

## CREATING A VPG

You protect one or more virtual machines in a VPG. The VPG must include at least one virtual machine. After creating a VPG, you can add more virtual machines later.

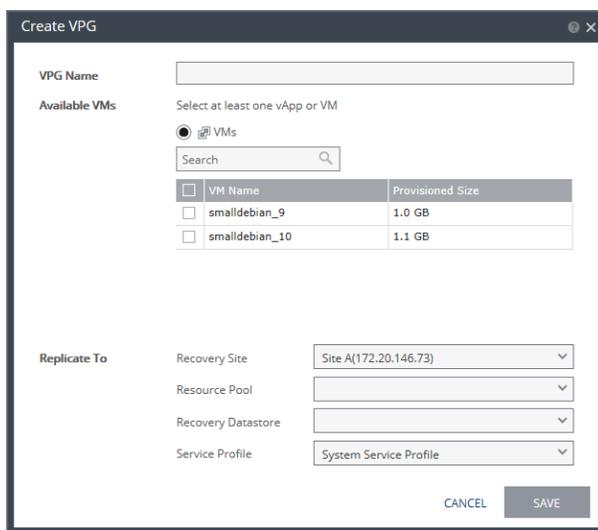
The VPG includes the virtual machines to recover. These virtual machines must be defined under a resource pool or be in a vApp.

Before you can create a VPG all the vCenter cloud resources must have been defined in Zerto Cloud Manager, such as the available networks.

To create a VPG:

1. In the Zerto Self-Service Portal ( <https://zerto.renovodata.com> ), click the *NEW VPG* button.

The *Create VPG* dialog is displayed.



VM Name	Provisioned Size
<input type="checkbox"/> smalldebian_9	1.0 GB
<input type="checkbox"/> smalldebian_10	1.1 GB

You select what you want to protect from the list, either one or more virtual machines or a vApp. If there are no vApps to be protected the vApp option is not displayed. If there are no virtual machines nor vApps to be protected a message is displayed to this effect.

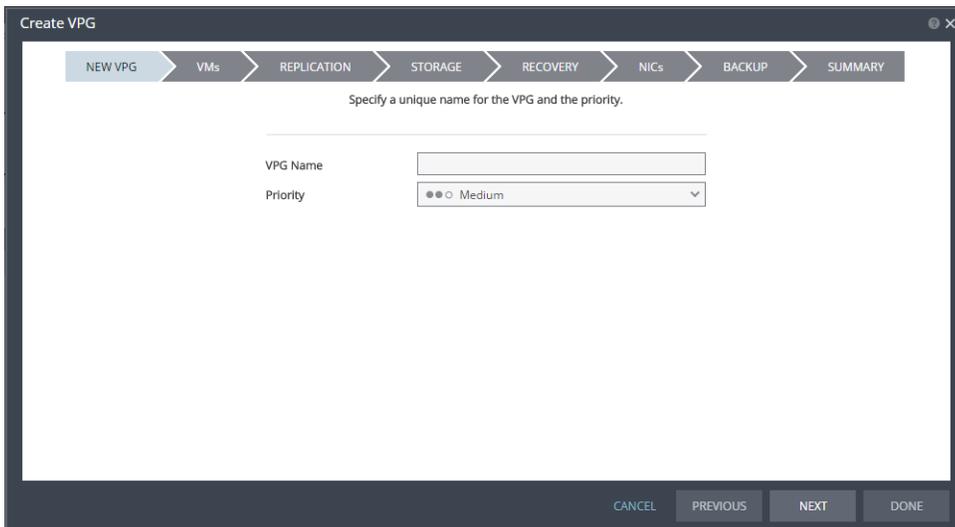
If the protected site is a vCD site, the vCD vApps are displayed:

Note: The virtual machines must be defined under a resource pool.

2. Select the virtual machines or vApp and specify the *Recovery Site*.
3. Specify the recovery resource pool, datastore and service profile.
4. If the specified service profile is not *Custom*, click *SAVE* to create the VPG, otherwise you can click *Advanced Configuration* to customize the VPG definition.

Note: Clicking *SAVE*, saves the VPG using the recovery datastore for all the virtual machines in the VPG. The failover/move and test failover networks are set to the network the recovery site uses. Even if you specify the *Custom* service profile, you can save the VPG using the default values.

