Zerto

Zerto Virtual Manager Administration Guide

AWS Environment

Version 5.5 Update 3

ZVR-ADVA-5.5U3 Rev01 Dec2017

ZVR-ADVA-5.5U3 Rev01 Dec2017

Zerto

TABLE OF CONTENTS

ABOUT THIS GUIDE

Zerto Virtual Replication provides a business continuity (BC) and disaster recovery (DR) solution in a virtual environment, providing near real-time replication, with write-order fidelity, with minimal impact on product workloads. Fully automated orchestration delivers failover and failback in one click. Non-disruptive disaster recovery testing gives you confidence that your DR solution will work predictably and consistently. Protection groups ensure that all virtual machines that comprise an application are protected in the exact same manner no matter where they are in the environment.

With support for different hypervisors, such as vSphere or Hyper-V, workloads can be protected, migrated, and recovered, either within the same hypervisor environment or across hypervisor environments.

This guide describes how to configure and manage Zerto Virtual Replication to implement business continuity and disaster recovery (DR) solutions in a VMware, Hyper-V, AWS, or mixed environment.

Intended Audience

This guide is for the use of experienced hypervisor administrators.

Overview of Content in This Guide

This guide contains the following chapters:

CHAPTER	TITLE	DESCRIPTION
1	Introduction to Zerto Virtual Replication	Describes the underlying concepts and architecture of Zerto Virtual Replication.
2	Accessing the Zerto User Interface	Describes how to access the Zerto User Interface.
3	Initial Configuration for Offsite Backup	Describes how to set up offsite backups for extended recovery.
4	Overview of Recovery Flows	Describes disaster recovery and offsite backup flows from the initial protection to the recovery of virtual machines. It also describes, at a high level, the file level recovery process.
5	Introduction to Protecting Virtual Machines	Describes how to set up protection for virtual machines.
6	Monitoring Zerto Virtual Replication	Describes how to monitor the protected virtual machines and the protected and AWS sites.
7	Managing VPGs	Describes how to manage VPGs using Zerto Virtual Replication.
8	Managing a Zerto Virtual Manager	Describes the processes available to manage the Zerto Virtual Manager.
8	Advanced Site Configuration	Describes the processes available to manage protected and recovery sites using Zerto Virtual Replication.
10	Overview of Disaster Recovery Operations	Describes the available recovery procedures and when they are used.
11	Testing Recovery to AWS	Describes how to test recovery to ensure the results you want.
12	Migrating a VPG to AWS	Describes the process of migrating protected virtual machines from the protected site to the recovery site.
13	Managing Failover to AWS	Describes the process of recovery from the protected site to the recovery site.

CHAPTER	TITLE	DESCRIPTION
14	Cloning a VPG to AWS	Describes the process of cloning protected virtual machines from the protected site to the recovery site in AWS.
15	Recovering Files and Folders	Describes the process of restoring files and folders from the recovery site.
16	Failing Back from AWS	Describes how to fail back recovered virtual machines to VMware vSphere.
17	Zerto Virtual Replication Reports	Describes the reporting and monitoring capabilities available with Zerto Virtual Replication.
18	Troubleshooting	Describes how to resolve problems, including generating logs.
19	The Zerto Virtual Manager User Interface	Describes the screens and dialogs in the Zerto User Interface
20	Glossary	A glossary of terms used throughout Zerto Virtual Replication.

Support and Feedback

Please send suggestions to improve the documentation to Zerto support.

CHAPTER 1: INTRODUCTION TO ZERTO VIRTUAL REPLICATION

Disaster recovery is the process of preparing for recovery or continuation of IT processing tasks that support critical business processes in the event of a threat to the IT infrastructure. Zerto Offsite Backup is the additional process of enabling recovery of IT processing tasks after an extended period. This chapter describes Zerto Virtual Replication general concepts to enable replication and recovery in a virtual environment.

The following topics are described in this chapter:

What is Zerto Virtual Replication?

Zerto Virtual Replication provides a business continuity (BC) and disaster recovery (DR) solution in a virtual environment, providing near real-time replication, with write-order fidelity, with minimal impact on product workloads. Fully automated orchestration delivers failover, and failback in one click. Non-disruptive disaster recovery testing gives you confidence that your DR solution will work predictably and consistently. Consistency groups ensure all virtual machines that comprise an application are protected in the exact same manner no matter where they are in the environment.

With support for different hypervisors, such as vSphere or Hyper-V, workloads can be protected, migrated, and recovered, either within the same hypervisor environment or across hypervisor environments.

Zerto Virtual Replication is installed in both the protected and the recovery sites. The disaster recovery across these sites is managed by a browser-based user interface. Managing Zerto Virtual Replication is also possible programmatically, either via a set of RESTful APIs or PowerShell cmdlets.

Recovery that does rely on native replication functionality, such as recovery available with Microsoft Active Directory or SQL Server, can also be replicated using Zerto Virtual Replication, and whether the native replication functionality is used or not is determined by site considerations, such as increased complexity of having multiple points of control and possible additional costs incurred when using vendor native replication.

You configure replication by first pairing the site with virtual machines to be protected with a recovery site. You then define what virtual machines you want replicated in consistency groups, where the virtual machines in a group comprise the application and data you want to protect. You can group different virtual machines together or keep them separate. By creating different replication groups, you can customize the replication requirements for each group to better optimize the recovery plan.

Disaster recovery is based on the premise that you will want to recover with a minimum RPO. However, to enable full recovery in cases such as virus attacks, Zerto Virtual Replication provides the ability to recover to a point in time up to 30 days prior to the disaster. When recovery earlier than 30days is required, Zerto Virtual Replication provides an extended recovery, using an offsite backup mechanism that enables you to recover to a recovery site based on a daily or weekly backup going as far back as a year. The majority of the processing for both disaster recovery and extended recovery is done at the recovery site, minimizing the impact on the production site.

Zerto Virtual Replication Architecture

Zerto Virtual Replication provides disaster recovery between VMware ESX/ESXi hosts managed by vCenter Servers or Microsoft Hyper-V hosts managed by SCVMM to a public cloud, such as Amazon Web Services.

Zerto Virtual Replication comprises the following components:**Zerto Virtual Manager** (ZVM) – A Windows service that manages everything required for the replication between the protected site and AWS, except for the actual replication of data. Each Zerto Virtual Manager can manage up to 5000 virtual machines, either being protected or recovered to that site.

Virtual Replication Appliance (VRA) – A Windows service that manages the replication of data from protected virtual machines to AWS. A VRA can manage a maximum of 500 volumes.

Virtual Backup Appliance (VBA) – A Windows service that manages back-ups within Zerto Virtual Replication and is responsible for 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 by the Zerto User Interface in a web browser. The following diagram shows how the main Zerto Virtual Replication components are deployed to provide disaster recovery across these sites.[1]

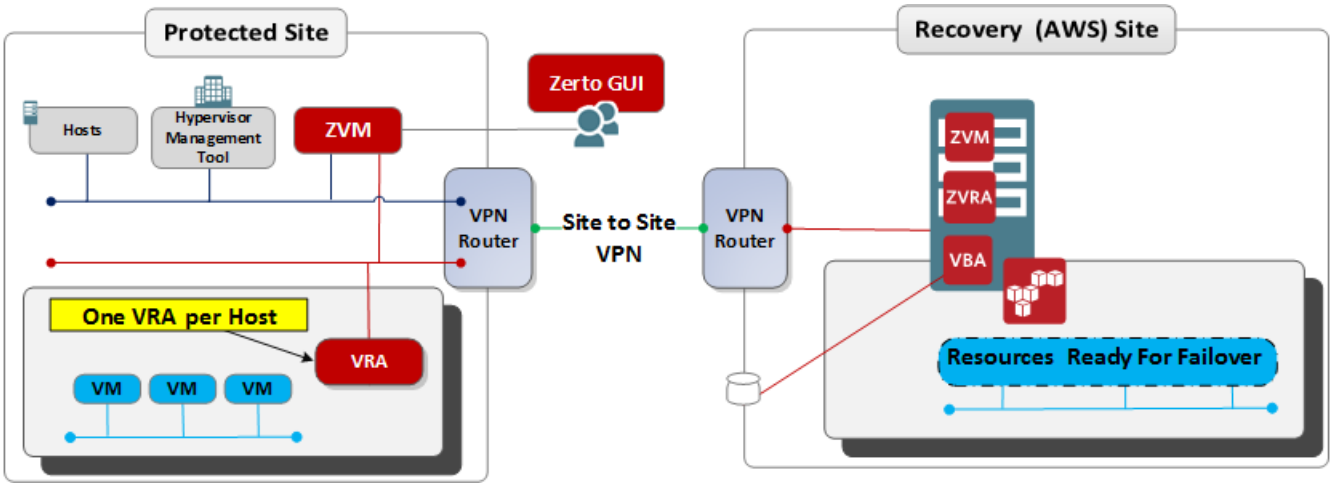

When you plan to recover the enterprise site to a public cloud, Zerto Virtual Replication is installed in the cloud environment. Zerto Virtual Replication comprises the same components but the VRA runs as a service, so that the ZVM, VRA, and VBA all run as services on a single virtual machine instance in the public cloud.

How Zerto Virtual Replication Recovery Works

Installing Zerto Virtual Replication installs the Zerto Virtual Manager, which sits in the hypervisor layer on the protected site and the Zerto Cloud Appliance which sits in AWS on the recovery site. You manage disaster recovery using the Zerto User Interface.

Zerto also provides a set of RESTful APIs and PowerShell cmdlets to enable incorporating some of the disaster recovery functionality within scripts or programs.

In the protected site you define the virtual machines that you want to replicate, either individually or together, as a virtual protection group (VPG). The virtual machines that you include in the VPG can come from one or more hypervisor hosts. In this way, you can protect applications that run on multiple virtual machines and disks as a single unit – a VPG. An example of an application that runs on multiple virtual machines includes software that requires a web server and database, both of which run on virtual machines different than the virtual machine where the application software runs.

A virtual machine can be included in several VPGs so that you can recover it to several sites, depending on the needs of the organization. For example the same workload can be protected to a local or a remote location as well as to the cloud. Using several recovery sites also enables migrating disaster recovery datacenters from one location to another.

1. For the architecture diagrams when one of the sites is a cloud service provider, see *Zerto Cloud Manager Administration Guide*.

Every write is copied by Zerto Virtual Replication and sent, asynchronously, to the recovery site, while the write continues to be processed on the protected site. For greater efficiency and performance, the write can be compressed before being sent to the recovery site with throttling techniques being used to prioritize network traffic.

On the recovery site the write is written to a journal managed by a Virtual Replication Appliance (VRA). Each protected virtual machine has its own journal. Every few seconds, a checkpoint is also written to each journal. These checkpoints ensure write order fidelity and crash-consistency to each checkpoint. During recovery you pick one of these crash-consistent checkpoints and recover to this point. Additionally, checkpoints can be manually added by the administrator, with a description of the checkpoint. For example, when an event is going to take place that might result in the need to perform a recovery, you can pinpoint when this event occurs as a checkpoint written to each journal.

The VRA manages the journals for every virtual machine that will be recovered to the hypervisor hosting that VRA. It also manages images of the protected volumes for these virtual machines. During a failover, you can specify that you want to recover the virtual machines in the VPG using the last checkpoint or you can specify an earlier checkpoint, in which case the recovery of the mirror images under the VRA are synchronized to this checkpoint. Thus, you can recover the environment to the point before any corruption and ignore later writes in the journal that were corrupted, regardless of the cause of the corruption, such as a crash in the protected site or a virus attack.

In AWS, users are not able to start working until all the information stored in the S3 buckets has been copied to the EBS disks attached to the new instances.

When recovery to a point is required that is further in the past than the time saved in the journal, an offsite backup can be restored. Offsite backups are an extension of disaster recovery, with the virtual machine files, such as the configuration and virtual disk files, saved to a repository for up to one year. These files are then used to restore the virtual machines to the point of the stored offsite backup at the recovery site.

Benefits of Using Zerto Virtual Replication

Datacenter optimization and virtualization technologies have matured and are now commonly used in IT infrastructure. As more applications are deployed in a virtualized infrastructure, there is a growing need for recovery mechanisms that support mission critical application deployments while providing complete BC and DR.

Traditional replication and disaster recovery solutions were not conceived to deal with the demands created by the virtualization paradigm. For example, most replication solutions are not managed in the hypervisor layer, considering the virtual machines and disks, but at the physical disk level. Hence they are not truly virtualization aware.

The lack of virtualization awareness creates a huge operational and administrative burden. It also results in operational inflexibility. Zerto Virtual Replication has been designed to resolve these issues by being fully virtualization aware.

See the following topics:

- "Hardware Agnostic", on page 12
- "Focus is on the Application, Not the Physical Storage", on page 12
- "Compatibility Across Virtual Environments – Cross-Hypervisor Platform and Version Agnostic", on page 12
- "Efficient Asynchronous Replication", on page 12
- "One-Click Failover and Control of the Recovery Process", on page 12
- "One-Click Migration", on page 12
- "File and Folder Recovery", on page 13
- "Offsite Backup", on page 13
- "Policy-based", on page 13
- "Minimal RPO", on page 13
- "WAN Optimization Between Protected and Recovery Sites", on page 13
- "WAN Resilience on Both the Protected and Recovery Sites", on page 13
- "DR Management Anywhere", on page 13

Hardware Agnostic

Because Zerto Virtual Replication software manages recovery of virtual machines and virtual disks only, it does not matter what hardware is used in either the protected or recovery sites; it can be from the same vendor or different vendors. With Zerto Virtual Replication the logical storage is separated from the physical storage so that the vendor and actual type of storage hardware do not need to be considered.

Zerto Virtual Replication provides a workload mobility and protection layer providing seamless connectivity, portability, protection, orchestration, and application encapsulation of workloads across clouds without vendor lock-in. High scale, mission critical applications, and data are encapsulated, as well as features, specifications, and configurations, and can be seamlessly migrated across different servers, storage, hypervisors, and clouds without any disruption to business services.

With Zerto Virtual Replication, IT managers can choose the right infrastructure for the right use case for the right price. One application can leverage several different environments for disaster recovery, bursting, production, backup, testing, and development. With Zerto Virtual Replication there is no vendor lock-in to a cloud, technology, or vendor. Any choice, any cloud, any technology, any price, any service level is available in minutes for any workload.

Focus is on the Application, Not the Physical Storage

By considering the physical disk level and not the virtual disk level, traditional replication is not truly application aware. Even virtual replication recovers block writes at the SCSI level and not at the application level. Zerto Virtual Replication is truly application focused, replicating the writes from the application in a consistent manner.

Compatibility Across Virtual Environments – Cross-Hypervisor Platform and Version Agnostic

Zerto Virtual Replication enables replication across multiple hypervisor managers, such as VMware vCenter Server and Microsoft SCVMM, and to public clouds such as Amazon Web Services (AWS) or Microsoft Azure. You can protect virtual machines in one hypervisor platform and recover to a different hypervisor platform. This feature can also be used to migrate virtual machines to a different hypervisor platform.

Also, virtual machines running in one version a hypervisor can be recovered in a different version of the same type of hypervisor, as long as Zerto Virtual Replication supports the hypervisor versions, virtual machines can be protected across versions.

Efficient Asynchronous Replication

Writes are captured by the Zerto Virtual Replication software in the hypervisor level, before they are written to the physical disk at the protected site. These writes are sent to the recovery site asynchronously, thus avoiding long distance replication latency for the production applications.

Also, because these writes are captured and sent to the recovery site, it is only the delta changes and not the whole file or disk that is sent to the recovery site, reducing the amount of network traffic, which reduces WAN requirements and significantly improves the ability to meet both RPO and RTO targets.

One-Click Failover and Control of the Recovery Process

When recovery is required, the administrator clicks on a button in the Zerto User Interface to initiate failover. This means that controlling the start of a recovery remains in the hands of the administrator, who can decide when to initiate the recovery and, by selecting a checkpoint, to what point-in-time to recover to.

One-Click Migration

Application migrations can be resource intensive projects that take weeks of planning, execution, and downtime. With Zerto Virtual Replication migrations are greatly simplified and can be completed without extended outages or maintenance windows and across different types of hardware and even different hypervisors, such as VMware ESXi or Microsoft Hyper-V. Migrations across different versions within a type of hypervisor, such as from a VMware vCenter environment to a vCloud environment or even cross hypervisor migration, such as migration from a vCenter environment to a Hyper-V environment is as easy as a migration from one site to another using the same hypervisor infrastructure.

