not deleted, the VRA name has a number appended to it.

Add a VRA to every host that hosts virtual machines that you want replicated. Zerto recommends installing a VRA on every listed host. An alert is issued after the first VRA is installed in a cluster that tells you to install a VRA on the other hosts in the cluster. The alert is automatically removed when all the hosts in the cluster have VRAs installed.

Note: VRAs are configured and managed by the Zerto Virtual Manager. You cannot take snapshots of VRAs as snapshots cause operational problems for the VRAs.

Pairing Sites to Enable Replicating From One Site to Another Site


Zerto Virtual Replication is installed on both the protected and recovery sites and these two sites are paired to enable disaster recovery across the sites.

If a single vCenter is used, for example with remote branch offices, when replicating from one datacenter to another datacenter, both managed by the same vCenter Server, you have to enable replication to the same vCenter Server and pairing is not required. For details, see [“Enabling Replication to the Same Site”, on page 8](#).

To pair sites:

1. In the Zerto User Interface, in the *SITES* tab click *PAIR*.

The *Add Site* dialog is displayed.



2. Specify the following:

Remote Site ZVM IP Address – IP address or fully qualified DNS host name of the remote site Zerto Virtual Manager to pair to.

Port – The TCP port communication between the sites. Enter the port that was specified during installation. The default port during the installation is 9081.

3. Click *PAIR*.

The sites are paired meaning that the Zerto Virtual Manager for the local vCenter site is connected to the Zerto Virtual Manager on the remote vCenter site.

After the pairing completes the content of the *SITES* tab changes to include summary information about the paired site.

Setting Up the Second Site

After pairing the sites you can access the second site without entering a license and then install VRAs in the site.

Install VRAs on hosts in the second site by repeating the procedure, [“Installing Virtual Replication Appliances”, on page 5](#).

Enabling Replication to the Same Site

When a single vCenter is used, for example with remote branch offices, when replicating from one datacenter to another datacenter, both managed by the same vCenter Server, you must enable replication to the same vCenter Server and pairing is not required. In this case, replication to the same vCenter must be set in the *Site Settings* dialog.

To enable replication to the same vCenter Server:

1. In the Zerto User Interface, click *SETTING* (☰) in the top right of the header and select *Site Settings*. The *Site Settings* dialog is displayed.

2. Click *Policies*.

3. Check the *Enable Replication to Self* checkbox.
4. Click *SAVE*.

When used to protect to itself, the Zerto Virtual Manager can protect up to 5000 virtual machines.

Protecting Virtual Machines

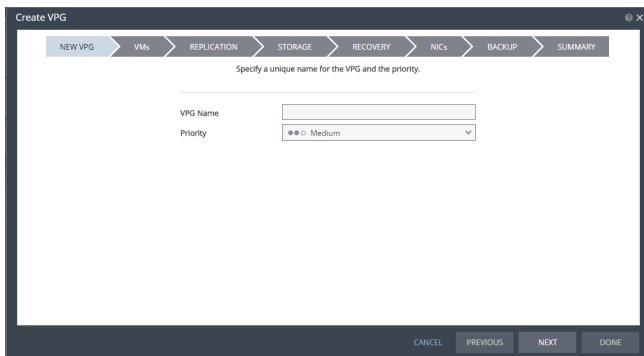
You can protect virtual machines to a recovery site vCenter Server. The procedure is the same whether you intend to protect one virtual machine or multiple virtual machines.

Note: You cannot protect virtual machines with VirtualEthernetCardLegacyNetworkBackingInfo NICs or with IDE devices.

To create a virtual protection group (VPG):

1. In the Zerto User Interface, select **ACTIONS > CREATE VPG**.

The **NEW VPG** step of the *Create VPG* wizard is displayed.



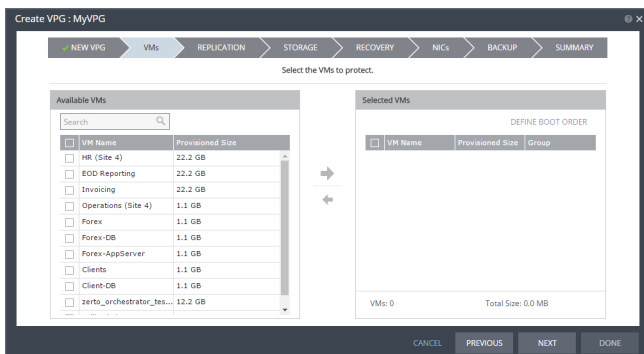
2. Specify the name of the VPG and the priority of the VPG.