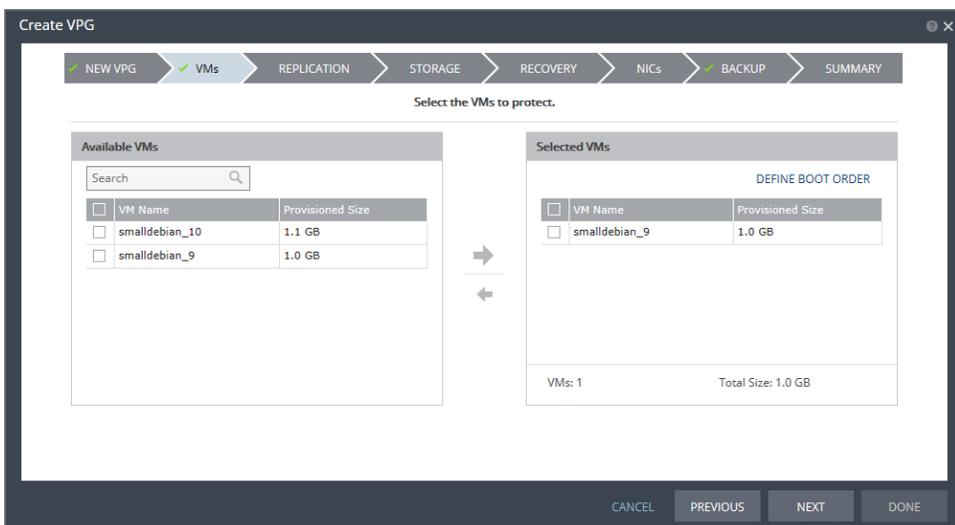
The *NEW VPG* step of the *Create VPG* wizard is displayed.



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. Note that updates to the protected virtual machines are always sent across the WAN before synchronization data, such as during a bitmap or delta sync. During a synchronization, only after updates to the virtual machines are sent over the WAN, based on the VPG priority, is synchronization data from the VPG sent, and the synchronization data from the VPG with the highest priority is passed over the WAN before data from medium and low priority VPGs.

5. Click *NEXT*.

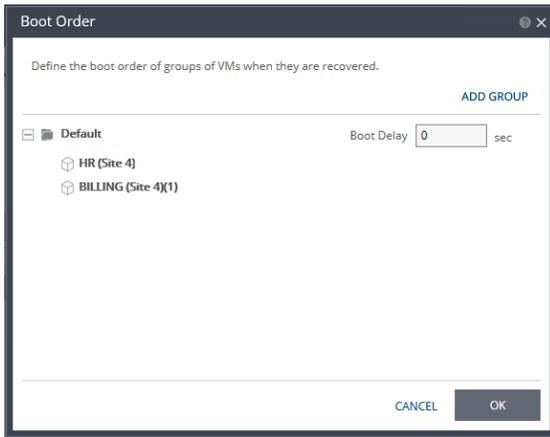
The *VMs* step is displayed.



6. Add VMs that will be part of this VPG and click the arrow pointing right to include these VMs in the VPG.
7. 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.

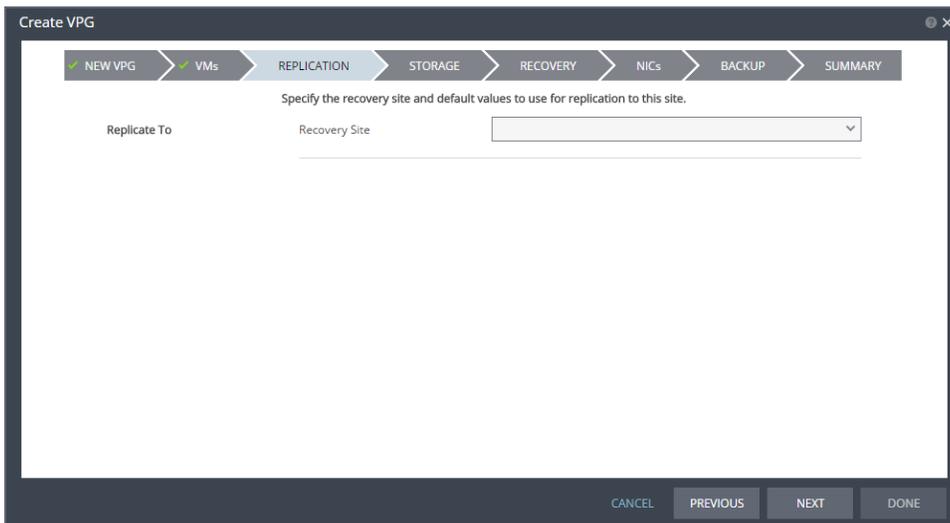
Initially, virtual machines in the VPG are displayed together under the *Default* group. If you want specific machines to start before other virtual machines, define new groups with one or more virtual machines in each group.