File and Folder Recovery

You can recover specific files and folders from the recovery site for virtual machines that are being protected by Zerto Virtual Replication and running Windows or Linux operating systems. You can recover the files and folders from a specific point-in-time.

You can choose to recover one or several files or folders from the recovery site.

Offsite Backup

Zerto Virtual Replication provides an offsite back up option that enables saving the protected virtual machines offsite for up to one year in a state where they can be easily deployed. Because the backups use the same mechanism used for disaster recovery, there is no performance impact on the production site, since the processing is performed on the recovery site. The offsite backups are fixed points saved daily, weekly or monthly.

Note: Zerto recommends **weekly** backups.

Policy-based

In the protected site you define the virtual machines that you want to recover, either individually or as groups, as a virtual protection group (VPG). The virtual machines that you include in the VPG can come from one or more hypervisor hosts. In this way, you can protect applications that run on multiple virtual machines and disks as a single unit, in a single VPG.

Minimal RPO

Zerto Virtual Replication utilizes continuous data protection, sending a record of every write in the virtual protection group to the recovery site. The transfer of this information is done over an optimized WAN asynchronously. If recovery is required, all the data that was transferred to the recovery site is available resulting in recovery within the requested RPO.

WAN Optimization Between Protected and Recovery Sites

Using compression to minimize bandwidth and other techniques such as throttling to prioritize network traffic to reduce the impact on day-to-day operations, you can make sure that the communication between the protected and recovery sites is fully optimized.

Zerto Virtual Replication also uses signature matching to reduce the amount of data sent across the WAN. During synchronization of the protected site and recovery site for every virtual machine in a VPG, Zerto Virtual Replication maintains a map of disk sectors so that if there is a need to resynchronize sites, the map signatures can be used to ensure that only data where changes occurred are passed over the WAN.

WAN Resilience on Both the Protected and Recovery Sites

Zerto Virtual Replication is highly resilient to WAN interruptions. In order to reduce storage overhead used for replication purposes, on WAN failure or when the load over the WAN is too great for the WAN to handle, 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. The bitmap is small and scales dynamically, but does not contain any actual IO data, just references to the areas of the protected disk that have changed. The bitmap is stored locally on the VRA within the available resources. Once the WAN connection resumes or the load returns to normal traffic, Zerto Virtual Replication uses this bitmap to check whether there were updates to the protected disks and if there were updates to the disks, these updates are sent to the recovery site.

DR Management Anywhere

With Zerto Virtual Replication everything is managed from a standalone browser-base user interface, enabling disaster recovery management from anywhere using any device.

Import Methods for AWS

During recovery operations, Zerto uses a combination of the following APIs and methods to convert the Amazon S3 objects into recovery disks in EC2 as EBS disks:

- **AWS Import:**
  - **Import-instance:** for the boot volume
  - **Import-volume:** for data volumes

  For more information see the relevant AWS documentation:
  - API_ImportInstance
  - API_ImportVolume
- **Zerto Import - zImport,** an import method that does not have the same limitations as the AWS APIs. It creates an AWS EC2 instance per protected VM volume, called zImporter, to convert the S3 objects and write them to a zImport local disk. When all the data has been imported and its disk have been attached to the recovered instance, the zImport instance is terminated.

  Notes:
  - zImporter is based on an official AWS Linux AMI (Amazon Machine Image), into which a script is injected to perform the import. The script is located online and downloaded to the zImporter, and thus the zImporter requires internet access in order to access and download the script. The zImport instance is therefore created with a public IP.
  - The only network in the customer environment that is certain to have internet access is the network that the ZCA is connected to.
  - To ensure that the zImport instance cannot be accessed from the outside world, a security group is created. During a recovery operation the zImport instance is connected to this security group. All inbound traffic is blocked and only outbound traffic to access the script online is allowed. The security group is deleted at the end of the recovery operation.
  - The default zImporter instance type is c4.8xlarge and the AWS EC2 default maximum instance quota is 20. If during the creation of zImport instances the maximum EC2 instance quota is reached, the creation of the next and subsequent zImport instances will be queued, increasing the RTO. If during recovery operations, the ZVM identifies a VPG with the potential to exceed the EC2 instance quota, the user will receive an alert with advice to contact AWS support to increase the service limits in order to improve RTO.
  - GPT formatted disks are supported for data volumes only, when using either of the zImport methods.
  - When using either of the zImport methods, each volume is created with EBS disk of type io1 with maximum 1000 EBS Provision IOPS allocated. EBS disk type can be changed post recovery without downtime, see the relevant for more information see the relevant AWS documentation. The minimum disk size for io1 is 4GB.
  - The default Max EBS Provision IOPS quota in a region across all io1 disks is 40000 EBS Provision IOPS, meaning that with 1000 EBS Provision IOPS per volume, the maximum possible number of volumes is 40. If the Max EBS Provision IOPS quota is reached, the failover process will switch to using slower gp2 disks. An event will notify the user of this, and recommend that the user contact AWS support to increase the Max EBS Provision IOPS quota.
- Depending on the desired RTO during recovery operations, or when testing failover, the user can select an import method per VPG or per virtual machine from the following options:
  - "Zerto Import for Data Volumes", on page 14
  - "Zerto Import for All Volumes", on page 15
  - "AWS Import", on page 16

Zerto Import for Data Volumes

This method is the **default setting** and has a faster RTO than AWS Import. This method uses a **combination** of the **AWS import-instance** API for the boot volume, and the **zImport** method for data volumes.

- **Each machine that you intend to protect** must have at least **250MB free space.** This is because AWS adds files to the recovered machines during failover, move, test failover, and clone operations.
- **Protected boot volumes** are recovered in EC2 as EBS disks with magnetic disk type. Virtual machines with disks that are less than 1GB are recovered with disks of 1GB. Temporary disks may be created based on the selected instance size.
- Temporary disks may be created based on the selected instance size.

- The **maximum** protected **data volume** size is **16TB**, while the **boot volume** can be up to **1TB**.
- The AWS ImportInstance API only supports single volume VMs. The boot volume of the protected virtual machine should not be attached to any other volume to successfully boot. For more information, see http://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_ImportInstance.html

Zerto Import for All Volumes

This method uses the **zImport** method for all volumes and ensures the fastest RTO.

This method creates an AWS EC2 instance per protected VM volume, called zImporter, to convert the S3 objects and write them to a zImport local disk. When all the data has been imported and its disk have been attached to the recovered instance, the zImport instance is terminated.

- Temporary disks may be created based on the selected instance size.
- This import method cannot be used for the recovery of protected virtual machines running Windows Server 2008R2. To support this operating system please contact Zerto support.
- The **maximum** protected **data volume** size is **16TB**, while the **boot volume** can be up to **2047GiB**.

  Note: Some VMs use the MBR partitioning scheme, which only supports up to 2047 GiB boot volumes. If your instance does not boot with a boot volume that is 2TB or larger, the VM you are using may be limited to a 2047 GiB boot volume size. Non-boot volumes do not have this limitation. See AWS Documentation for more information: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

When using this import method, if the protected virtual machine using this import method is running **Windows 2012**, **Windows 2012R2** or **Windows 2016**, the following drivers **must** be installed on the **protected** virtual machine **before starting recovery operations**:

- Windows PV (Paravirtualization) Drivers
- Windows ENA (Elastic Network Adapter) Drivers

The following steps **must be performed** to ensure that the virtual machine will be able to run on the recovery site:

1. **Download and Install Windows PV Drivers:**
   a) Go to https://www.xenproject.org/downloads/windows-pv-drivers/winpv-drivers-81/winpv-drivers-820.html
   b) Follow the instructions at the site for downloading and installing **all** the Windows PV Drivers 8.2.0 drivers.
2. **Download and Install Windows ENA Drivers:**

   If you are running Windows 2012 or Windows 2016 on any of the following AWS instance types:

   - C3
   - C4
   - D2
   - I2
   - R3
   - M4 (excluding M4.16xlarge)

   a) Go to http://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/sriov-networking.html#enable-enhanced-networking
   b) Follow the instructions at the site for downloading and installing the Windows ENA Drivers.

Note: If these drivers are installed on a VM running Windows 2012R2, the other AWS import methods will fail. To overcome this, you must uninstall the drivers before using the other AWS import methods.

AWS Import

This method uses a combination of the **AWS import-instance** and **import-volume** APIs for the boot and data volumes respectively. This was the only method supported until version 5.5.

- **Each machine that you intend to protect** must have at least **250MB free space.** This is because AWS adds files to the recovered machines during failover, move, test failover, and clone operations.
- **Protected boot volumes** are recovered in EC2 as EBS disks with magnetic disk type. Virtual machines with disks that are less than 1GB are recovered with disks of 1GB. Additional volumes might be created in the recovered instance, dependent on the instance type used for the recovery. These volumes can be ignored.
- **Protected volumes** are recovered in EC2 as EBS disks with magnetic disk type. Virtual machines with disks that are less than 1GB are recovered with disks of 1GB. Additional volumes might be created in the recovered instance, dependent on the instance type used for the recovery. These volumes can be ignored. Temporary disks may be created based on the selected instance size.
- The **maximum** protected **data volume** and **boot disk** size is **1TB**.
- The AWS ImportInstance API only supports single volume VMs. The boot volume of the protected virtual machine should not be attached to any other volume to successfully boot. For more information, see http://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_ImportInstance.html

CHAPTER 2: ACCESSING THE ZERTO USER INTERFACE

You manage the protection and replication of virtual machines between the protected site and AWS using the Zerto User Interface. On first access to the Zerto User Interface, you might have to add a security certificate to set up secure communication. Zerto also provides a set of RESTful APIs and PowerShell cmdlets to enable incorporating some of the disaster recovery functionality within scripts or programs.

Note: 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.

Note: 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.

The following topics are described in this chapter:

Using the Zerto Virtual Manager Web Client

1. 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.
2. Log in using the user name and password of the instance on AWS on which you installed the Zerto Cloud Appliance.

Adding a Security Certificate for the Zerto User Interface

Communication between the Zerto Virtual Manager and the user interface uses HTTPS. On the first login to the Zerto User Interface, you must install a security certificate in order to be able to continue working without each login requiring acceptance of the security.

To install a security certificate for the Zerto User Interface:

On first access to the Zerto User Interface, if you haven't installed the security certificate, a security alert is issued.

Note the following:

- To run this procedure run Microsoft Internet Explorer as administrator. The procedure is similar for Google Chrome and for Mozilla Firefox.
- Access the Zerto User Interface using the IP and not the name of the machine where Zerto Virtual Replication is installed.

1. Click *View Certificate*.

   The *Certificate* dialog is displayed.
2. Click *Install Certificate*.

   The *Certificate Import wizard* dialog is displayed.
3. Follow the wizard: Place all the certificates in the `Trusted Root Certification Authorities` store: Select the `Place all certificates in the following store` option and browse to select the `Trusted Root Certification Authorities` store.

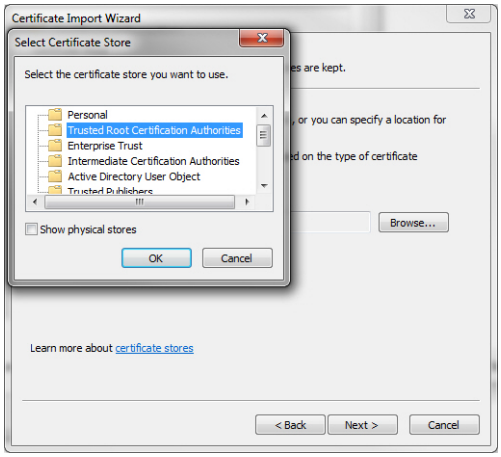

4. Continue to the end of the wizard. Click *Yes* when the Security Warning is displayed.

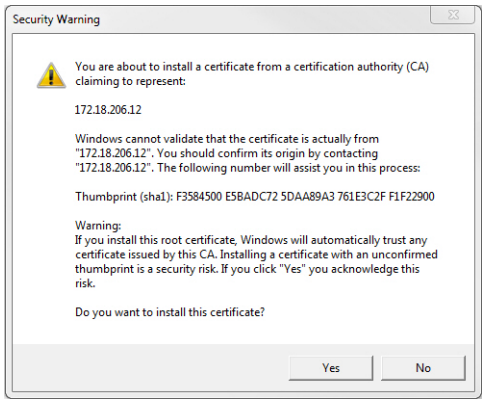

5. Click *OK* that the installation was successful.
6. Click *OK* when prompted and then *Yes* in the *Security Alert* dialog to continue.

Working With the Zerto User Interface

After logging on to the Zerto User Interface for the first time, the dashboard is displayed. The dashboard provides summary information about the status of the site, as shown in the following diagram:

Use the tabs to access the specific information you want:

DASHBOARD – General information about the site, including the status of the VPGs being protected or recovered to the site.

VPGs – All the VPGs from both the local and remote sites and provides summary details of each VPG.

VMs – All the protected virtual machines from both the local and remote sites and provides summary details of each virtual machine.

SITES – Details of the paired sites. This tab lists all the paired sites to the local site and provides summary details of each paired site.

SETUP – Details about VRAs, storage and repositories.

OFFSITE BACKUP – Details of the offsite backup jobs either by VPG or virtual machine. This tab lists all the defined offsite backups and their statuses.

MONITORING – Details about the alerts, events and tasks for the site.

REPORTS – General reports.

Subtabs

The SETUP, OFFSITE BACKUP and MONITORING tabs and details of a specific VPG and VRA can be viewed from different perspectives via subtabs. For example, under SETUP you can manage VRAs, storage and repositories via subtabs.

Views

Lists can be displayed with different views. For each view you can filter the information in columns via the filter icon next to each column title. Clicking the column title enables sorting the column in ascending to descending order.

You can customize the default views or add a new view by clicking the view configuration button.

Customize a default view by selecting `Show/Hide Columns` and then checking the columns you want displayed. Create a new view by selecting `Create View`.

CHAPTER 3: OVERVIEW OF RECOVERY FLOWS

Zerto Virtual Replication enables protecting virtual machines, for both disaster recovery or for extended, longer term recovery from an offsite backup, by protecting the relevant virtual machines in virtual protection groups. A virtual protection group (VPG) is a group comprised of virtual machines that are grouped together for recovery 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 for the Web Server require that all three virtual machines be replicated to maintain data integrity.

The following topics are described in this chapter:

Once a VPG has been created, each virtual machine in the VPG can be replicated on the recovery site under the VRA on the host specified in the VPG definition as the host for the recovery of the virtual machine.

Every write to the protected virtual machine in a VPG is copied by the VRA on the same host as the protected machine and passed to the VRA in AWS. The VRA writes these transactions to buckets in S3. This information is ready to be written to EBS disks in EC2 when recovering from a disaster or when recovering an offsite backup.

Flow for a Disaster Recovery Operation

Disaster recovery using Zerto Virtual Replication enables recovering from a disaster to any point between the moment just before the disaster and a specified amount of time in the past up to 30 days. The recovery is done in real time at the recovery site with a minimal RTO.

A recovery operation is one of the following:

- A failover.
- A planned move of the protected virtual machines from the protected site to the recovery site.
- A clone of the protected virtual machine to the recovery site.

Virtual machines are protected in VPGs, which are defined in the protected site. Once a VPG is created, Zerto Virtual Replication creates a copy of the protected virtual machines under the management of a Virtual Replication Appliance, VRA, on the AWS recovery site. The data managed by the VRA is saved in an S3 bucket.

When a recovery operation is performed, the VRA creates the virtual machines defined in the VPG in EC2 and imports data from the S3 bucket as EBS disks in EC2.

After initializing the VPG, all writes to the protected virtual machines are sent by the VRA on the relevant host for each virtual machine on the protected site to the VRA on the AWS recovery site. The information is saved in the journal for the virtual machine with a timestamp, ensuring write-fidelity. Every few seconds the Zerto Virtual Manager causes a checkpoint to be written to every journal on the recovery site for every virtual machine in the VPG, ensuring crash-consistency.

The data remains in the journal until the time specified for the journal when it is moved to the relevant mirror disks in S3, also managed by the VRA for the virtual machine. In this way, you can recover the virtual machines using the mirror disks and the data from the journal to include the final few hours of data for each virtual machine. Refer to "The Role of the Journal During Protection", on page 24 for more details about the journal.

The following references the procedures to recover virtual machines protected in a VPG:

- "Overview of Disaster Recovery Operations", on page 73
- "Managing Failover to AWS", on page 97
- "Migrating a VPG to AWS", on page 91
- "Cloning a VPG to AWS", on page 103

Flow for a Test Failover Operation

When testing that the recovery works as planned, the Zerto Virtual Manager (ZVM) creates the virtual machines defined in the VPG in EC2. The following references the procedure to recover virtual machines:

- "Overview of Disaster Recovery Operations", on page 73
- "Testing Recovery to AWS", on page 83

Flow for a File or Folder Level Restore Operation

You can recover specific files and folders from the recovery site for virtual machines that are being protected by Zerto Virtual Replication and running Windows operating systems. You can recover the files and folders from a specific point-in-time.