VPG Name – The VPG name must be unique.

Priority – Determine the priority for transferring data from the protected site to the recovery site when there is limited bandwidth and more than one VPG is defined on the protected site. When there are updates to virtual machines protected in VPGs with different priorities, first the updates from the VPG with the highest priority are passed over the WAN. Medium priority VPGs will only be able to use whatever bandwidth is left after the high priority VPGs have used it. This is also true between medium and low priorities.

3. Click **NEXT**.

The **VMs** step is displayed.



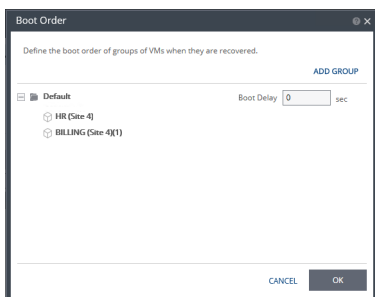
4. Select the VMs that will be part of this VPG and click the arrow pointing right to include these VMs in the VPG.

Zerto Virtual Replication uses the SCSI protocol. Only virtual machines with disks that support this protocol can be specified.

5. If you want to define the boot order of the virtual machines in the VPG, click **DEFINE BOOT ORDER**, otherwise, go to the next step.

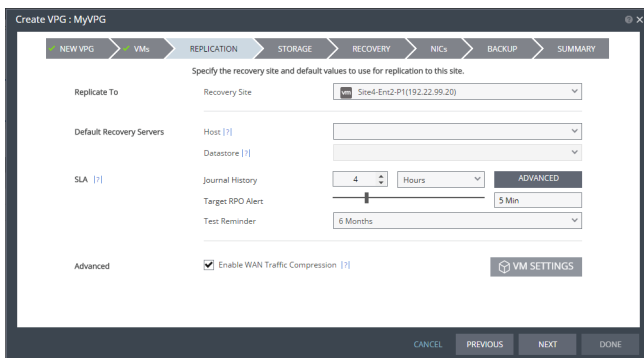
When virtual machines in a VPG are started in the recovery site, by default these machines are not started up in a particular order. If you want specific virtual machines to start before other machines, you can specify a boot order. The virtual machines are defined in groups and the boot order applies to the groups and not to individual virtual machines in the groups. You can specify a delay between groups during startup.

Initially, virtual machines in the VPG are displayed together under the *Default* group. If you want specific machines to start before other virtual machines, define new groups with one or more virtual machines in each group.



- a) Click *ADD GROUP* to add a new group.
 - b) To change the name of a group, click the Pencil icon next to the group. To delete a group, click the delete icon on the right side. You cannot delete the *Default* group nor a group that contains a virtual machine.
 - c) Drag virtual machines to move them from one group to another.
 - d) Drag groups to change the order the groups are started.
 - e) Optionally, in *Boot Delay*, specify a time delay between starting up the virtual machines in the group and starting up the virtual machines in the next group. For example, assume three groups, *Default*, *Server*, and *Client*, defined in this order. The boot delay defined for the *Default* group is 10, for the *Server* group is 100, and for the *Client* group 0. The virtual machines in the *Default* group are started together and after 10 seconds the virtual machines in the *Server* group are started. After 100 seconds the virtual machines in the *Client* group are started.
 - f) Click *OK*.
6. Click *NEXT*.

The *REPLICATION* step is displayed.



Note: If the protected site is paired with only one recovery site, the recovery step is displayed with the *Recovery Site* field automatically filled in and defaults set for the *SLA* and *Advanced* settings.

7. Specify the values to use when replicating to this site.

Recovery Site – The site to which you want to recover the virtual machines. After specifying the recovery site, the host and datastore on the site to use for replication can be specified. If you specified that replication is possible to the same site, as described in “[Enabling Replication to the Same Site](#)”, on page 8, then you can specify the local site here.