- a) Click *ADD* 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* to save the boot order.

8. Click *NEXT*.

The *REPLICATION* step is displayed.



9. Specify the recovery site.

**Recovery Site** – The site to which you want to recover the virtual machines. After specifying the recovery site, other fields are displayed including the host and datastore to use for replication.

**ZORG** – Specify the name the cloud service provider uses to identify you as a Zerto Organization, ZORG. For details about Zerto Cloud Manager, refer to *Zerto Cloud Manager Administration Guide*.

**Host** – The default resource pool 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.

**Datstore** – The datstore volume to use for all recovered virtual machine files as well as for their data volumes. Every datstore for the selected recovery host is included in the drop-down list. If a cluster or resource pool is selected for the host, only datstores 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 datstore in the cluster.

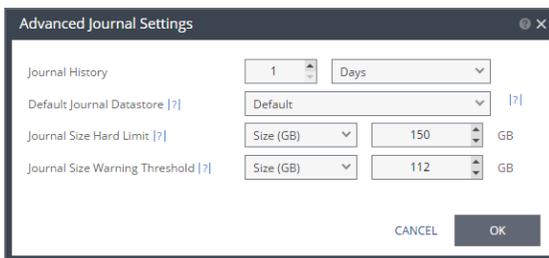
Optionally, change the default SLA values:

**Service Profile** – The name of the service profile to use which determines the VPG SLA settings for the group, which apply to every virtual machine in the group. To change the VPG SLA settings, select the *Custom Service Profile*.

**Default Journal History** – The time for which all write commands are saved in the journal. The longer the information is saved in the journal, the more space is required for each journal in the VPG.

For additional journal-related fields, click *ADVANCED*.

The *Advanced Journal Settings* dialog is displayed.



**Default Journal Datstore** – The datstore used for the journal data for each virtual machine in the VPG. Select a datstore accessible to the host. When you select a specific journal datstore, the journals for each virtual machine in the VPG are stored in this datstore, regardless of where the recovery datstores are for each virtual machine. In this case, all protected virtual machines must be recovered to the hosts that can access the specified journal datstore.

**Journal Size Hard Limit** – The maximum size that the journal can grow, either as a percentage or a fixed amount. The minimum journal size, set by Zerto Virtual Replication, is 8GB. The journal is always thin-provisioned.

**Unlimited** – The size of the journal is unlimited and it can grow to the size of the recovery datstore.

**Size (GB)** – The maximum journal size in GB.

**Percentage** – The percentage of the virtual machine volume size the journal can grow to.

**Journal Size Warning Threshold** – The size of the journal that triggers a warning that the journal is nearing its hard limit.

**Unlimited** – The size of the journal is unlimited and it can grow to the size of the recovery datstore.

**Size (GB)** – The size in GB that will generate a warning.

**Percentage** – The percentage of the virtual machine volume size that will generate a warning.

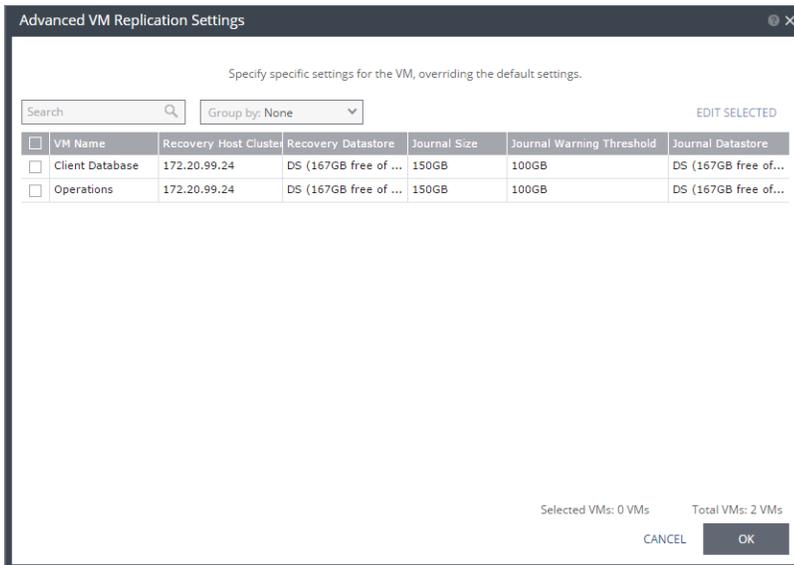
Both the value of Size and Percentage must be less than the configured hard limit so that the warning will be generated when needed. In addition to the warning threshold, Zerto Virtual Replication will issue a message when the free space available for the journal is almost full.