To recover files and folders, see "Recovering Files and Folders", on page 107.

Flow for an Offsite Backup and Restore Operation

If there is a requirement to extend the recovery ability to more than the 30 days that are available with disaster recovery, Zerto Virtual Replication provides an offsite backup option that enables saving the protected virtual machines offsite for up to one year in a state where they can easily be deployed. The recovery virtual machines are saved in a repository of offsite backups that can extend as far back as a year. These offsite backups are fixed points saved either daily or weekly. To save space the offsite backups can be compressed before they are stored in the repository.

When an offsite backup job starts, the Virtual Backup Appliance (VBA) on the recovery site communicates with the VRA on the recovery site to create the backup files of the virtual machines in the VPG, including the data in the journal and saves these files in the repository.

To set up repositories to protect virtual machines in a VPG with offsite backup, see "Offsite Backup Configuration", on page 123.

Setting up offsite backups is part of defining a VPG.

After initializing the VPG, Zerto Virtual Replication periodically checks that the time to run an offsite backup has not passed. At the scheduled backup time, the offsite backup is run and the offsite backup file stored in the specified repository.

Offsite backups are kept for the retention period specified in the VPG. Over time the number of stored offsite backups is reduced to save space.

Note: You cannot restore a backup in AWS.

CHAPTER 4: INTRODUCTION TO PROTECTING VIRTUAL MACHINES

Virtual machines are protected in virtual protection groups (VPGs). A VPG is a group of virtual machines that you group together for recovery 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 for the Web Server require that all three virtual machines be replicated to maintain data integrity.

Once a virtual machine is protected, all changes made on the machine are sent to the remote site. The remote site can be recovered to any point in time defined for the VPG or if a period further in the past is required, an offsite backup can be restored.

Any virtual machine whose operating system is supported in both the protected site and AWS can be protected in a VPG.

When a VPG is created, application data and the data required to recreate the protected virtual machines are copied to AWS during a process of synchronization. This synchronization between the protected site and AWS takes time, depending on the size of the virtual machines.

After the initial synchronization completes, only the writes to disk from the virtual machines in the protected site are sent to AWS. These writes are stored by the Virtual Replication Appliance (VRA) in the journals in an S3 bucket for a specified period, after which they are promoted to the replica virtual disks managed by the VRA, which are also in an S3 bucket.

The number of VPGs that can be defined on a site is limited only by the number of virtual machines that can be protected. Each site can manage a maximum of 5000 virtual machines.

Note: If the total number of protected virtual machines on the paired sites is 5000, then any additional machines are not protected.

Any virtual machine that is supported by the protected site hypervisor can be protected. The protected virtual machines must also be supported by the recovery AWS site.

The following topics are described in this chapter:

Configuring Virtual Protection Groups

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 or remove virtual machines as required. You can only protect a virtual machine in a VPG when the virtual machine has no more than 60 disks.

VPGs must be created on the protected site; they cannot be created on AWS, the recovery site. The virtual machines on the protected site can be defined under a single hypervisor host or under multiple hosts.

To create a VPG that will be recovered to AWS, you must have a virtual instance in AWS with a Zerto Cloud Appliance installed on it. This virtual instance must be paired with the protected site.

The VPG definition consists of the following:

- **General:** A name to identify the VPG and the priority to assign to the VPG.
- **Virtual machines:** The list of virtual machines being protected as well as the boot order and boot delay to apply to the virtual protection groups during recovery.
- **Replication:** The recovery site settings and the VPG SLA. SLA information includes the default journal history settings and how often tests should be performed on the VPG. These settings are applied to every virtual machine in the VPG but can be overridden per virtual machine, as required.
- **Storage:** The default storage volume to use for the recovered virtual machine files and for their data volumes. If a cluster is selected for the host, only storage accessible.
- **Recovery:** The networks, subnets, security groups, instance families, and instances types to use for failover/move and failover test procedures and the scripts, if any, that should run at the start or end of a recovery operation.
- **Backup:** The properties that govern the VPG backup including the repository where the backups are saved.
- **Summary:** The details of the VPG configuration defined in the previous components.

Requirements for AWS Environments

- Only virtual machines that are supported by AWS can be protected by Zerto Virtual Replication. Refer to AWS documentation for the supported operating systems.
- **A VPC must exist,** and a security group and subnet must be assigned to it and to all other VPCs you want to use for recovered virtual machines.
- The following **limitations** apply when protecting to AWS:
  - For **Linux**, AWS supports virtual machines with up to **12 volumes**, including the boot volume.
  - For **Windows**, AWS supports virtual machines with up to **22 volumes**, including the boot volume.
  - GBT formatted disks are supported for data volumes only.
  - The following table describes the limitations per Import Method:

IMPORT METHOD						
	AWS Import		**zImport for Data Volumes**		**zImport for all volumes**	
OS	Boot Volume	Additional Volume	Boot Volume	Additional Volume	Boot Volumes	Additional Volumes
Linux	1 TB	1 TB	1 TB	16 TB	2047 GiB*	16 TB
Windows	1 TB	1 TB	1 TB	16 TB	2047 GiB*	16 TB

* Some VMs use the MBR partitioning scheme, which only supports up to **2047 GiB** boot volumes. If your instance does not boot with a boot volume that is 2 TB or larger,, the VM you are using may be limited to a 2047 GiB boot volume.
See the relevant AWS documentation for more information: http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

  - For the **AWS Import** and **zImport for Data Volumes** import methods, the **AWS ImportInstance API** only supports single volume VMs. The boot volume of the protected virtual machine should not be attached to any other volume to successfully boot. For more information, see http://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_ImportInstance.html

See also: "Import Methods for AWS", on page 14.

The Role of the Journal During Protection

After defining a VPG, the protected virtual machine disks are synced with the recovery site. After initial synchronization, every write to a protected virtual machine is copied by Zerto Virtual Replication to the recovery site. The write continues to be processed normally on the protected site and the copy is sent asynchronously to the recovery site and written to a journal in a bucket in S3 managed by a Virtual Replication Appliance (VRA). Each protected virtual machine has its own journal.

In addition to the writes, every few seconds all journals are updated with a checkpoint time-stamp. Checkpoints are used to ensure write order fidelity and crash-consistency. A recovery can be done to the last checkpoint or to a user-selected, crash-consistent checkpoint. This enables recovering the virtual machines, either to the last crash-consistent point-in-time or, for example, when the virtual machine is attacked by a virus, to a point-in-time before the virus attack.

Data and checkpoints are written to the journal until the specified journal history size is reached, which is the optimum situation. At this point, as new writes and checkpoints are written to a journal, the older writes are written to the virtual machine recovery virtual disks. When specifying a checkpoint to recover to, the checkpoint must still be in the journal. For example, if the value specified is 24 hours then recovery can be specified to any checkpoint up to 24 hours. After the time specified, the mirror virtual disk volumes maintained by the VRA are updated.

During recovery, the virtual machines at the recovery site are created and the recovery disks for each instance, managed by the VRA, are attached to the recovered virtual machines. Information in the journal is promoted to the virtual instances to bring them up to the date and time of the selected checkpoint.

Each protected virtual machine has its own dedicated journal.

What Happens After the VPG is Defined

After defining a VPG, the VPG is created. The VRA in AWS is updated with information about the VPG and its protected virtual machines. Until a recovery operation is performed, all data managed by the VRA is stored in an S3 bucket.

The synchronization process can take some time, depending on the size of the virtual machines, the amount of data in its volumes, and the bandwidth between the sites. During this synchronization, you cannot perform any replication task, such as adding a checkpoint.

For synchronization to work, the protected virtual machines must be powered on. The VRA requires an active IO stack to access the virtual machine data to be synchronized across the sites. If the virtual machine is not powered on, there is no IO stack to use to access the protected data to replicate to the target recovery disks, and an alert is issued.

Once synchronized, the VRA in AWS 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.

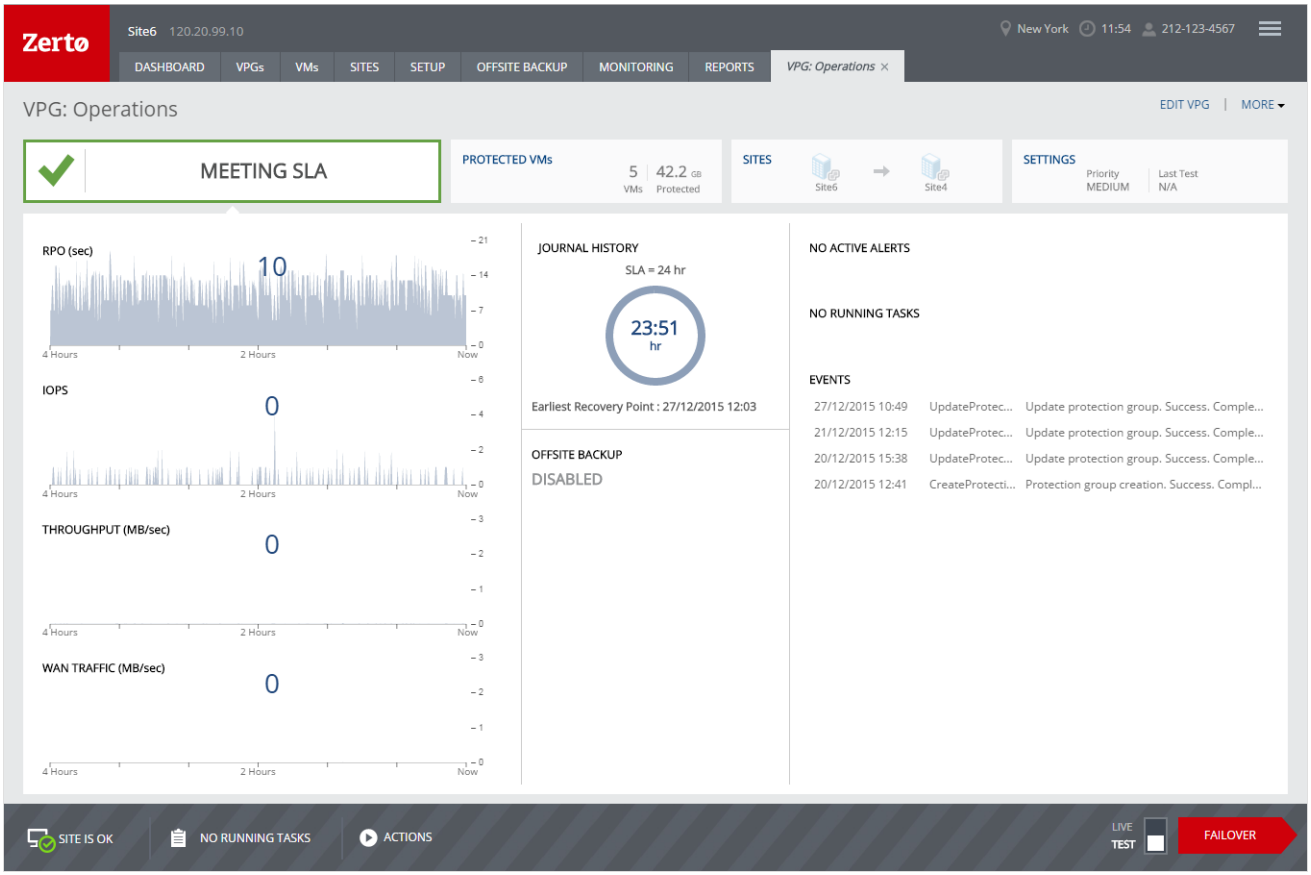

For details of the screen, see "Monitoring a Single VPG", on page 33.

Recovery

After initializing the VPG, all writes to the protected virtual machines are sent by the VRA on the relevant host for each virtual machine on the protected site to the VRA on AWS. The information is saved in the journal for the virtual machine with a timestamp, ensuring write-fidelity. Every few seconds the Zerto Virtual Manager writes a checkpoint to every journal on AWS for every virtual machine in the VPG, ensuring crash-consistency.

The data remains in the journal, in an S3 bucket, for the time defined by the journal history configuration, after which it is moved to the relevant mirror disks, also in the S3 bucket, for each virtual machine. Both the journal and the mirror disks are managed by the VRA.

When recovering, either a failover or move, or testing failover or cloning protected virtual machines in the recovery site, you specify the checkpoint at which you want the recovered virtual machines to be recovered. The mirror disks and journal are used to recover the virtual machines to this point-in-time. The recovered virtual machines are created as new instances in EC2.

File and Folder Recovery

After initializing the VPG, instead of recovering a virtual machine, you can recover specific files and folders in the protected virtual machines from a checkpoint.

Offsite Backups

After initializing the VPG, Zerto Virtual Replication periodically checks that the schedule to run an offsite backup has not been passed, either a daily or weekly. At the scheduled backup time, the offsite backup is run and the offsite backup file stored in the specified repository.

Offsite backups are kept on a ZCA configured repository for the retention period specified in the VPG. However, over time the number of stored offsite backups is reduced to save space.

The number of stored offsite backups for daily backups is as follows:

RETENTION PERIOD	DAILY	WEEKLY	MONTHLY	NUMBER OF BACKUPS	MAXIMUM NUMBER OF DAYS TO OLDEST BACKUP
1 week	7	0	0	7	7
1 month	7	4	0	11	35
3 months	7	4	2	13	91
6 months	7	4	5	16	175
9 months	7	4	8	19	259
12 months	7	4	11	22	343

That is, an offsite backup is kept for each day for the current week and then the oldest offsite backup for the previous week is kept for the previous four weeks and then the oldest monthly backup is kept for the rest of the retention period.

The number of stored offsite backups for weekly backups is as follows:

RETENTION PERIOD	WEEKLY	MONTHLY	NUMBER OF BACKUPS	MAXIMUM NUMBER OF DAYS TO OLDEST BACKUP
1 week	1	0	1	7
1 month	4	1	5	58
3 months	4	3	7	121
6 months	4	6	10	205
9 months	4	9	13	289
12 months	4	12	16	373

That is, an offsite backup is kept for each week for the current month and then the oldest backup for the month is kept and then the oldest monthly backup is kept for the rest of the retention period.

CHAPTER 5: MONITORING ZERTO VIRTUAL REPLICATION

You can monitor information about all the VPGs either protected at the local site or recovered to the local site in the *VPGs* tab. You can also drill-down to monitor information about a specific VPG displayed in the *VPGs* tab or about the virtual machines being protected by VPGs. You can also view summary details of the protected and AWS site in either the protected site or AWS, as well as monitor the status of each virtual protection group and any of the virtual machines being protected in either site.

The following general monitoring options are described in this chapter:

- "The DASHBOARD Tab", below
- "Monitoring VPGs – The VPGs Tab", on page 29
- "Monitoring a Single VPG", on page 33
- "Monitoring Tasks", on page 36
- "Monitoring Protected Virtual Machines – The VMs Tab", on page 38

The following site monitoring option is described in this chapter:

- "Monitoring Peer Sites - The SITES Tab", on page 40

The following offsite backup monitoring options are described in this chapter:

- "Monitoring Repositories – The SETUP Tab – The REPOSITORIES Tab", on page 41
- "Monitoring Offsite Backups – The OFFSITE BACKUP Tab", on page 42

For details about monitoring Zerto Virtual Manager alerts and events, refer to *Zerto Virtual Replication Guide to Alarms, Alerts and Events*.

The DASHBOARD Tab

The DASHBOARD provides an overview of the sites and VPGs being protected at the protected site or at AWS.