Host – The default cluster, resource pool or host in the recovery site that handles the replicated data.

When a resource pool is specified, Zerto Virtual Replication checks that the resource pool capacity is enough for any virtual machines specified in the VPG.

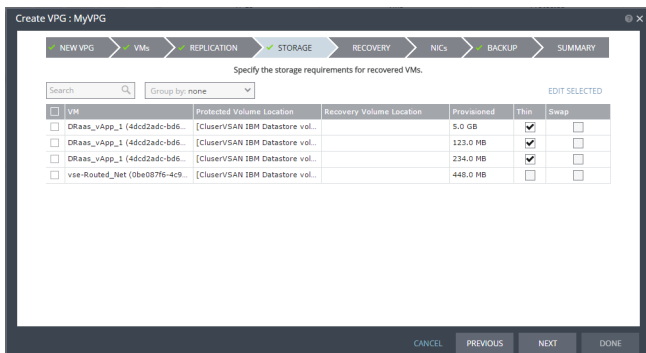
All resource pool checks are made at the level of the VPG and do not take into account multiple VPGs using the same resource pool. If the resource pool CPU resources are specified as unlimited, the actual limit is inherited from the parent but if this inherited value is too small, failover, move, and failover test operations can fail, even without a warning alert being issued by Zerto Virtual Manager.

Note that if a resource pool is specified and DRS is disabled for the site later on, all the resource pools are removed by VMware and recovery will be to any one of the hosts in the recovery site with a VRA installed on it.

Datastore – The datastore volume to use for all recovered virtual machine files as well as for their data volumes. Every datastore for the selected recovery host is included in the drop-down list. If a cluster or resource pool is selected for the host, only datastores that are accessible by every host in the cluster or resource pool are displayed. When specifying the recovery storage for a virtual machine with a storage cluster, you must specify a datastore in the cluster.

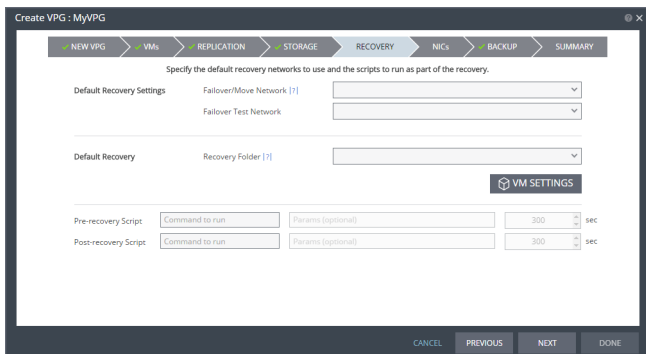
- The following settings can be changed later by editing the VPG definition. For your first VPG, leave the default values and click *NEXT*.

After clicking *NEXT*, the *STORAGE* step is displayed. By default the storage used for the virtual machine definition is also used for the virtual machine data. For each virtual machine in the VPG, Zerto Virtual Replication displays its storage-related information.



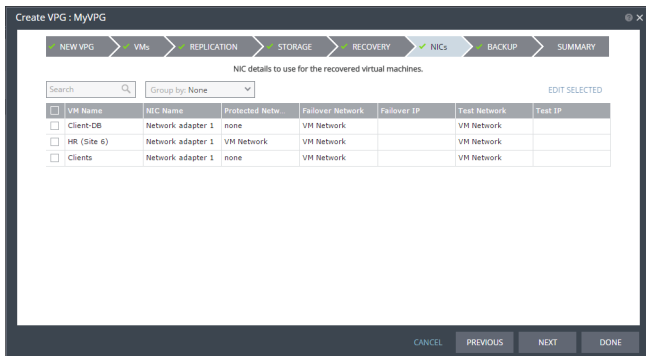
Note: Steps that do not require input are marked with a check mark. You can jump directly to a step that has been marked with a check mark to edit the values for that step. Every step must be marked with a check mark before you can click *DONE* to create the VPG.

- For your first VPG, leave the default values. These settings can be changed later by editing the VPG definition. Click *NEXT*. The *RECOVERY* step is displayed. Recovery details include the networks to use for failover, move, and for testing failover, and whether scripts should run as part of the recovery operation.