**Target RPO Alert** – The maximum desired time between each automatic checkpoint write to the journal before an alert is issued. To increase the value, move the slider right; to decrease the value, move the slider left.

**Test Reminder** – The time recommended between testing the integrity of the VPG. A warning is issued if a test is not done within this time frame.

10. If you want to change the replication settings per virtual machine, click *VM SETTINGS*.

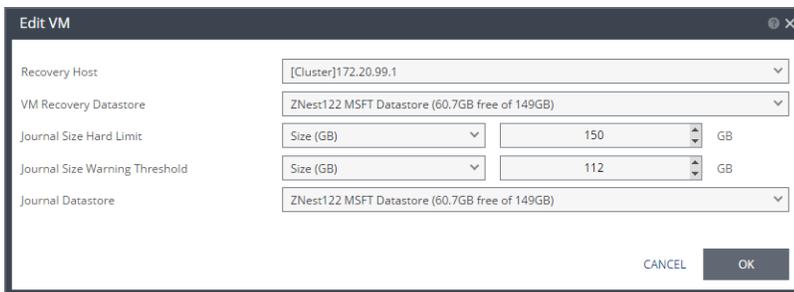
The *Advanced VM Replication Settings* dialog is displayed.



In this dialog, you can edit the values of one or more of the virtual machines in the VPG.

11. If you want to edit information in one field, click the field and update the information. If you want to edit information for several virtual machines at the same time, select the virtual machines and click *EDIT SELECTED*.

The *Edit VM* dialog is displayed.



**Recovery Host** – The cluster, resource pool, or host that will host the recovered virtual machine. If the site is defined in Zerto Cloud Manager, only a resource pool can be specified and the resource pool must also have been defined in Zerto Cloud Manager. For details about Zerto Cloud Manager, see *Zerto Cloud Manager Administration Guide*.

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

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 is to any one of the hosts in the recovery site with a VRA installed on it.

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 defined 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.

**Recovery Datastore** – The datastore where the VMware metadata files for the virtual machine are stored, such as the vmx file. If a cluster or resource pool is selected for the host, only datastores that are accessible by every ESX/ESXi host in the cluster or resource pool are displayed. This is also the datastore where RDM backing files for recovery volumes are located.

**Journal Size Hard Limit** – The maximum size that the journal can grow, either as a percentage or a fixed amount. The minimum journal size, set by Zerto Virtual Replication, is 8GB. The journal is always thin-provisioned.

Unlimited – The size of the journal is unlimited and it can grow to the size of the recovery datastore.

Size (GB) – The maximum journal size in GB.

Percentage – The percentage of the virtual machine volume size the journal can grow to.

**Journal Size Warning Threshold** – The size of the journal that triggers a warning that the journal is nearing its hard limit.

Unlimited – The size of the journal is unlimited and it can grow to the size of the recovery datastore.

Size (GB) – The size in GB that will generate a warning.

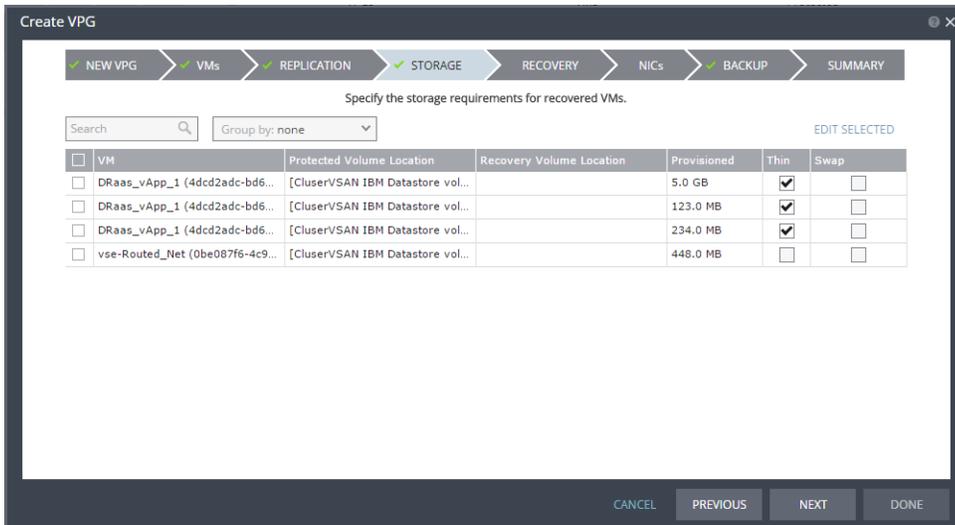
Percentage – The percentage of the virtual machine volume size that will generate a warning.