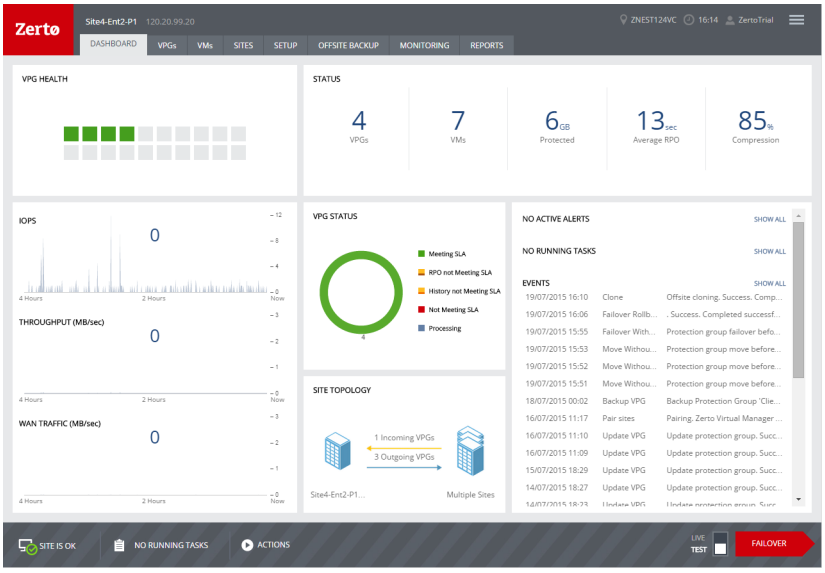

The following information is displayed:

VPG HEALTH

The VPGs being recovered to AWS with the health of each VPG, represented by a colored block, where the color represents the following:

Green – The VPG is being replicated, including syncing the VPG between the sites.

Orange – The VPG is being replicated but there are problems, such as an RPO value larger than the target RPO value specified for the VPG.

Red – The VPG is not being replicated, for example because communication with AWS is down.

Positioning the mouse over a block displays the VPG name as a tooltip. Clicking the block opens the details tab for the VPG.

STATUS

The status of the site, including the following:

- The number of VPGs and virtual machines being protected or recovered.
- The amount of storage being protected.
- The average RPO.
- The percentage compression of data passed between the site and peer sites.

Performance Graphs

The current site performance, which includes the following information:

IOPS (IO per second) – The IO between all the applications running on the virtual machines being protected and the VRA that sends a copy to AWS for replication.

Throughput (MB/sec) – The MBs for all the applications running on the virtual machines being protected. There can be a high IO rate with lots of small writes resulting in a small throughput as well as a small IO with a large throughput. Thus, both the `IOPS` and `Throughput` values together provide a more accurate indication of performance.

WAN Traffic (MB/sec) – The VPG related outgoing traffic between the sites.

VPG STATUS

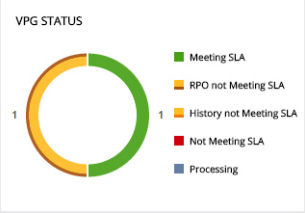

The status of the VPGs displayed as a pie chart. The legend describes what the pie chart colors represent.

SITE TOPOLOGY

A graphical display of the sites including the number of VPGs.

ACTIVE ALERTS, RUNNING TASKS, and EVENTS

A listing of the currently active alerts and running tasks, and the events run during the last few hours.

User input, for example, stopping a failover test or committing or rolling back a Move or Failover operation, can be initiated from the relevant task displayed in the RUNNING TASKS section.

Monitoring VPGs – The VPGs Tab

View details of all VPGs in the *VPGs* tab. This tab lists all the VPGs from both the local and remote sites and provides summary details of each VPG.

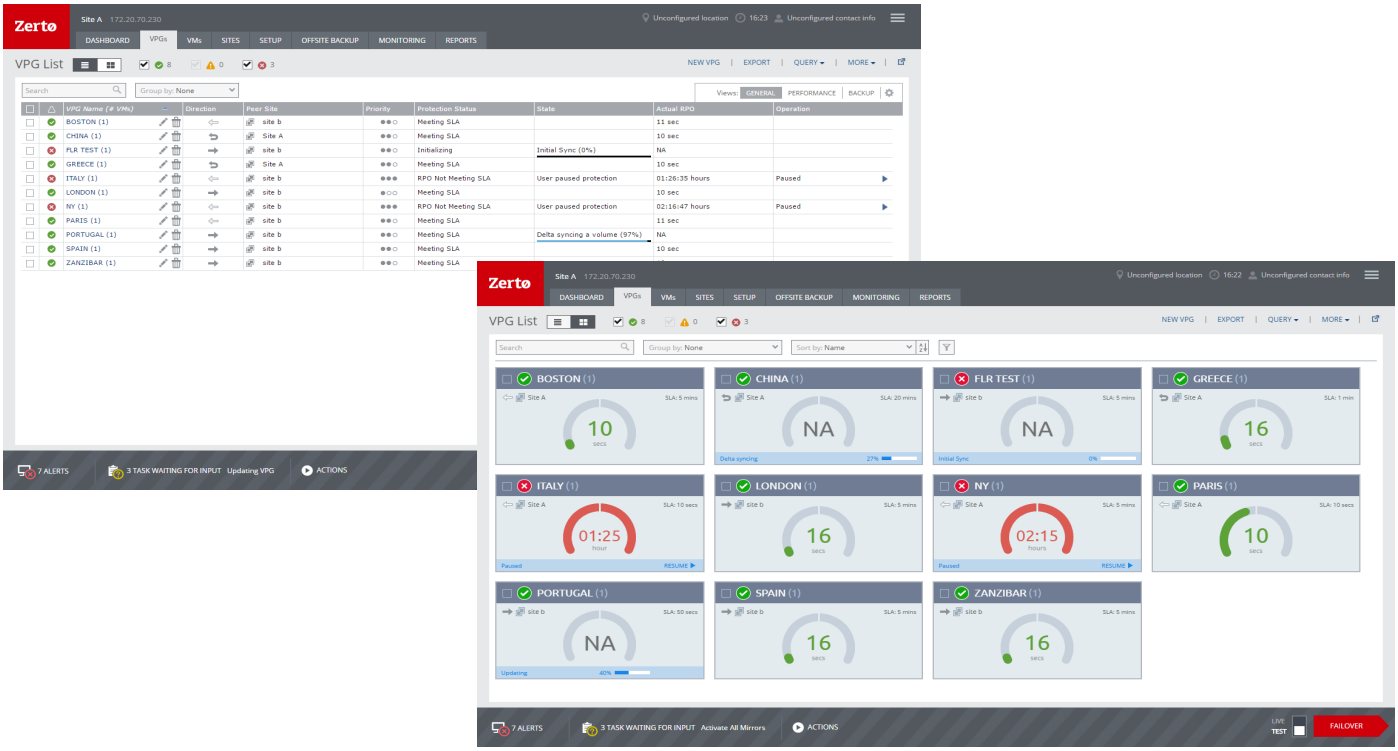

You can create a query using the view buttons ( )to display VPG information in a list or as a grid. In both list and grid views you can filter the VPGs that will be displayed according to their status by checking the checkboxes alongside the VPG status icons (3 1 1). The query can be customized by adding and removing filters.

The *QUERY* option allows you to save or run a personal query, or set the VPG tab back to its default view.

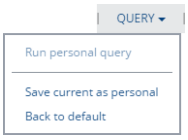

List View - GENERAL

The following information is displayed in the *GENERAL* view:

Alert status indicator – The color indicates the status of the VPG. Hovering over the alert displays a popup of all active alerts with descriptions:

Green – The VPG is being replicated, including syncing the VPG between the sites.

Orange – The VPG is being replicated but there are problems, such as an RPO value larger than the `Target RPO Alert` value specified for the VPG.

Red – The VPG is not being replicated, for example, because communication with the remote site is down.

Move the cursor over the *Alert status indicator* to display details of the alert.

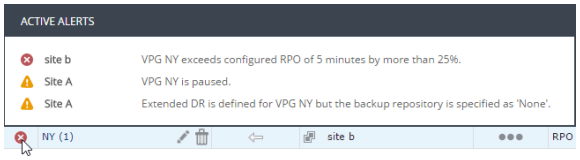

VPG Name (#VMs) – The name of the VPG. The name is a link: Click the VPG name to drill-down to more specific details about the VPG that are displayed in a dynamic tab. The number of VMs protected in the VPG is displayed in parentheses.

Direction – The direction of the replication, from this site to the remote site or from the remote site to this site.

Peer Site – The name of the site with which this site is paired: the site where the VPG is protected or will be recovered to.

Priority – The priority of the VPG.

Protection Status – The current status of the VPG, such as *Meeting SLA*. Where appropriate, the percentage of the operation completed, such as syncing, is displayed.

State – The current substatus of the VPG, such as *Delta syncing*. Where appropriate, the percentage of the operation completed, such as syncing, is displayed.

Actual RPO – The time since the last checkpoint was written to the journal. This should be less than the `Target RPO Alert` value specified for the VPG.

Operation – The operation, such as Move, that is currently being performed.

List View - PERFORMANCE

The following information is displayed in the *PERFORMANCE* view:

IO – The IO per second between all the applications running on the virtual machines in the VPG and the VRA that sends a copy to the remote site for replication.

Throughput – The MB per second for all the applications running on the virtual machines being protected. There can be a high IO rate with lots of small writes resulting in a small throughput as well as a small IO with a large throughput. Thus, both the `IOPS` and `Throughput` values together provide a more accurate indication of performance.

Network – The amount of WAN traffic.

Provisioned Storage (not shown by default) – The provisioned storage for all the virtual machines in the VPG. This value is the sum of the values that are used in the vSphere Client console per virtual machine in the *Virtual Machines* tab for the root vCenter Server node. Each value is the sum of both the hard disk and memory. Thus, a virtual machine with 1GB hard disk and 4GB memory will show 5GB provisioned storage.

Used Storage – The storage used by all of the virtual machines in the VPG. This value is the sum of the values that are used in the vSphere Client console per virtual machine in the *Virtual Machines* tab for the root vCenter Server node.

List View - BACKUP

The following information is displayed in the *BACKUP* view:

Alert status indicator – The color indicates the status of the VPG. Hovering over the alert displays a popup of all active alerts with descriptions:

Green – The VPG is being replicated, including syncing the VPG between the sites.

Orange – The VPG is being replicated but there are problems, such as an RPO value larger than the `Target RPO Alert` value specified for the VPG.

Red – The VPG is not being replicated, for example, because communication with the remote site is down.

Move the cursor over the *Alert status indicator* to display details of the alert.

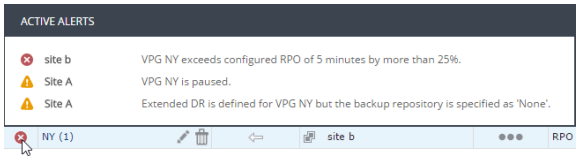

VPG Name (#VMs) – The name of the VPG. The name is a link: Click the VPG name to drill-down to more specific details about the VPG that are displayed in a dynamic tab.

Retention Policy – Whether the VPG is protected against a disaster only with the ability to recover to a point in time up to 30 days before the disaster, or protection is extended to include offsite backups of the virtual machines, going back for a maximum of one year.

Backup Status – The status of the backup.

Backup Repository – The name of the repository where the jobs are stored.

Restore Point Range – The restore points for the backup jobs out of the total backup jobs run for the VPG.

Backup Scheduling – The schedule for offsite backups.

Additional Fields and Options

In the *GENERAL*, *PERFORMANCE*, and *BACKUP* views you can:

- *Show/Hide Columns*, *Create View* and *Reset Columns* using the settings (⚙) menu.
- Sort the list by a column by clicking in the column title.
- Filter information in the columns by clicking the filter button (▽) that is displayed when the mouse cursor is moved into the column title. Active filters are displayed with a yellow background.

Grid View

In the grid view each VPG is displayed as a card.

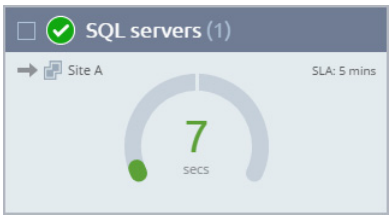

The default view is of all the VPG cards, un-grouped and sorted by VPG name.

The cards displayed can be filtered by clicking the filter button(▽). The default filters are *Direction* and *Protection Status*. You can click the **ADD** button to open the filters drop-down, and select additional filters. Active filters are displayed with a yellow background.

Each card contains the following:

Alert status indicator – The color indicates the status of the VPG. Hovering over the alert displays a popup of all active alerts with descriptions:

Green – The VPG is being replicated, including syncing the VPG between the sites.

Orange – The VPG is being replicated but there are problems, such as an RPO value larger than the `Target RPO Alert` value specified for the VPG.

Red – The VPG is not being replicated, for example, because communication with the remote site is down.

Move the cursor over the *Alert status indicator* to display details of the alert.

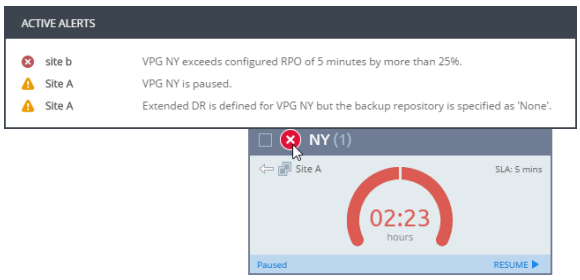

VPG Name (#VMs) – The name of the VPG. The name is a link: Click the VPG name to drill-down to more specific details about the VPG that are displayed in a dynamic tab. The number of VMs protected in the VPG is displayed in parentheses.

Direction – The direction of the replication, from this site to the remote site or from the remote site to this site.

Peer Site – The name of the site with which this site is paired: the site where the VPG is protected or will be recovered to.

State – The current substatus of the VPG, such as *Delta syncing*. Where appropriate, the percentage of the operation completed, such as syncing, is displayed.

Actual RPO – The time since the last checkpoint was written to the journal. This should be less than the `Target RPO Alert` value specified for the VPG.

Operation – The operation, such as Move, that is currently being performed.

Saving Details of Virtual Protection Groups to a File

You can save details of every VPG displayed in the `VPGs` tab to a `CSV` file, which can be opened using programs such as Microsoft Excel.

In the *VPGs* tab, click *EXPORT* and specify where to save the VPG details.

Monitoring a Single VPG

You can monitor the status of a specific VPG by clicking the VPG name in the *VPGs* tab or clicking the VPG name in the *VMs* tab. The VPG details are displayed in a dynamic tab.

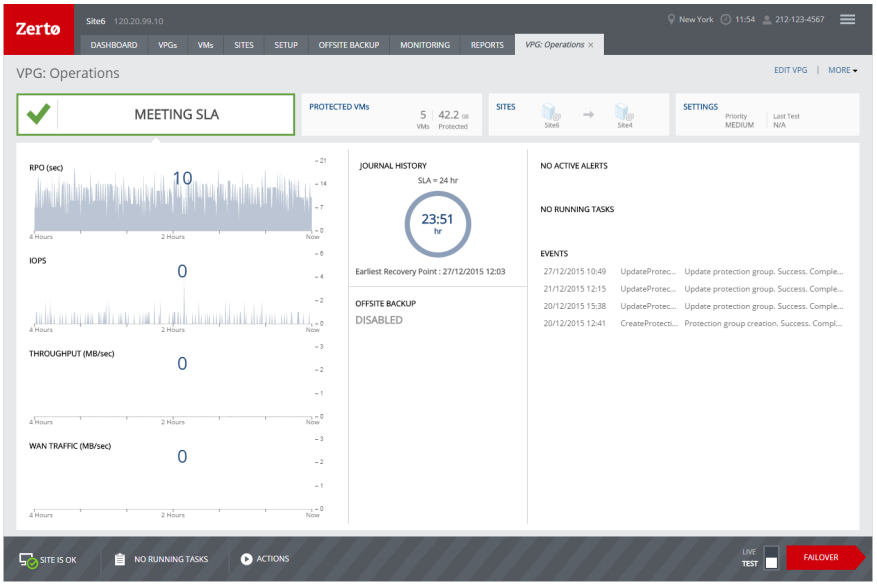

General Tab

The tab on the left side shows the status of the VPG. The following information is displayed in this tab:

Performance Graphs

The current VPG performance, which includes the following information:

RPO (sec) – The time since the last checkpoint was written to the journal. This should be less than the `Target RPO Alert` value specified for the VPG.

IOPS – The IO per second between all the applications running on the virtual machines in the VPG and the VRA that sends a copy to the remote site for replication.

Throughput (MB/sec) – The MB per second for all the applications running on the virtual machines being protected. There can be a high IO rate with lots of small writes resulting in a small throughput as well as a small IO with a large throughput. Thus, both the `IOPS` and `Throughput` values together provide a more accurate indication of performance.

WAN TRAFFIC (MB/sec) – The outgoing traffic between the sites.

JOURNAL HISTORY

The journal history shows:

- The SLA defined for the VPG.
- The amount of time currently covered by information in the journal.
- The earliest—oldest—checkpoint currently in the journal that can be used for a recovery operation.

OFFSITE BACKUP

If backup is enabled, the following backup details are displayed:

Retention Policy – Whether the VPG is protected against a disaster only with the ability to recover to a point in time up to 30 days before the disaster, or protection is extended to include offsite backups of the virtual machines, going back for a maximum of one year.