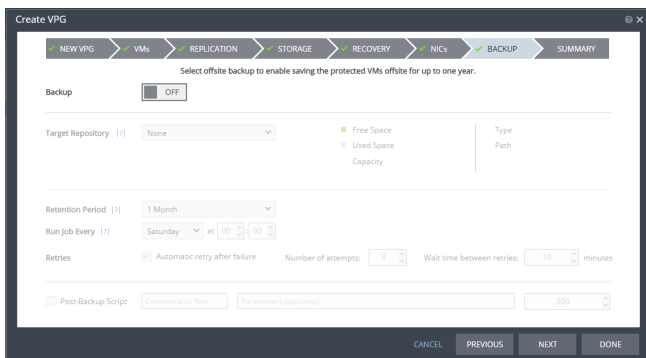
- Select the recovery settings.
 - Failover/Move Network** – The network to use during a failover or move operation in which the recovered virtual machines will run.
 - Failover Test Network** – The network to use when testing the failover of virtual machines in the recovery site. Zerto recommends using a fenced-out network so as not to impact the production network at this site.
 - Recovery Folder** – The folder to which the virtual machine is recovered.
- You can run scripts as part of the recovery process, and these scripts are defined in the VPG. Script settings can be added later by editing the VPG definition. For your first VPG, leave the default values and click *NEXT*.

The *NICs* step is displayed. In this step, you can specify the NIC details to use for the recovered virtual machines after a failover, a test failover, or migration.



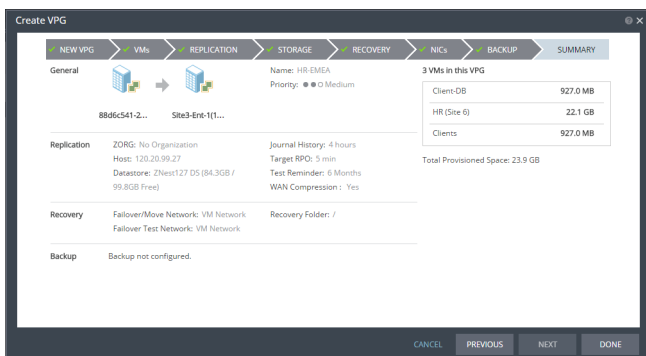
12. Again, leave the defaults and click *NEXT*.

The *BACKUP* step is displayed. Backup properties govern the VPG backup, including the repository where the backups are saved. Backup extends the ability to recover virtual machines in a VPG going back one year.



13. Again, leave the defaults and click *NEXT*.

The *SUMMARY* step is displayed. It shows the VPG configuration that you defined in previous tabs.



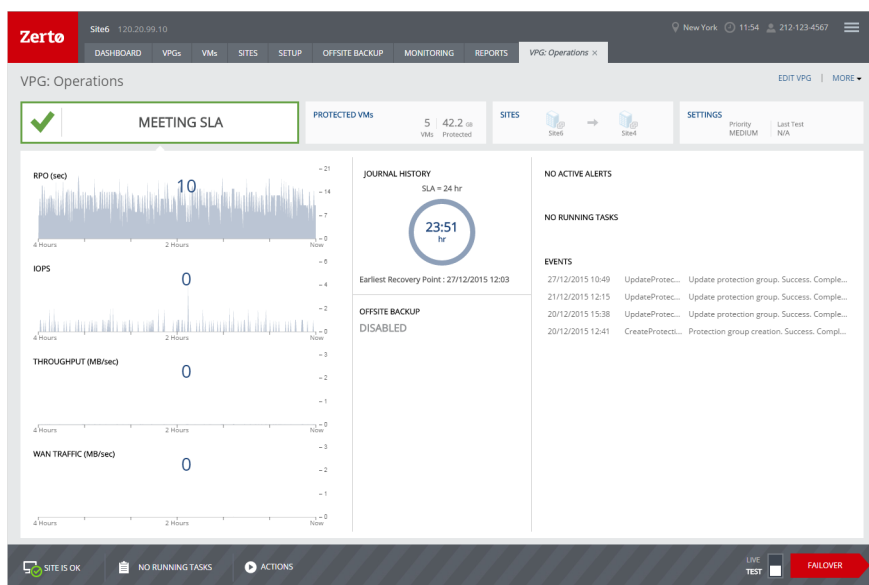
14. Click *DONE*.