Both the value of Size and Percentage must be less than the configured hard limit so that the warning will be generated when needed. In addition to the warning threshold, Zerto Virtual Replication will issue a message when the free space available for the journal is almost full.

**Journal Datastore** – The datastore used for the journal data for each virtual machine in the VPG. To change the default, specify a host and then select one of the datastores accessible by this host to be used as the journal datastore. When you select specific journal datastore, the journals for each virtual machine in the VPG are stored in this datastore, regardless of where the recovery datastores are for each virtual machine. In this case, all the protected virtual machines must be recovered to hosts that can access the specified journal datastore.

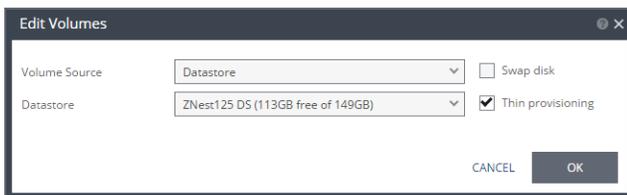
12. Click **SAVE**.
13. In the *Advanced VM Replication Settings* dialog, click **SAVE**.
14. Click **NEXT**.

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.

15. If you want to edit storage information for one of the virtual machines, select the machine and click **EDIT SELECTED**. The *Edit Volumes* dialog is displayed.



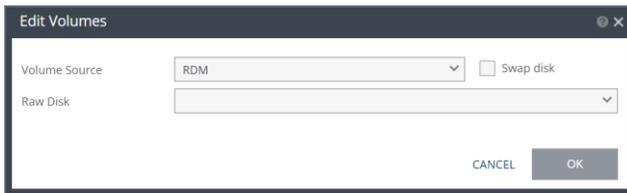
16. Specify the volume source for recovery from one of the following options.
17. Specify the volume source for recovery from one of the following options.

**Datastore** – A new volume is used for replicated data. Specify the datastore to use to create disks for the replicated data. If the source disk is thin provisioned, the default for the recovery volume is that it is also thin provisioned.

The datastore specified for replication must have at least the same amount of space as the protected volume and an additional amount for the journal. The amount of additional space needed for the journal can be fixed by specifying a maximum size for the journal, or can be calculated as the average change rate for the virtual machines in the VPG, multiplied by the length of time specified for the journal history.

Zerto Virtual Replication supports the SCSI protocol. Only disks that support this protocol can be specified.

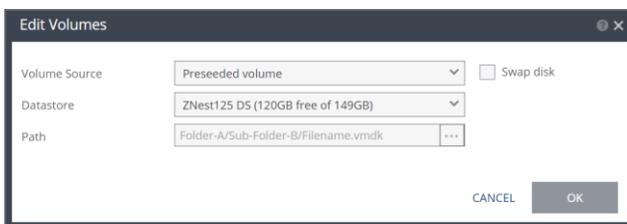
**Raw Disk** – The VMware RDM (Raw Device Mapping) to use for the replication: By default, RDM is recovered as thin-provisioned VMDK in the datastore specified in the *VM Recovery Datastore* field in the *Edit VM* dialog, and not to RDM.



You cannot define an RDM disk if the virtual machine uses a BusLogic SCSI controller, nor when protecting or recovering virtual machines in an environment running vCenter Server 5.x with ESX/ESXi version 4.1 hosts. Only a raw disk with the same size as the protected disk can be selected from the list of available raw disks. Other raw disks with different sizes are not available for selection. The RDM is always stored in the recovery datastore used for the virtual machine. The following limitations apply to protecting RDM disks:

- RDM disks with an even number of blocks can replicate to RDM disks of the same size with an even number of blocks and to VMDKs.
- RDM disks with an odd number of blocks can only replicate to RDM disks of the same size with an odd number of blocks and not to VMDKs.

**Preseeded volume** – Whether to copy the protected data to a virtual disk in the recovery site. Zerto recommends using this option particularly for large disks so that the initial synchronization will be faster since a *Delta Sync* can be used to synchronize any changes written to the recovery site after the creation of the preseeded disk. When not using a preseeded disk, the initial synchronization phase must copy the whole disk over the WAN. When using a preseeded virtual disk, you select the datastore and exact location, folder, and name of the preseeded disk, which cannot be an IDE disk. Zerto Virtual Replication takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. Only disks with the same size as the protected disk can be selected when browsing for a preseeded disk. The datastore where the preseeded disk is placed is also used as the recovery datastore for the replicated data.