Backup Status – The status of the backup.

Backup Repository – The name of the repository where the jobs are stored.

Restore Point Range – The restore points for the backup jobs out of the total backup jobs run for the VPG.

Backup Scheduling – The schedule for offsite backups.

ACTIVE ALERTS, RUNNING TASKS, and EVENTS

A listing of the currently active alerts and running tasks, and the events run during the last few hours.

User input, for example, stopping a failover test or committing or rolling back a Move or Failover operation, can be initiated from the relevant task displayed in the RUNNING TASKS section.

PROTECTED VMs Tab

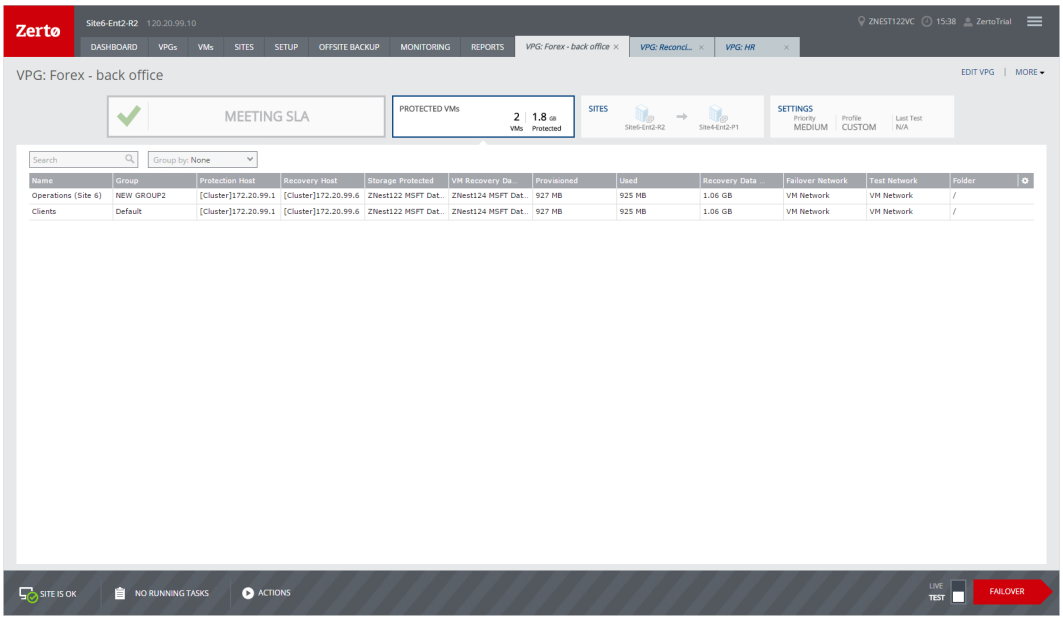

The *PROTECTED VMs* tab shows details about the protected virtual machines:

Name – The name of the virtual machine.

Group – The boot order group to which the virtual machine belongs.

Protection Host – The protected virtual machine host.

Storage Protected – The name of the protected storage.

Provisioned – The protected virtual machine provisioned storage.

Used – The amount of data used on the recovery site for this virtual machine.

Recovery Data Size – The total size of the data on the recovery site.

Failover Network – The failover network used when recovering this virtual machine.

Test Network – The test network used when testing the recovery of this virtual machine.

The following details are displayed with a vSphere recovery site:

Recovery Host – The host to use for recovery.

VM Recovery Datastore – The name of the recovery datastore.

Folder – The folder where the virtual machine is recovered to.

The following details are displayed with a Hyper-V recovery site:

Recovery Host – The host to use for recovery.

VM Recovery Storage – The name of the recovery storage.

The following details are displayed with an AWS recovery site:

Failover/Move VPC – The virtual network dedicated to your AWS account during a failover or move operation. A security group and subnet must be assigned to this VPC.

Failover/Move Subnet – The subnet mask for the VPC network during a failover or move operation.

Failover/Move Security Groups – The AWS security to be associated with the virtual machines in this VPG during a failover or move operation.

Test VPC – The virtual network dedicated to your AWS account during a failover test operation. A security group and subnet must be assigned to this VPC.

Test Subnet – The subnet mask for the VPC network during a failover test operation.

Test Security Groups – The AWS security to be associated with the virtual machines in this VPG. during a failover test operation.

Folder – The folder where the virtual machine is recovered to.

SITES Tab

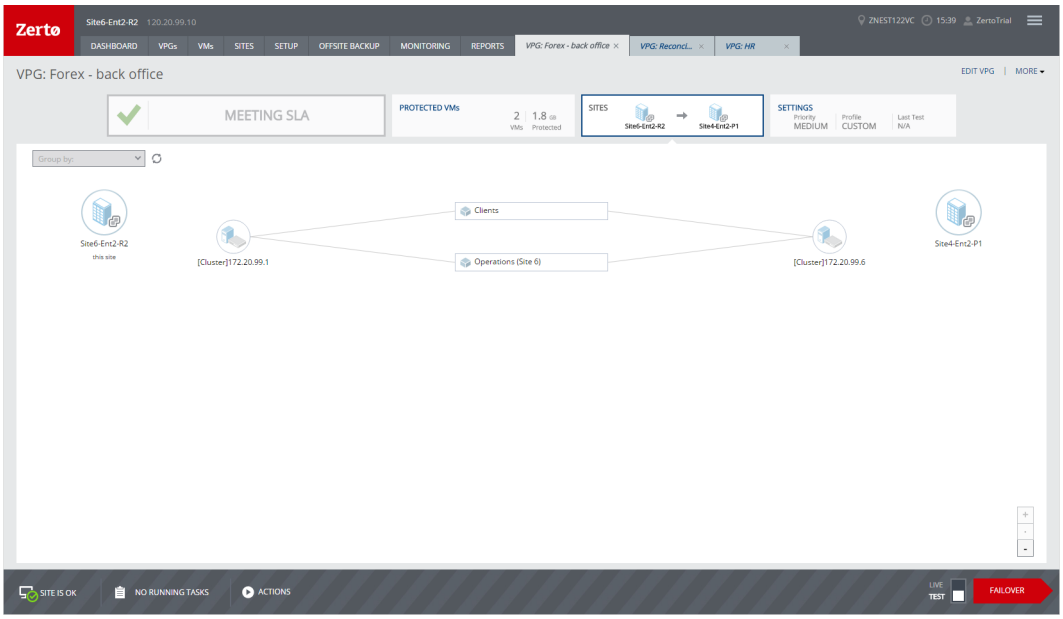

The *SITES* tab shows the topology of the VPG, including both the protected and recovery sites.

SETTINGS Tab

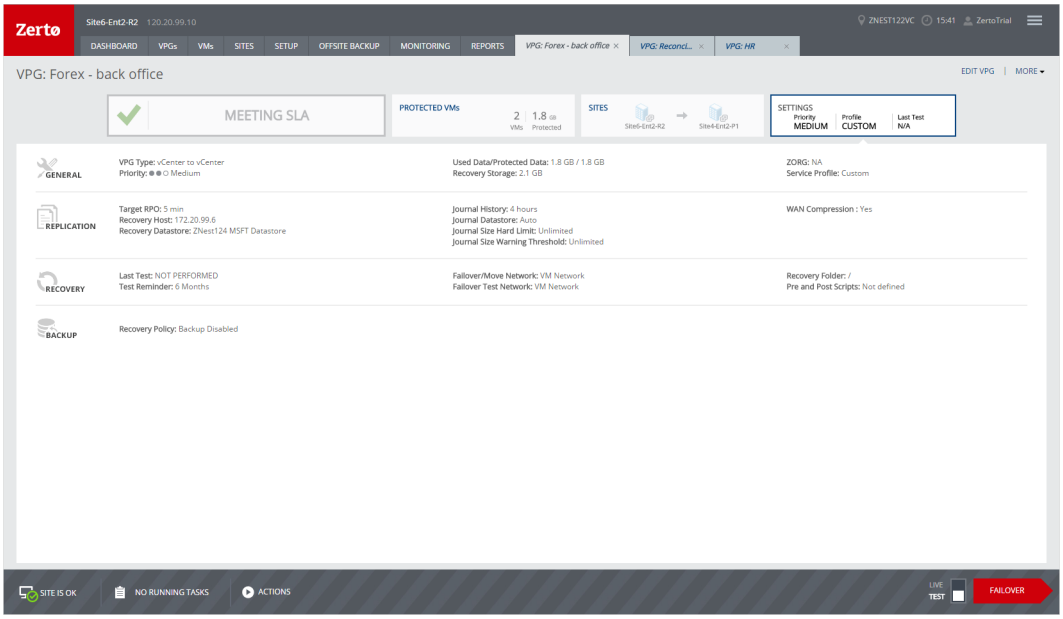

The *SETTINGS* tab shows details about the VPG settings, divided into general, replication, recovery, and backup categories.

Monitoring Tasks

Recent tasks can also be reviewed for a site by clicking the TASKS area in the status bar at the bottom of the user interface.

The following information is displayed for each task:

Status – The task status.

Name – The name of the task.

Description – A description of the task.

Action – The ability to perform an action directly. For example, stop a failover test, or commit or rollback a move or failover operation.

The full details of the tasks can be monitored in the TASKS subtab under the MONITORING tab.

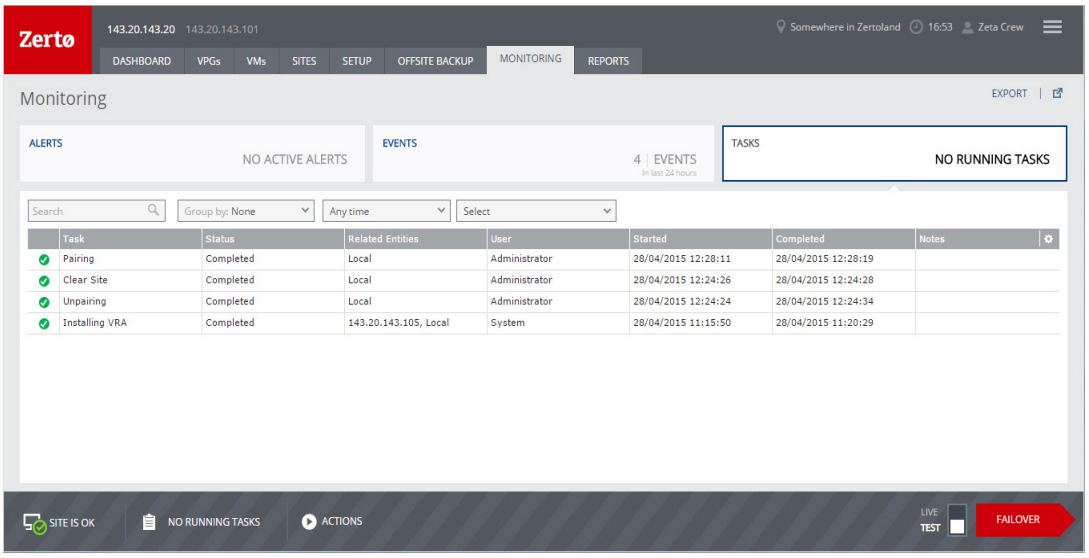

The following information is displayed for each task:

Task status indicator – The color indicates the status of the task. The following statuses exist for each task:

Green – The task was completed successfully.

Red – The task failed.

Task – The task name.

Status – The task status.

Related Entities – The sites which were effected by the task.

User – The user who initiated the task.

Started – The date and time the task started.

Completed – The date and time the task completed.

Notes – Notes added at the completion of a failover test.

Monitoring Protected Virtual Machines – The VMs Tab

View details of the protected VMs in the *VMs* tab. This tab lists all the protected virtual machines from both the local and remote sites and provides summary details of each virtual machine.

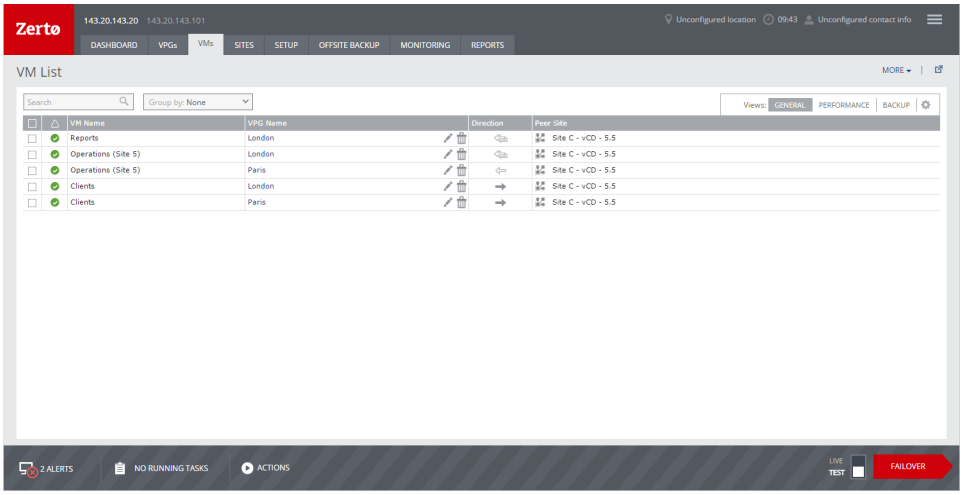

You can filter information in columns via the filter icon next to each column title. You can also sort the list by each column.

GENERAL View

The following information is displayed in the *GENERAL* view:

Alert status indicator – The color indicates the status of the VPG:

Green – The VPG is being replicated, including syncing the VPG between the sites.

Orange – The VPG is being replicated but there are problems, such as an RPO value larger than the `Target RPO Alert` value specified for the VPG.

Red – The VPG is not being replicated, for example, because communication with the remote site is down.

VM Name – The name of the virtual machine. The name is a link.

VPG Name – The name of the VPG. The name is a link: Click the VPG name to drill-down to more specific details about the VPG that are displayed in a dynamic tab.

Direction – The direction of the replication, from this site to the remote site or from the remote site to this site.

Peer Site – The name of the site with which this site is paired: the site where the VPG is protected or will be recovered to.

Priority – The priority of the VPG.

Protection Status – The current status of the virtual machine, such as *Meeting SLA*. Where appropriate, the percentage of the operation completed, such as syncing, is displayed.

State – The current substatus of the VPG, such as *Delta syncing*. Where appropriate, the percentage of the operation completed, such as syncing, is displayed.

Actual RPO – The time since the last checkpoint was written to the journal. This should be less than the `Target RPO Alert` value specified for the VPG.

Operation – The operation, such as Move, that is currently being performed.

PERFORMANCE View

The following information is displayed in the *PERFORMANCE* view:

IO – The IO per second between all the applications running on the virtual machine and the VRA that sends a copy to AWS.

Throughput – The MB per second for all the applications running on the virtual machines being protected. There can be a high IO rate with lots of small writes resulting in a small throughput as well as a small IO with a large throughput. Thus, both the `IOPS` and `Throughput` values together provide a more accurate indication of performance.

Network – The amount of WAN traffic.

Provisioned Storage – The provisioned storage for the virtual machine in the recovery site. Thus, a virtual machine with 1GB hard disk and 4GB memory will show 5GB provisioned storage.

Used Storage – The storage used by the virtual machine in the recovery site.

BACKUP View

The following information is displayed in the *BACKUP* view:

Alert status indicator – The color indicates the status of the VPG. Hovering over the alert displays a popup of all active alerts with descriptions:

Green – The VPG is being replicated, including syncing the VPG between the sites.

Orange – The VPG is being replicated but there are problems, such as an RPO value larger than the `Target RPO Alert` value specified for the VPG.

Red – The VPG is not being replicated, for example, because communication with the remote site is down.

Move the cursor over the *Alert status indicator* to display details of the alert.

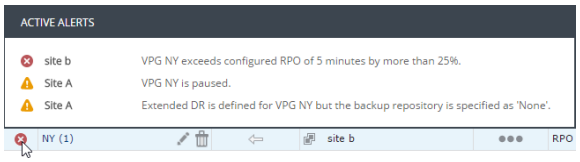

VPG Name (#VMs) – The name of the VPG. The name is a link: Click the VPG name to drill-down to more specific details about the VPG that are displayed in a dynamic tab.

Retention Policy – Whether the VPG is protected against a disaster only with the ability to recover to a point in time up to 30 days before the disaster, or protection is extended to include offsite backups of the virtual machines, going back for a maximum of one year.

Backup Status – The status of the backup.

Backup Repository – The name of the repository where the jobs are stored.

Restore Point Range – The restore points for the backup jobs out of the total backup jobs run for the VPG.

Backup Scheduling – The schedule for offsite backups.