The VPG is created.

The VRA in the recovery site is updated with information about the VPG and then the data on the protected virtual machines are synchronized with the replication virtual machines managed by the VRA on the recovery site. This process can take some time, depending on the size of the VMs and the bandwidth between the sites.

Note: For synchronization to work, the protected virtual machines must be powered on.

Once synchronized, the VRA on the recovery site includes a complete copy of every virtual machine in the VPG. After synchronization, the virtual machines in the VPG are fully protected, meeting their SLA, and the delta changes to these virtual machines are sent to the recovery site.



In order to verify that the disaster recovery that you have planned is the one that will be implemented, Zerto recommends testing the recovery of the VPGs defined in the protected site to the recovery site.

Testing Disaster Recovery

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 creating virtual machines in 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 storage defined for the journal.

The Failover Test operation has the following basic steps:

- Starting the test.
 - The test virtual machines are created at the remote site using the network specified for testing in the VPG settings and configured to the checkpoint specified for the recovery. The test virtual machines are created without CD-ROM drives, even if the protected virtual machines have CD-ROM drives.
 - The virtual machines are powered on, making them available to the user. If applicable, the boot order defined in the VPG settings is used to power on the machines.
- Testing. The virtual machines in the VPG are created as test machines in a sandbox and powered on for testing using the test network specified in the VPG definition and 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 storage defined for the journal. Using scratch volumes makes cleaning up the test failover more efficient.

Note: You must not delete, clone, migrate to another host or change the disk properties of any of the test virtual machines.
- Stopping the test.
 - The test virtual machines are powered off and removed from the inventory.
 - The following tag is added to the checkpoint specified for the test: `Tested at startDateAndTimeOfTest`. The tagged checkpoint can be used to identify the point-in-time to restore the virtual machines in the VPG during a failover.

Testing that recovery is accomplished successfully should be done periodically so that you can verify that a failover will work. Zerto also recommends testing all the VPGs being recovered to the same cluster together. For example, in a cluster, if the HA configuration in a cluster includes admission control to prevent virtual machines being started if they violate availability constraints, testing the failover of every VPG configured for recovery to this cluster, at the same time, will show whether the constraints are violated or not.

When configuring a VPG, specify the period between tests for that VPG in the `Test Reminder` field in the `REPLICATION` step of the Create VPG wizard.

Starting a Failover Test

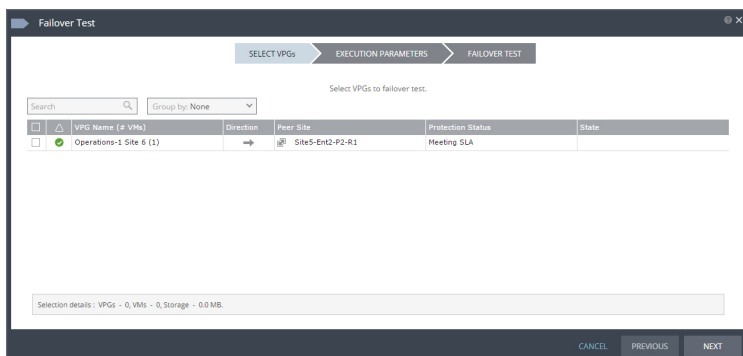
You can test a single VPG or multiple VPGs to make sure that if an actual failover is needed, the failover will perform as expected.

Note: You can initiate the failover test from either the protected site or recovery site.

To test failover:

1. In the Zerto User Interface set the operation to `TEST` and click `FAILOVER`.

The *Failover Test* wizard is displayed.



2. Select the VPGs to test. 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.



By default, the last checkpoint added to the journal is displayed. The checkpoints determine the RPO and ensure crash consistency and write-fidelity when the virtual machines in a VPG are recovered. These checkpoints are written every few seconds and you can recover to any of the available checkpoints.

4. Click `NEXT`.
5. To start the test, click `START FAILOVER TEST`.

The test starts for the selected VPGs. The test begins with an initialization period during which the virtual machines are created in the recovery site.

After Starting a Test, What Happens?