If the preseeded disk is greater than 1TB on NFS storage, the VPG creation might fail. This is a known VMware problem when the NFS client does not wait for sufficient time for the NFS storage array to initialize the virtual disk after the RPC parameter of the NFS client times out. The timeout default value is 10 seconds. See the VMware documentation, [http://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=1027919](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1027919), which describes the configuration option to tune the RPC timeout parameter using the `esxconfig-advcfg -s <Timeout> /NFS/SetAttrRPCTimeout` command.

**Datastore** – The datastore where the preseeded disk is located.

**Path** – The full path to the preseeded disk.

Note the following conditions:

- If the protected disks are non-default geometry, configure the VPG using preseeded volumes.
- If the protected disk is an RDM disk, it can be used to preseed to a recovery VMDK disk. Zerto Virtual Replication makes sure that the VMDK disk size is a correct match for the RDM disk.
- If the VPG is being defined for a Zerto Organization, ZORG, the location of the preseeded disk must be defined in the Zerto Cloud Manager. For details, see *Zerto Cloud Manager Administration Guide*.

**Swap disk** – If the virtual machine to be replicated includes a swap disk as part of its configuration, mark the recovery disk for this disk as a swap disk. In this case, data is not replicated to the swap disk after initial synchronization.

**Datastore** – The datastore to use to create disks for the replicated data. Specify whether the target is thin provisioned. If the source disk is thin provisioned, the default for the recovery volume is that it is also thin provisioned.

The datastore specified for the replication must have at least the same amount of space as the protected volume and then an additional amount for the journal. The amount of additional space needed for the journal can be fixed by specifying a maximum size for the journal, or can be calculated as the average change rate for the virtual machines in the VPG, multiplied by the length of time specified for the journal history.

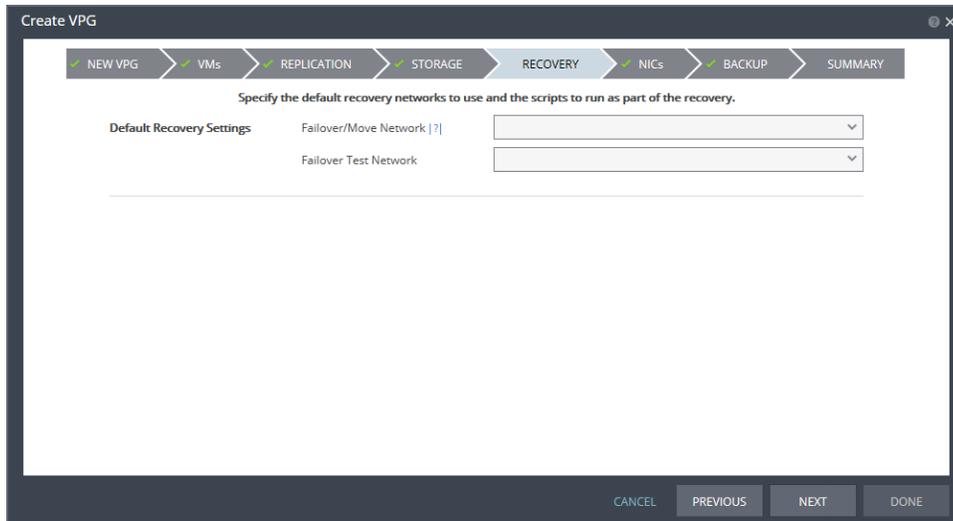
Note: You can use the vSphere Client console *Performance* tab for each virtual machine to help estimate the change rate.

**Thin Provisioning** – If the recovery volumes are thin-provisioned or not.

18. Click **SAVE**.

19. 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.