Additional Fields

In the *GENERAL*, *PERFORMANCE*, and *BACKUP* views you can:

- *Show/Hide Columns*, *Create View* and *Reset Columns* using the settings (⚙) menu.
- Sort the list by a column by clicking in the column title.
- Filter information in the columns by clicking the filter button (▽) that is displayed when the mouse cursor is moved into the column title. Active filters are displayed with a yellow background.

Monitoring Peer Sites - The SITES Tab

View details of the paired sites in the *SITES* tab. This tab lists all the sites paired to the local site and provides summary details of each paired site.

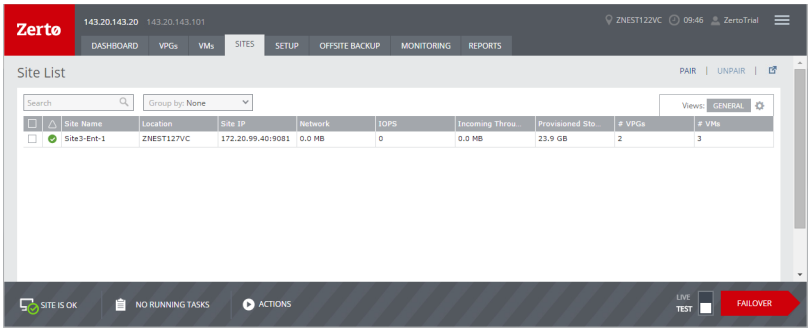

You can filter information in columns via the filter icon next to each column title. You can also sort the list by each column.

GENERAL View

The following information is displayed in the *GENERAL* view:

Alert status indicator – The color indicates the alert status of the site:

Green – The Zerto Virtual Manager for the site is running without problems.

Orange – The Zerto Virtual Manager for the site has a problem that does not stop the protection of virtual machines, such as an RPO value larger than the `Target RPO Alert` value for a VPG.

Red – The Zerto Virtual Manager for the site is not running correctly, for example, because communication with the site is down.

Site Name – The name specified for the paired site during installation or in the *Site Settings* dialog.

Location – The location specified for the paired site during installation or in the *Site Settings* dialog.

Site IP – The IP of the peer site.

Network – The amount of WAN traffic.

IOPS – The IO per second between all the applications running on the virtual machine in the VPG and the VRA that sends a copy to the remote site for replication.

Incoming Throughput – The MBs for all the applications running on the virtual machine being protected. There can be a high IO rate with lots of small writes resulting in a small throughput as well as a small IO with a large throughput. Thus, both the `IO` and `Incoming Throughput` values together provide a more accurate indication of performance.

Provisioned Storage (GB) – The maximum storage that can be protected.

 # VPGs – The total number of VPGs being protected by the site and replicated to the site.

 # VMs – The total number of virtual machines being protected by the site and replicated to the site.

Additional Fields

There are additional fields that you can display that are listed when you select *Show/Hide Columns* from the dropdown list shown by clicking the configuration icon (⚙):

Used Storage (GB) – The name of the protected site.

ZORG Name – A name given to the organization by a cloud service provider. For details refer to *Zerto Cloud Manager Administration Guide*.

Version – The Zerto Virtual Replication version installed at this site.

Monitoring Repositories – The SETUP Tab – The REPOSITORIES Tab

View details of the repositories that can be used for offsite backup jobs in the *REPOSITORIES* subtab, under the *SETUP* tab. This tab lists all the repositories created for the site.

You can filter information in columns via the filter icon next to each column title. You can also sort the list by each column.

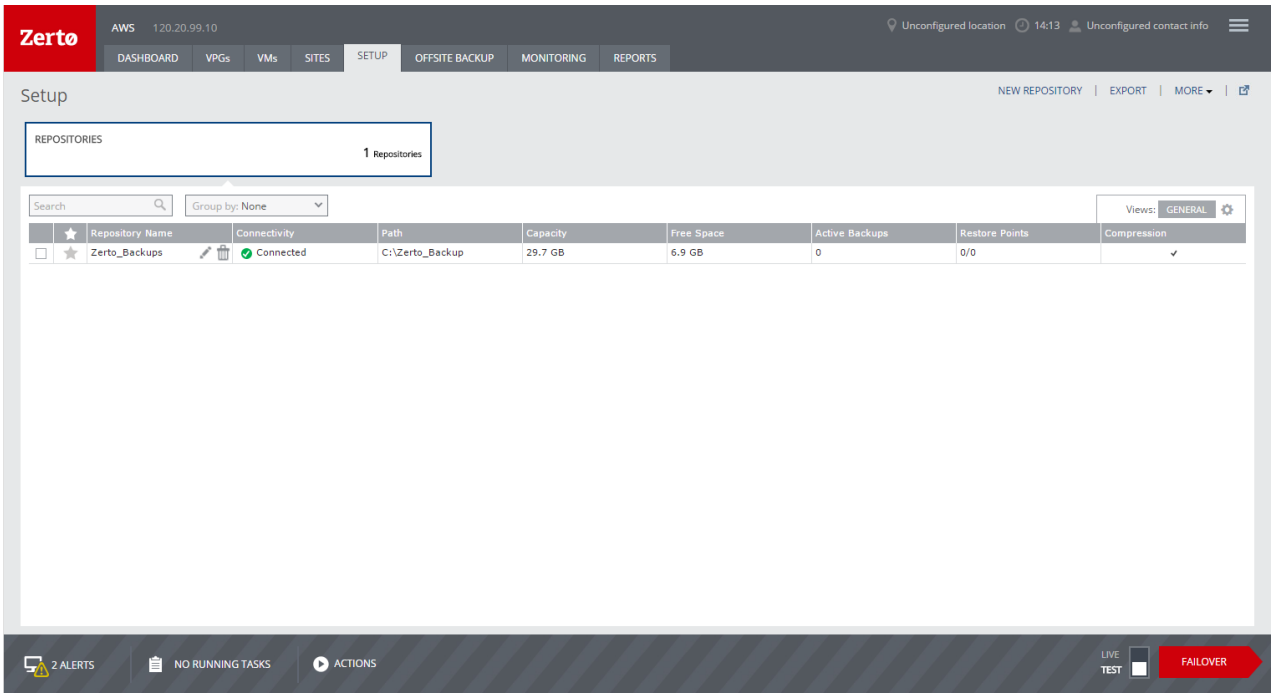

GENERAL View

In this view, the number of repositories is displayed in the *REPOSITORIES* tab. The following information is displayed in this view:

Delete Repository – Links to edit or delete a repository.

Repository Name – The name of the repository. This field contains icons that you can click to edit or delete the repository.

Repository Type – The type of repository. The options are Local or Network Share (SMB).

Connectivity – Whether the repository is connected or not.

Path – The path to the repository.

Capacity – The overall capacity of the repository.

Free Space – The amount of free space currently available on the repository.

Active Backups – The number of backup jobs currently active that are stored in the repository.

Restore Points – The restore points for the backup jobs out of the total backup jobs saved to the repository.

Compression – A check in this field means that the backups stored in the repository are compressed.

Click *NEW REPOSITORY* to display the *New Repository* dialog that you can use to create a new repository.

Monitoring Offsite Backups – The OFFSITE BACKUP Tab

View details of the offsite backup jobs in the *OFFSITE BACKUP* tab either by VPG or virtual machine. This tab lists all the defined offsite backups and their statuses.

See also:

- "VPGs Tab", on page 42
- "VMs Tab", on page 43

VPGs Tab

View details of the offsite backup jobs by VPG.

You can filter information in columns via the filter icon next to each column title. You can also sort the list by each column.

GENERAL View

The following information is displayed in the *GENERAL* view:

VPG Name – The name of the VPG.

Backup Site – The site where the VPG is backed up. The backup jobs are stored either locally at this site or on a network shared drive which is accessible from this site.

Status – The status of the job: `Running` or `Scheduled`.

Repository Name – The name of the repository where the job is stored.

VPG Size – The size of the VPG in the last run stored on disk.

Result of Last Run – The result of the last run: `Full success`, `Partial success`, or `Failed`.

Restore Points – The restore points for the backup jobs out of the total backup jobs run for the VPG.

RUN DETAILS View

The following information is displayed in the *RUN DETAILS* view:

VM Name – The name of the Virtual machine.

VPG Name – The name of the VPG.

Result of Last Run – The result of the last run: `Full success`, `Partial success`, or `Failed`.

Time of Last Run – The time of the last run.

Next Scheduled Run – The time of the next scheduled run.

Last Full Backup – The date and time of the last full backup.

Additional Fields

There are additional fields that you can display that are listed when you select *Show/Hide Columns* from the dropdown list shown by clicking the configuration icon (⚙):

Protected Site – The name of the site.

Last Backup Size – The size of the last backup performed by Zerto Virtual Manager.

ZORG – A name given to an organization by a cloud service provider. For details refer to *Zerto Cloud Manager Administration Guide*.

 # VMs – The total number of virtual machines protected by the VPG.

# of Volumes – The number of volumes protected by the VPG.

VMs Tab

View details of the offsite backup jobs by virtual machine.

You can filter information in columns via the filter icon next to each column title. You can also sort the list by each column.

GENERAL View

The following information is displayed in the *GENERAL* view:

VM Name – The name of the virtual machine.

VPG Name – The name of the VPG.

Protected Site – The name of the site where the VPG is protected.

Backup Site – The site where the virtual machines are backed up. The backup jobs are stored either locally at this site or on a network shared drive which is accessible from this site.

Status – The status of the job.

Repository Name – The name of the repository where the job is stored.

VM Size – The size of the VMs stored on disk.

Result of Last Run – The result of the last run: `Full success`, `Partial success`, or `Failed`.

Restore Points – The restore points for the backup jobs out of the total backup jobs run for the VPG.

RUN DETAILS View

The following information is displayed in the *RUN DETAILS* view:

VM Name – The name of the Virtual machine.

VPG Name – The name of the VPG.

Result of Last Run – The result of the last run: `Full success`, `Partial success`, or `Failed`.

Time of Last Run – The time of the last run.

Next Scheduled Run – The time of the next scheduled run.

Last Full Backup – The date and time of the last full backup.

MORE Options

Click *MORE > Edit* to edit the backup parameters of the VPG.

Click *MORE > Abort Backup* to abort a running job. Any virtual machine volumes already stored in the repository are not removed and the job status is partial if there are any stored volumes.

Click *MORE > Run Backup* to start a job for a selected VPG, outside of the schedule for that VPG.

Click *EXPORT* to export the backup list as a Microsoft Excel worksheet.

Additional Fields

There are additional fields that you can display that are listed when you select *Show/Hide Columns* from the dropdown list shown by clicking the configuration icon ( ):

Last Backup Size – The size of the last backup performed by Zerto Virtual Manager.

ZORG – A name given to an organization by a cloud service provider. For details refer to *Zerto Cloud Manager Administration Guide*.

# of Volumes – The number of volumes associated with the VM.

CHAPTER 6: MANAGING VPGS

After defining virtual protection groups (VPGs) the virtual machines specified as part of each VPG are protected. There are a number of ongoing management tasks that you can perform on a VPG, such as specifying a checkpoint to enable recovery to that specific point or you can modify the configurations of existing VPGs.

The following VPG management options are described in this chapter:

Monitoring VPGs and the VMs that are protected is described in "Monitoring Zerto Virtual Replication", on page 27.

Note: You cannot add a virtual machine to a VPG from the AWS site.

Editing a VPG

You can edit a VPG definition, including changing the information about how virtual machines are recovered.

Note: You cannot edit the VPG while a backup job is running.

After modifying the VPG, the definition is updated.

While the VPG definition is being updated, you cannot perform any operations on the VPG, such as adding a checkpoint, editing the VPG properties, or failing the VPG.

After the definition is updated, the VPG is synchronized with the recovery site.

To modify a VPG:

1. In the *VPGs* tab in the Zerto User Interface, select the VPG to be edited and click *MORE* > *Edit VPG*. You can also select the VPG, display the VPG details, and click *EDIT VPG*.

   The *Edit VPG* wizard is displayed, enabling editing the VPG, including adding and removing virtual machines from the VPG.

   Note: If the VPG was previously viewed, and the tab for this VPG is still displayed, you can access the details by selecting the tab.
2. Make any required changes to the VPG definition. You can jump directly to a step to make a change in that step, for example, the *REPLICATION* step or the *RECOVERY* step, by clicking the step. Steps that have been completed are marked with a check.
3. Click *DONE*.

   Note: The changed values are not applied to existing virtual machines but only to new virtual machines added to the VPG, When a virtual machine is removed from a VPG, a warning is displayed. Another message is displayed when trying to save the VPG, if a virtual machine is added to the VPG.

The VPG is updated and then synchronized with the recovery site, if required, for example when the host was changed.

See also:

- "Modifying the Journal Size Hard Limit", on page 46
- "Modifying the Retention Period for Offsite Backups", on page 46

Modifying the Journal Size Hard Limit

If the journal size hard limit is reduced, and if the current size is greater than the newly defined size, the journal remains at the current size. When the amount of the journal used falls below the hard limit value it will not grow greater than the new hard limit. Unused journal volumes from the added volumes are marked for removal and removed after the time equivalent to three times the amount specified for the journal history, or twenty-four hours, whichever is more.

Note: If the `Journal Size Hard Limit` or `Journal Size Warning Threshold` in the VPG SLA settings are changed, the changed values are not applied to existing virtual machines but only to new virtual machines added to the VPG.

Modifying the Retention Period for Offsite Backups

If the retention period was shortened, the number of backup jobs older than the new retention period are deleted from the repository.

Pausing the Protection of a VPG

During periods when the WAN bandwidth is utilized to its maximum, you can pause the protection of a VPG, to free up some of this bandwidth. After pausing the protection, the VPG can still be recovered to the last checkpoint written to the journal before the pause operation.

Note:

- Zerto recommends adding a checkpoint to the VPG immediately before pausing protection, if you might want to recover the VPG to the latest point in time before the pause.
- You cannot pause a VPG while a backup job is running.

To pause the protection of VPGs:

1. In the Zerto User Interface, click the *VPGs* or *VMs* tab and select one or more VPGs to pause protection.
2. Click *MORE* > *PAUSE*.

   A warning is displayed. If you click *PROCEED* in this warning, the VPG protection is paused.

   Note: If the VPG was previously viewed, and the tab for this VPG is still displayed, you can access the details by selecting the tab.

   The VPG protection is paused until you click *Resume VPGs*.

To resume the protection of VPGs:

1. In the Zerto User Interface, click the *VPGs* or *VMs* tab and select one or more VPGs to resume protection.
2. Click *MORE* > *Resume*.

   After resuming protection, a `Bitmap Sync` will most probably be performed to synchronize the protection and recovery sites.

Forcing the Synchronization of a VPG

If the protected virtual machines are updated such that they are no longer synchronized with their mirror machines in the recovery site, you can force the resynchronization of the machines. An example of when the machines can be out-of-sync is when there is a rollback of a virtual machine to a VMware snapshot. In this case, the recovery virtual machine will include changes that have been rolled back in the protected machine, so that they are no longer synchronized.

You can force the synchronization of the machines in a VPG to remedy this type of situation.

Note: You cannot force the synchronization of a VPG while a backup job is running.

To forcibly synchronize a VPG:

1. In the Zerto User Interface, select the VPGs or *VMs* tab and click the VPG to display the VPG details.
2. Click *MORE* > *Force Sync*.

   Note: If the VPG was previously viewed, and the tab for this VPG is still displayed, you can access the details by selecting the tab.

The VPG starts to synchronize with the recovery site. As the journal fills up during the synchronization, older checkpoints are deleted from the journal to make room for the new data and the data prior to these checkpoints are promoted to the virtual machine virtual disks. Thus, during the synchronization, you can recover the virtual machine to any checkpoint still in the journal, but as time progresses the list of checkpoints available can lessen. If the journal is not big enough to complete the synchronization without leaving at least ten minutes worth of checkpoints, the synchronization pauses for the time specified in the `Replication Pause Time` value for the VPG, to enable intervention to ensure recovery to a checkpoint remains available. The intervention can be, for example, increasing the size of the journal, or cloning the journal as described in "Deleting a VPG", below.

Deleting a VPG

You can delete a VPG. Any offsite backups stored for the VPG are not deleted and the virtual machines that were backed up can be restored.

Note: You cannot delete a VPG while a backup job is running.

To delete a VPG:

1. In the Zerto User Interface, click the *VPGs* or *VMs* tab and select one or more VPGs to delete.
2. Click *MORE* > *Delete*.

   The *Delete* dialog is displayed.
3. Check `Keep target disks at the peer site` if you might reprotect the virtual machines. When protecting to AWS, you cannot save the disks for preseeding.
4. Click *DONE* to delete the VPG.