During the initiation phase, the virtual machines in the virtual protection group are created at the recovery site with the suffix *testing recovery*.


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 storage defined for the journal. Using these test scratch volumes makes cleaning up the test failover more efficient.

While a test is running:

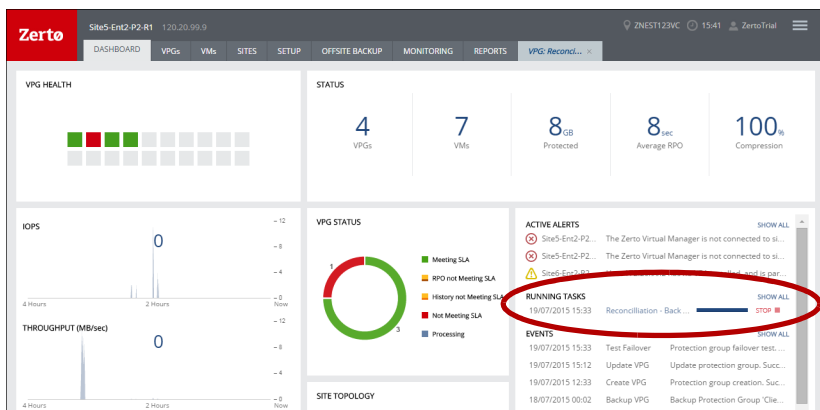
- The virtual machines in the VPGs continue to be protected.
- You can add checkpoints to the VPGs, and if necessary fail over the VPGs.
- You cannot take a snapshot of a test machine, since the virtual machine volumes are still managed by the VRA and not by the virtual machine. Using a snapshot of a test machine will create a corrupted virtual machine.
- You cannot move VPGs being tested.
- You cannot initiate a failover while a test is being initialized or closed.

Monitor the status of a failover test by doing the following:

- In the Zerto User Interface, click the *VPGs* tab. The *Operation* field in the *GENERAL* view displays *Failover test* when a failover test is being performed.



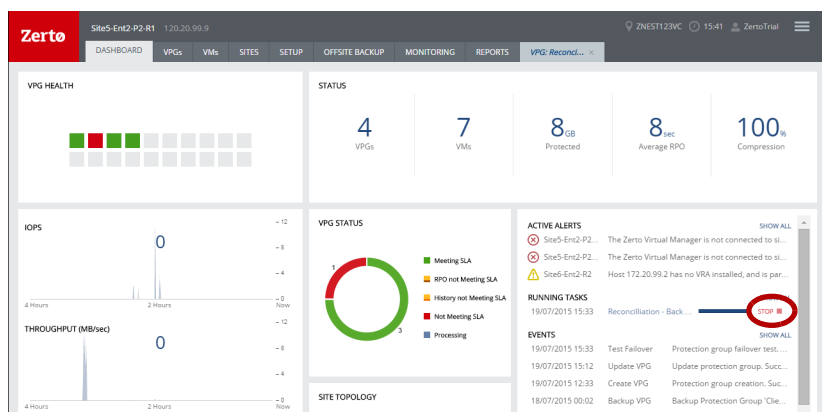
- In the Zerto User Interface, click the *VPGs* tab, and then click on the name of a VPG you are testing. A dynamic tab is created displaying the specific VPG details including the status of the failover test under *RUNNING TASKS*.




Stopping a Failover Test

To stop a failover test:

1. Click the *Stop test* icon to stop the test in the specific VPG tab.



You can also stop the test via the *TASKS* popup dialog in the status bar or under *MONITORING > TASKS*. The *Stop Test* dialog is displayed.



2. In the *Result* field specify whether the test succeeded or failed.
3. Optionally, in the *Notes* field, add a description of the test. For example, specify where external files that describe the tests performed are saved. Notes are limited to 255 characters.
4. Click *STOP*.

After stopping a test, the virtual machines in the recovery site are powered off and removed, and the checkpoint that was used for the test has the following tag added to identify the test: *Tested at startDateAndTimeOfTest*.

This checkpoint can be used to identify the point-in-time to use to restore the virtual machines in the VPG during a failover.

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.

www.zerto.com

For further assistance using Zerto Virtual Replication, contact Zerto support at support@zerto.com.