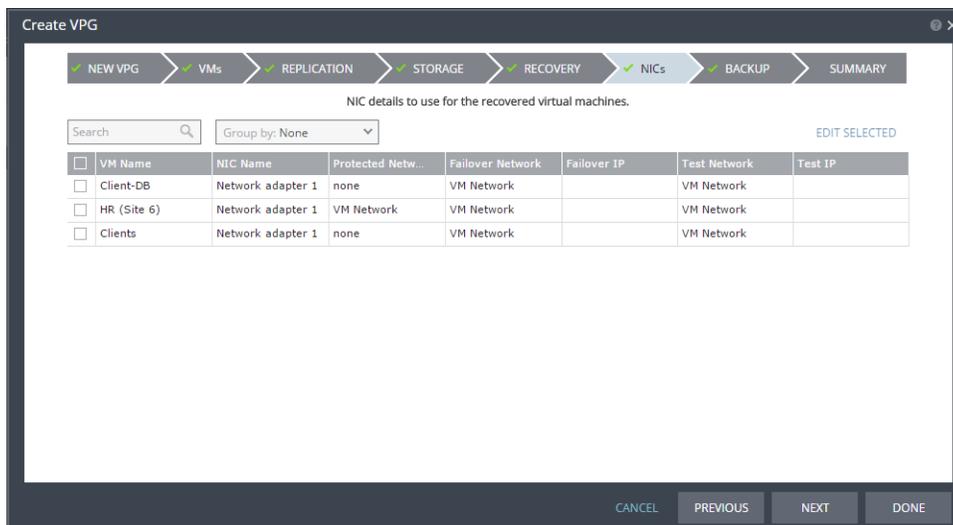
20. Select the default recovery settings. These are applied to every virtual machine in the VPG.

**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.

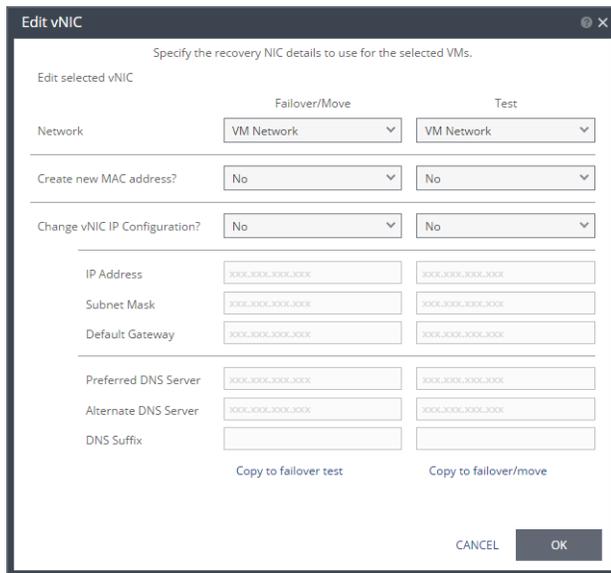
21. Click *NEXT*.

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.



22. If you want to edit information in one field, click the field and update the information. If you want to edit information for several virtual machines at the same time, select the virtual machines and click *EDIT SELECTED*. Otherwise, go to step Click .

The *Edit vNIC* dialog is displayed.



- Specify the network details to use for the recovered virtual machines after a failover or move operation, in the *Failover/Move* column, and for the recovered virtual machines when testing replication, in the *Test* column.

In each column, specify the following:

**Network:** The network to use for this virtual machine.

**Create New MAC Address:** Whether the Media Access Control address (MAC address) used on the protected site should be replicated on the recovery site. The default is to use the same MAC address on both sites. Note that if you check this option, to create a new MAC address, and the current IP address is not specified, the protected virtual machine static IP address might not be used for the recovered virtual machine.

**Change vNIC IP Configuration:** Whether or not to keep the default virtual NIC (vNIC) IP configuration. You can only change the vNIC IP for virtual machines with VMware Tools running for the following operating systems: Windows 2003 and higher, Red Hat Enterprise Linux versions 5-7.x, SUSE Linux Enterprise versions 10-11, Ubuntu 12.04, 12.10, 13.04, 13.10, 14.04, 14.10, 15.04, CentOS versions 5-7.x, and Oracle Linux versions 5.9-6.6.

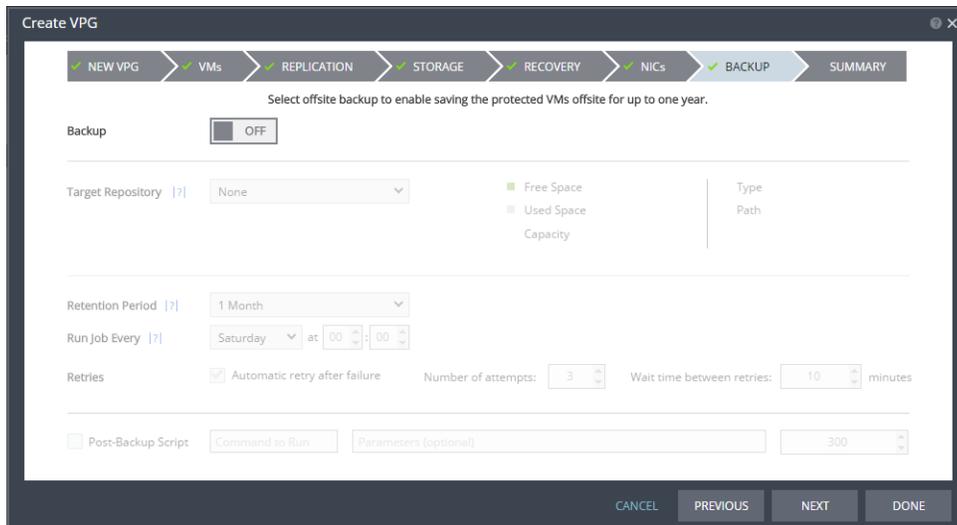
To change the vNIC IP, select *Yes* in the *Failover/Move* or *Test* column. If you select to use a static IP connection, set the IP address, subnet mask, and default gateway. Optionally, change the preferred and alternate DNS server IPs and the DNS suffix. If you leave the DNS server and suffix entries empty, or select to use DHCP, the IP configuration and DNS server configurations are assigned automatically, to match the protected virtual machine. You can change the DNS suffix.