   The VPG configuration is deleted. The VRA on the recovery site that handles the replication for the VPG is updated including keeping or removing the replicated data for the deleted VPG, dependent on the `Keep target disks at the peer site` setting during the deletion.

   The locations of the saved target disks are specified in the description of the event for the virtual machines being removed, event EV0040, displayed in *MONITORING* > *EVENTS*.

See also "Deleting a VPG When the Status is Deleting", on page 47.

Deleting a VPG When the Status is Deleting

If, for some reason, the VPG cannot be deleted, the VPG status changes to `Deleting` and the substatus is `VPG waiting to be removed`. Attempting to delete the VPG a second time causes the following to be displayed:

Retry – Retry deleting the VPG.

Force Delete – Forcibly delete the VPG.

Cancel – Cancel the delete operation.

Running an Unscheduled Offsite Backup

After initializing the VPG, Zerto Virtual Replication periodically checks that the schedule to run an offsite backup - either daily or weekly - has not passed. At the scheduled backup time, the offsite backup is run and the offsite backup file stored in the specified repository.

To run an unscheduled offsite backup:

1. In the Zerto User Interface, click the VPGs or *VMs* tabs and select one or more VPGs to be backed up.

   Note: You can also start from the *OFFSITE BACKUP* tab.
2. Click *MORE* > *Run Backup*.

   Note: If the VPG was previously viewed, and the tab for this VPG is still displayed, you can access the details by selecting the tab.
3. Click *OK*.

The offsite backup starts. You can monitor the progress in the Offsite Backup tab and the tasks pane. During the backup job you cannot perform any other operation on the VPG without first aborting the job. You can start a live failover and you are then prompted to abort the job.

Scheduled backup runs for the VPG are skipped until the unscheduled run ends.

If the job runs out of the configured backup window, the virtual machines that are already stored in the repository are kept but remaining virtual machines in the VPG are not backed up. The job is reported as a partial backup.

Ensuring Application Consistency – Checkpoints

Checkpoints are **recorded automatically** every **few seconds** in the journal. These checkpoints **ensure crash-consistency**, and are written to the virtual machines journals by the Zerto Virtual Manager.

Each checkpoint has the **same timestamp** which is set by the **Zerto Virtual Manager**.

During recovery you **pick a checkpoint** in the journal and **recover to this point**. The crash-consistent checkpoints **guarantee write order fidelity**.

For Example:

If **write A** on a virtual machine in the VPG occurred **before write B** on a virtual machine in the VPG, then when a checkpoint is written, the journal will contain:

- Neither of the writes
- Both writes, and if they overlap the B data takes precedence
- Only A – indicating the checkpoint occurred between A and B

The coordination is done by the Zerto Virtual Manager.

You can also integrate Microsoft Volume Shadow Copy Service (VSS) with Zerto Virtual Replication to ensure transaction consistency in a Microsoft Windows server environment.

You can also use a **script** to place the application in a quiesced mode, such as Oracle Hot Backup mode, and execute the Zerto Virtual Replication **PowerShell** cmdlet **Set-Checkpoint**, then release the quiesced mode. For more information about Zerto Virtual Replication PowerShell cmdlets, see *Zerto Virtual Replication Cmdlets*.

Note:

- To write application-consistent checkpoints, there is a performance impact on the virtual machine running the application as a result of the application-consistent mechanism used, such as VSS. This is because the guest operating system and any integrated applications will be quiesced.

  This impact on performance may be negligible and does not always happen since not all applications require these checkpoints in order to achieve successful application recovery. Also, Zerto Virtual Replication only requires the guest and application to quiesce for a brief moment, just long enough to add a checkpoint.

As previously mentioned, checkpoints are recorded every few seconds in the journal. After a while, the number of checkpoints available from which to choose a recovery point can be in excess of **thousands per VPG**.

When this threshold is reached, in order to **enable efficient management and use of the checkpoints**, the number of checkpoints is **diluted** with respect to time, as follows:

- Within the latest **2 hours**: All of the checkpoints are available for recovery.
- Between **2 and ~4.5 hours**: There are about two to three checkpoints every 15 minutes.
- From **4.5 hours and over**: 1 checkpoint is kept every 15 minutes.

  Note: Checkpoints which are either added manually, or added via the ZertoVssAgent, or marked as part of a Failover test are not diluted.

This section describes the different options available to ensure application consistency:

- "Adding a Checkpoint to Identify a Key Point", below.
- "Ensuring Transaction Consistency in Microsoft Windows Server Environments", on page 50.

Adding a Checkpoint to Identify a Key Point

In addition to the automatically generated checkpoints, you can add checkpoints manually to ensure application consistency and to identify events that might influence recovery, such as a planned switch-over to a secondary generator. You can recover the machines in a VPG to any checkpoint in the journal, to one added automatically or to one added manually. Thus, recovery is done to a point-in-time when the data integrity of the protected virtual machines is ensured.

Note:

- Adding a checkpoint manually does not guarantee transaction consistency.
- Changes to a VPG that result in re-synchronization of the VPG results in all checkpoints being removed. Adding checkpoints to the journal is resumed after synchronization completes. A forced synchronization of the VPG only removes checkpoints if the journal fills up during the synchronization.

To add a checkpoint to a VPG:

1. In the Zerto User Interface select *ACTIONS* > *ADD CHECKPOINT*.

   The *Add Checkpoint* dialog is displayed.

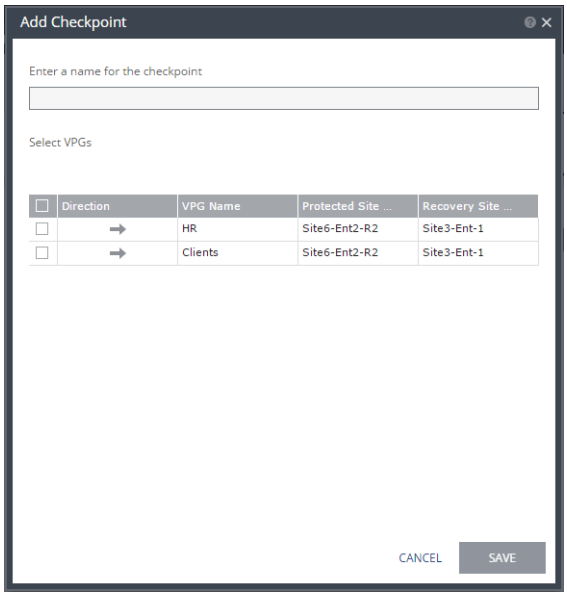

   A list of VPGs is displayed with the requested VPG selected. You can select more VPGs to add the same checkpoint to, for example, when something is happening at your site that affects multiple VPGs.

   Note: Crash-consistency is per VPG and not across VPGs, even if a checkpoint was added to multiple VPGs.

2. Enter a name for the checkpoint.

3. Click *SAVE*.

When testing a failover, as described in "Testing Recovery to AWS", on page 83, or actually performing a failover, as described in "Managing Failover to AWS", on page 97, you can choose the checkpoint as the point to recover to.

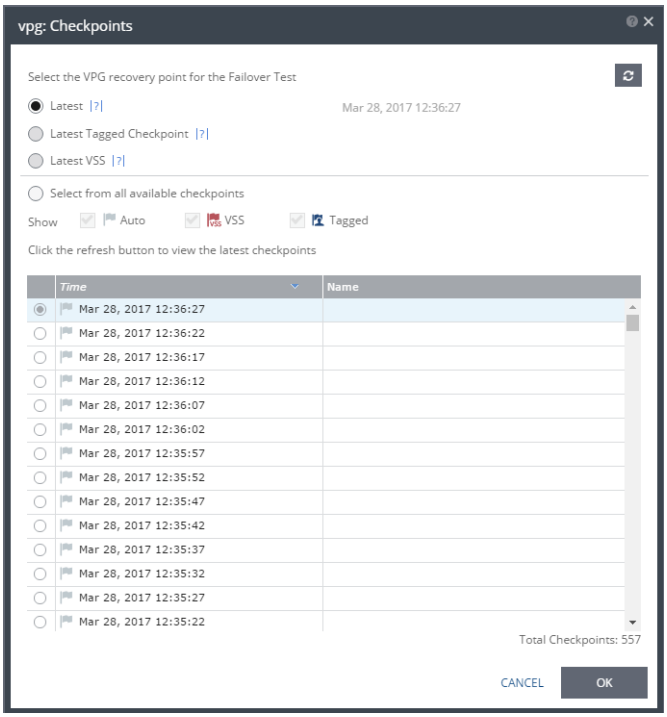

The checkpoints listed include checkpoints added via the *ZertoVssAgent*, as described in "Ensuring Transaction Consistency in Microsoft Windows Server Environments", below.

Ensuring Transaction Consistency in Microsoft Windows Server Environments

The Microsoft Volume Shadow Copy Service (VSS) enables taking manual or automatic offsite backup copies or snapshots of data, even if it has a lock, on a specific volume at a specific point-in-time over regular intervals. This ensures not just that the data is crash consistent but also transaction consistent if recovery is needed.

Zerto Virtual Replication enables adding checkpoints to the journal that are synchronized with VSS snapshots.

To use Zerto Virtual Replication with VSS to ensure application consistency you must install the *ZertoVssAgent* on every virtual machine that uses VSS and that you want to protect with Zerto Virtual Replication.

You can install the *ZertoVssAgent* on the following supported Windows operating systems:

OPERATING SYSTEMS
Windows Server 2008, all versions (SPs and R2)
Windows Server 2012, all versions (SPs and R2)

To install the ZertoVssAgent:

1. Download the ZertoVssAgent, *ZertoVss64Agent.msi*, from the Zerto Support Portal downloads page, on the virtual machines that use VSS and that you want to protect with Zerto Virtual Replication.
2. Run the ZertoVssAgent on the virtual machines that use VSS and that you want to protect.

   Note: Only a single virtual machine in a VPG can have application consistent checkpoints and the VSS checkpoint is only applied to the virtual machine where the ZertoVssAgent is installed. Thus, even if more than one virtual machine runs VSS, you only install the Zerto VssAgent on one of the virtual machines in the VPG. Also, the virtual machine where the

ZertoVssAgent is installed must have network connectivity to the local Zerto Virtual Manager in order to be able to add VSS checkpoints successfully.

3. Enter the license key and click *Validate*.

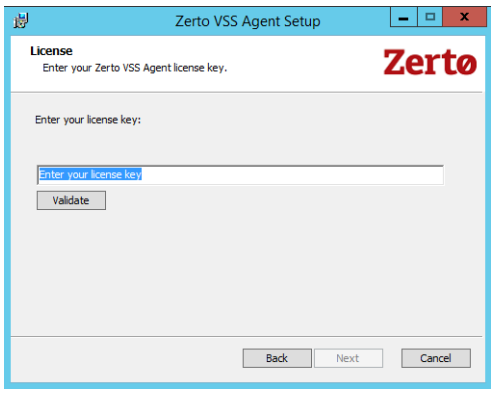

4. Follow the wizard through the installation.

   The `Zerto Virtual Manager Connections Settings` dialog is displayed.

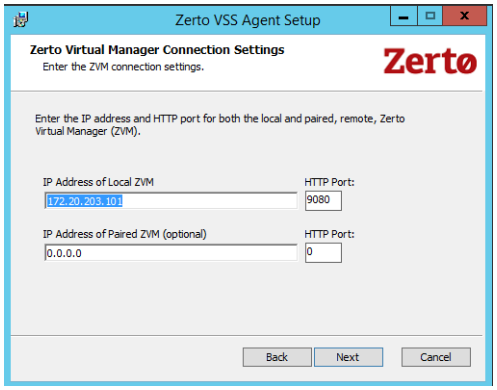

5. Specify the IP address and HTTP port number for the Zerto Virtual Managers managing the protection of the virtual machines, both for the local site and optionally, for the paired, remote site. If the same hypervisor manager is used both for protecting and recovering virtual machines, specify the IP address and HTTP port number for the single Zerto Virtual Manager installed.

   Note: The default HTTP port number when Zerto Virtual Replication is installed is 9080.

   If you enter a wrong IP address or port you can correct the address or port after the installation completes by editing the `ZertoVssAgentGUI.exe.conf` file in the `ZertoVssAgent` folder under the folder where the *ZertoVssAgent* is installed, for example, `C:\Program Files\Zerto`.

6. Click *OK*.

The *ZertoVssAgent* is installed and the `Add VSS Checkpoint` is placed on the desktop. The agent runs as a Windows service, *ZertoVssprovider*.

You can add a checkpoint to the Zerto Virtual Replication via the `Add VSS Checkpoint` dialog, via the command line or as a scheduled task. The *ZertoVssAgent* ensures that the virtual machine is in an application consistent state and then sends the checkpoint to the Zerto Virtual Manager, which then adds the checkpoint to the journals for the VPG containing that virtual machine.

The checkpoint is logged for the entire VPG, however any other virtual machine in the VPG will have a crash-consistent checkpoint.

To add a checkpoint while ensuring application consistency via the Add VSS Checkpoint dialog:

1. On a virtual machine where the *ZertoVssAgent* has been installed, click *Start* > *Programs* > *Zerto Virtual Replication* > *Add VSS Checkpoint* or double-click the *Add VSS Checkpoint* icon on the desktop.

   The `Add VSS Checkpoint` dialog is displayed.

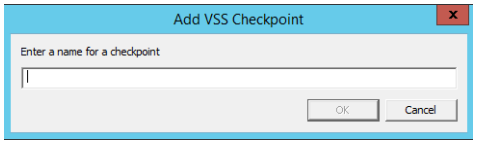

2. Enter a name for the checkpoint.
3. Click *OK*.

Note: A message that the process was completed is displayed on the machine where the *ZertoVssAgent* has been installed. The handling of the checkpoint by the Zerto Virtual Manager is done asynchronously and you can check via the recent tasks list in the Zerto User Interface that the checkpoint is added in the VPG.

To add a checkpoint while ensuring application consistency via the command line:

1. Open the command line dialog as an administrator.
2. Navigate to the directory where the *ZertoVssAgent* is installed. The default location is `C:\Program Files\Zerto\ZertoVssAgent\`
3. In the command line, run the following:

   `ZertoVssAgent.exe <localURL> <localPort> <remoteURL> <remotePort> <checkpoint>`

   where:

   localURL – The URL for the Zerto Virtual Manager that manages the protected site.

   localPort – The HTTP port for the Zerto Virtual Manager that manages the protected site.

   remoteURL – The URL for the Zerto Virtual Manager that manages the recovery site.

   remotePort – The HTTP port for the Zerto Virtual Manager that manages the recovery site.

   checkpoint – The name of the checkpoint.

Note: A message that the process was completed is displayed on the machine where the *ZertoVssAgent* is installed. The handling of the checkpoint by the Zerto Virtual Manager is done asynchronously and you can check via the recent tasks list in the Zerto User Interface that the checkpoint is added in the VPG.

To schedule checkpoints:

1. Open the Task Scheduler.
2. Under the *Actions* menu item, select `Create Task`.

   The `Create Task` dialog is displayed.

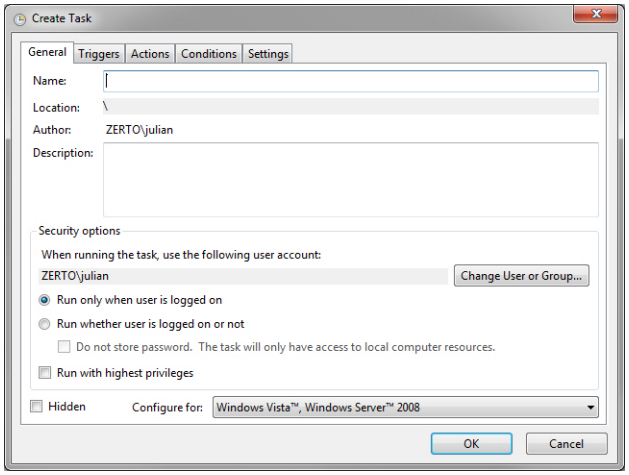

3. Enter the following:

   Name – A name for the task.

   Run whether the user is logged on or not – Make sure that this is checked.

   Run with highest privileges – Make sure that this is checked.

   The Windows Scheduled Task will be created and run by the currently logged in user. After the task is created, Zerto recommends changing this to `NT AUTHORITY\Network Service` permissions and follow the steps to allow the correct permissions as described in "To set COM permissions for VSS when "Access Denied" errors are received:", on page 56.

4. Select the *Triggers* tab and configure a new trigger.

   The `New Trigger` dialog is displayed.