If the virtual machine has multiple NICs but is configured to only have a single default gateway, fill in a 0 for each octet in the *Default gateway* field for the NICs with no default gateway.

**Note:** During a failover, move, or test failover, if the recovered virtual machine is assigned a different IP than the original IP, after the virtual machine has started it is automatically rebooted so that it starts up with the correct IP. If the same network is used for both production and test failovers, Zerto recommends changing the IP address for the virtual machines started for the test, so that there is no IP clash between the test machines and the production machines.

- Click *SAVE*.
- 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.



26. By default, backup is off. If you do not want to change this value, go to step Click . Otherwise, toggle OFF to ON and enter the following information:

**Target Repository** – The name of the repository where the offsite backups are written.

**Retention Period** – The length of time to keep offsite backups, up to a maximum of one year.

**Run Job Every** – The day and time to start the backup.

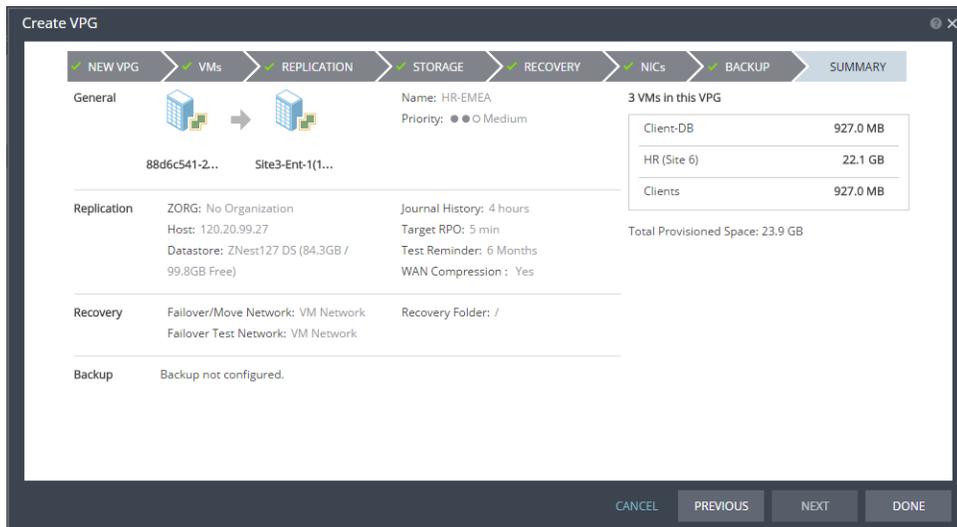
**Retries** – Whether to rerun the backup job automatically if the job fails. If you select this option, you must also define the number of retries that will be attempted and the time to wait after a job fails before running the backup job again.

**Post-Backup Script** – The information about a script that should run at the end of the recovery process. Enter the following information:

TEXT BOX	DESCRIPTION
Command to run	The full path of the script. The script must be located on the same machine as the Zerto Virtual Manager for the recovery site.
Parameters	The values of parameters to pass to the script. Separate parameters with a space.
Timeout	The time-out, in seconds, for the script to run. If the timeout value is reached, an alert is generated. The default time-out value is specified in the Performance and Throttling tab of the <i>Site Settings</i> dialog.

27. Click **NEXT**.

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



28. Click *DONE*.

The VPG is created.

The VPG is created. The protected virtual machines are synchronized with the replication virtual machines on the recovery site. This process takes some time, depending on the size of the VMs and the bandwidth between the sites. You can see the status in the *VPGs* tab or click the VPG link to open a tab specifically for the VPG.

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 and the delta changes to these virtual machines are sent to the recovery site.