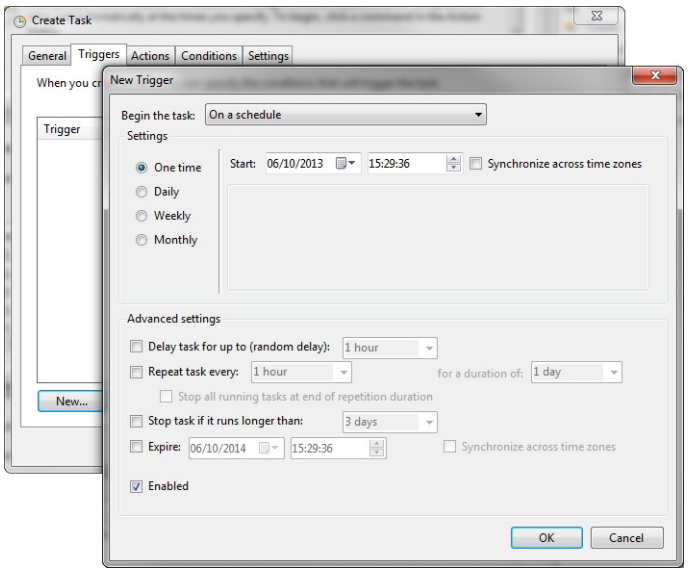

5. Select the *Actions* tab and create a new action to start the ZertoVssAgent with the IP address and port of the Zerto Virtual Manager and the checkpoint to use. For example:

   `C:\Program Files\Zerto\ZertoVssAgent\ZertoVssAgent.exe` and

   `106.18.206.10 9080 106.18.206.10 9080 "VSSTaskCP"`

   That is, with the format: `<protecting_ZVM_IP> 9080 <recovery_ZVM_IP> 9080 "<CP_name>"`

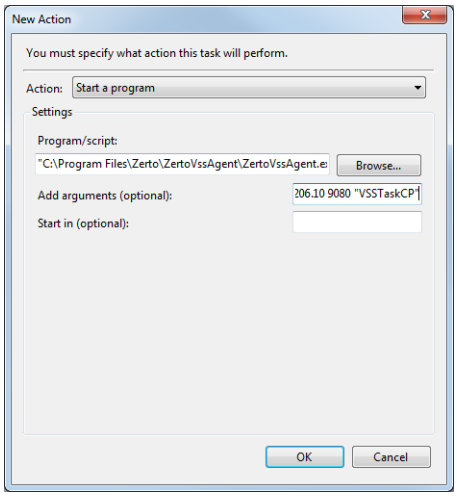

6. Click *OK*.
7. Select the *Settings* tab and make changes as required. Make sure `Stop the task if it runs longer than` is not selected.
8. Click *OK*.

There are certain permissions required for the Windows scheduled task to execute successfully. For example, you may see the following in the event logs:

```
Volume Shadow Copy Service error: Unexpected error querying for the IVssWriterCallback
interface. hr = 0x80070005
```

This is often caused by incorrect security settings in either the writer or requestor process.

If this is the case, the service which runs the Windows Scheduled Task must have `NT AUTHORITY\Network Service` permissions or be using the `SYSTEM` account to run the task. VSS operations are performed as `NT AUTHORITY\Network Service` which is not granted COM access by default on the service assigned to Windows Scheduled Tasks.

The following procedure is only required if the windows scheduled task is using the `Network Services` account.

The correct permissions can be assigned by using the `Component Services` application, accessed by running `dcomcnfg.exe`, in the windows guest.

To set COM permissions for VSS when "Access Denied" errors are received:

1. Run `dcomcnfg.exe`.
   The `Component Services` dialog is displayed.
2. Expand the Component Services node to `My Computer` and right-click to access the *Properties* menu.
   The `My Computer Properties` dialog is displayed.
3. Select the *COM Security* tab and click *Edit Limits* under `Access Permissions`.
4. Add the `NETWORK SERVICE` local access.

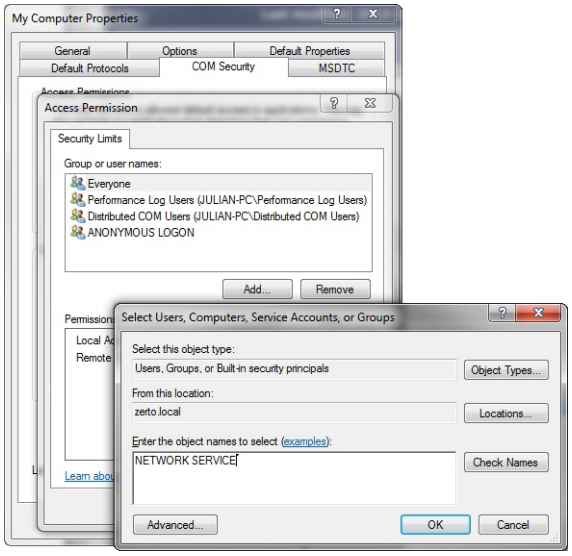

5. Click *OK* and verify that the user is now in the `Access Permission` list.

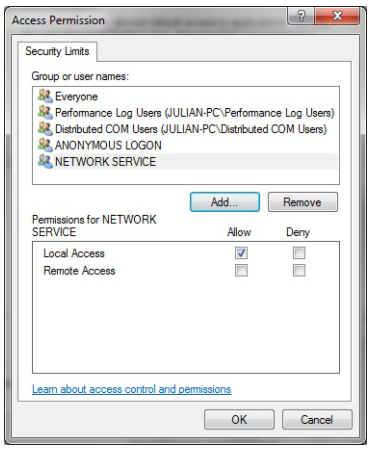

6. Click *OK* to commit these changes.

   `Access Denied` messages should no longer be written in the event viewer for VSS. Additionally, you can grant `Network Service` full control over `HKLM\SYSTEM\CurrentControlSet\Services\VSS\Diag`. You can also check this key `HKLM\SYSTEM\CurrentControlSet\Services\VSS\VssAccessControl` which should at least contain the `DWORD NT Authority\NetworkService` set to value `1`.

   You may also add a new `DWORD` like `DOMAIN\MyZertoServiceUserAccount` and set its value to `1`.

   The Windows Scheduled Task will be created and run by the currently logged in user. After the task is created, Zerto recommends changing this to `NT AUTHORITY\Network Service` permissions and follow the steps to allow the correct permissions as described in "To set COM permissions for VSS when "Access Denied" errors are received:", on page 56.

7. Select the *Triggers* tab and configure a new trigger.

   The `New Trigger` dialog is displayed.

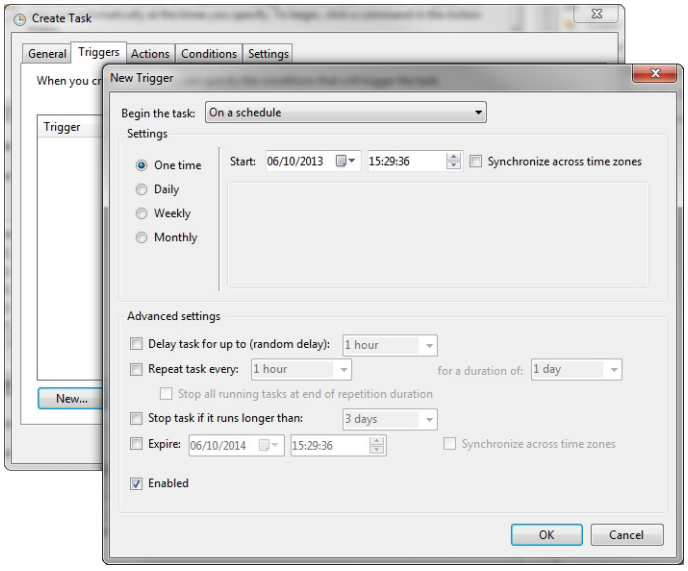

8. Select the *Actions* tab and create a new action to start the ZertoVssAgent with the IP address and port of the Zerto Virtual Manager and the checkpoint to use. For example:

   `C:\Program Files\Zerto\ZertoVssAgent\ZertoVssAgent.exe` and

   `106.18.206.10 9080 106.18.206.10 9080 "VSSTaskCP"`

   That is, with the format: `<protecting_ZVM_IP> 9080 <recovery_ZVM_IP> 9080 "<CP_name>"`

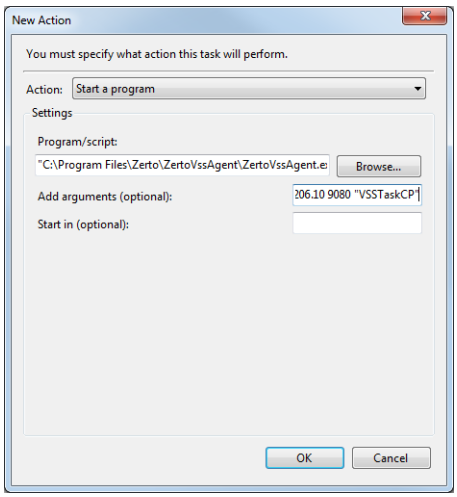

9. Click *OK*.
10. Select the *Settings* tab and make changes as required. Make sure `Stop the task if it runs longer than` is not selected.
11. Click *OK*.

There are certain permissions required for the Windows scheduled task to execute successfully. For example, you may see the following in the event logs:

```
Volume Shadow Copy Service error: Unexpected error querying for the IVssWriterCallback
interface. hr = 0x80070005
```

This is often caused by incorrect security settings in either the writer or requestor process.

If this is the case, the service which runs the Windows Scheduled Task must have `NT AUTHORITY\Network Service` permissions or be using the `SYSTEM` account to run the task. VSS operations are performed as `NT AUTHORITY\Network Service` which is not granted COM access by default on the service assigned to Windows Scheduled Tasks.

The following procedure is only required if the windows scheduled task is using the `Network Services` account.

The correct permissions can be assigned by using the `Component Services` application, accessed by running `dcomcnfg.exe`, in the windows guest.

To set COM permissions for VSS when "Access Denied" errors are received:

1. Run **dcomcnfg.exe**.
   The Component Services dialog is displayed.
2. Expand the Component Services node to `My Computer` and right-click to access the *Properties* menu.
   The My Computer Properties dialog is displayed.
3. Select the **COM Security** tab and under **Access Permissions**, click **Edit Limits**.

4. Add the **NETWORK SERVICE** local access.

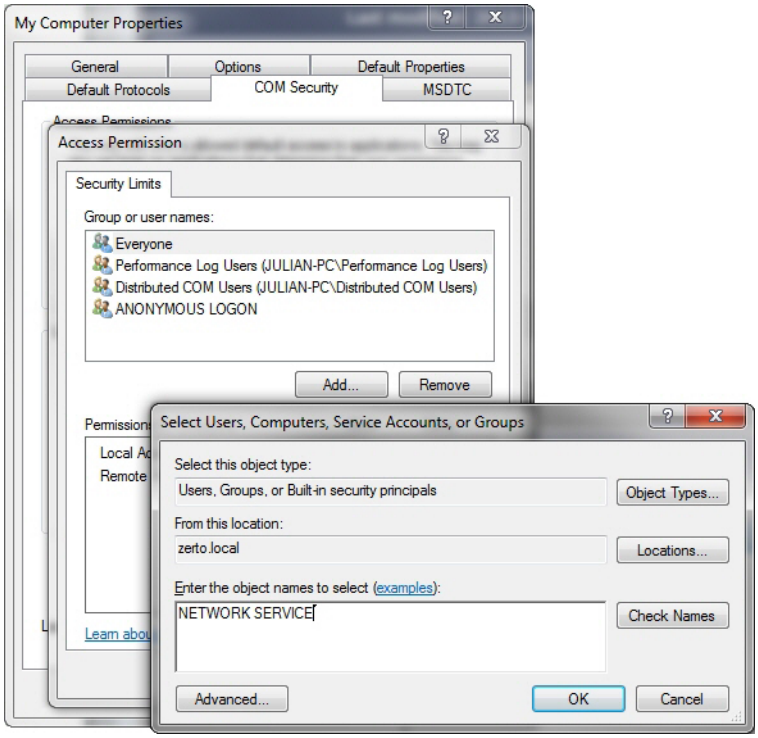

5. Click **OK** and verify that the user is now in the **Access Permission** list.

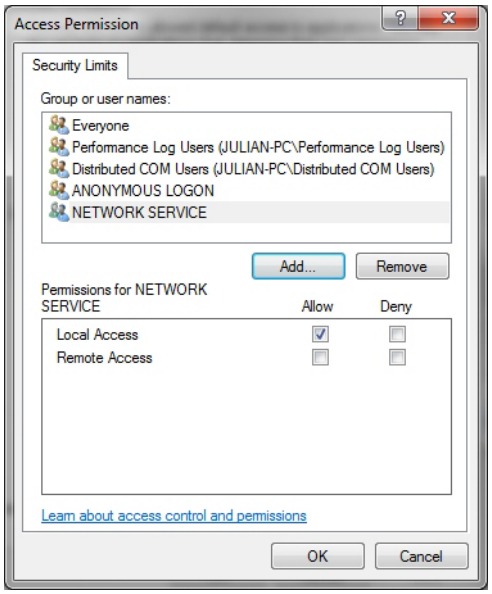

6. Click **OK** to commit these changes.

   `Access Denied` messages should no longer be written in the event viewer for VSS. Additionally, you can grant `Network Service` full control over `HKLM\SYSTEM\CurrentControlSet\Services\VSS\Diag`. You can also check this key `HKLM\SYSTEM\CurrentControlSet\Services\VSS\VssAccessControl` which should at least contain the `DWORD NT Authority\NetworkService` set to value `1`.

   You may also add a new `DWORD` like `DOMAIN\MyZertoServiceUserAccount` and set its value to `1`.

During recovery you can recover to the VSS checkpoint, ensuring both application consistency and that the data is crash-consistent for this virtual machine. For details, refer to "To test failover:" on page 84 and "To initiate a failover:", on page 98.

Changing the Zerto Virtual Manager Used by the ZertoVssAgent

When you install the *ZertoVssAgent*, you specify the Zerto Virtual Manager to use to manage the addition of checkpoints for the virtual machines that uses VSS and that you want to protect in VPGs. You can change the IP and port of the VPG that you specified during the installation either by rerunning the installation and selecting the `Repair ZertoVssAgent` option or by editing IP and port values in the `ZertoVssAgentGUI.exe.conf` file in the folder where the *ZertoVssAgent* is installed.

Running Scripts Before or After Recovering a VPG

Before and after executing a failover, move, or test failover, you can run executable scripts, such as Windows .bat files or PowerShell scripts. A pre-recovery script is always run at the beginning of the recovery operation. A post-recovery script is run after all the virtual machines are powered on at the recovery site.

The scripts must be saved to the machine where the remote Zerto Virtual Manager (ZVM) is installed.

Both pre-recovery and post-recovery scripts are run by the ZVM service on the ZVM machine. The account running the ZVM service is the account that will run the scripts when they are executed.

The scripts can include environment variables that can be included as part of the script itself, or passed to the script as parameters. When the script is passed an environment variable as a parameter, the variable is evaluated before executing the script. The following environment variables are available:

- %ZertoVPGName%: The name of the VPG. If the name includes a space, enclose the variable in double quotes (”). For example, the VPG MyVPG uses the format %ZertoVPGName% but the VPG My VPG uses the format “%ZertoVPGName%”.
- %ZertoOperation%: The operation being run: FailoverBeforeCommit, FailoverRollback, Test, MoveBeforeCommit, MoveRollback. Use the result returned for this variable to limit when the script runs, dependent on the operation. The scripts are run after all the virtual machines are powered on at the recovery site and the variable is set to FailoverBeforeCommit or MoveBeforeCommit. Use FailoverRollback or MoveRollback when rolling back the Failover or Move operation, to undo whatever changes a previous script has done (such as updating the DNS records).
- %ZertoHypervisorManagerIP%: The IP address of the hypervisor manager, VMware vCenter Server or Microsoft SCVMM, where the VPG is recovered.
- %ZertoHypervisorManagerPort%: The port used by the Zerto Virtual Manager to communicate with the hypervisor manager, VMware vCenter Server or Microsoft SCVMM.
- %ZertoForce%: A Boolean value, `Yes/No`, that dictates whether to abort the recovery operation if the script fails. For example, whether to rollback a Move operation when the script fails and returns a non-zero value.

For example, if a specific VPG should not be migrated, the pre-recovery script can determine whether to continue based on the values of the %ZertoOperation% and %ZertoVPGName%.

When specifying scripts in the definition of a VPG, enter values for the Pre-recovery Script and Post-recovery Script:

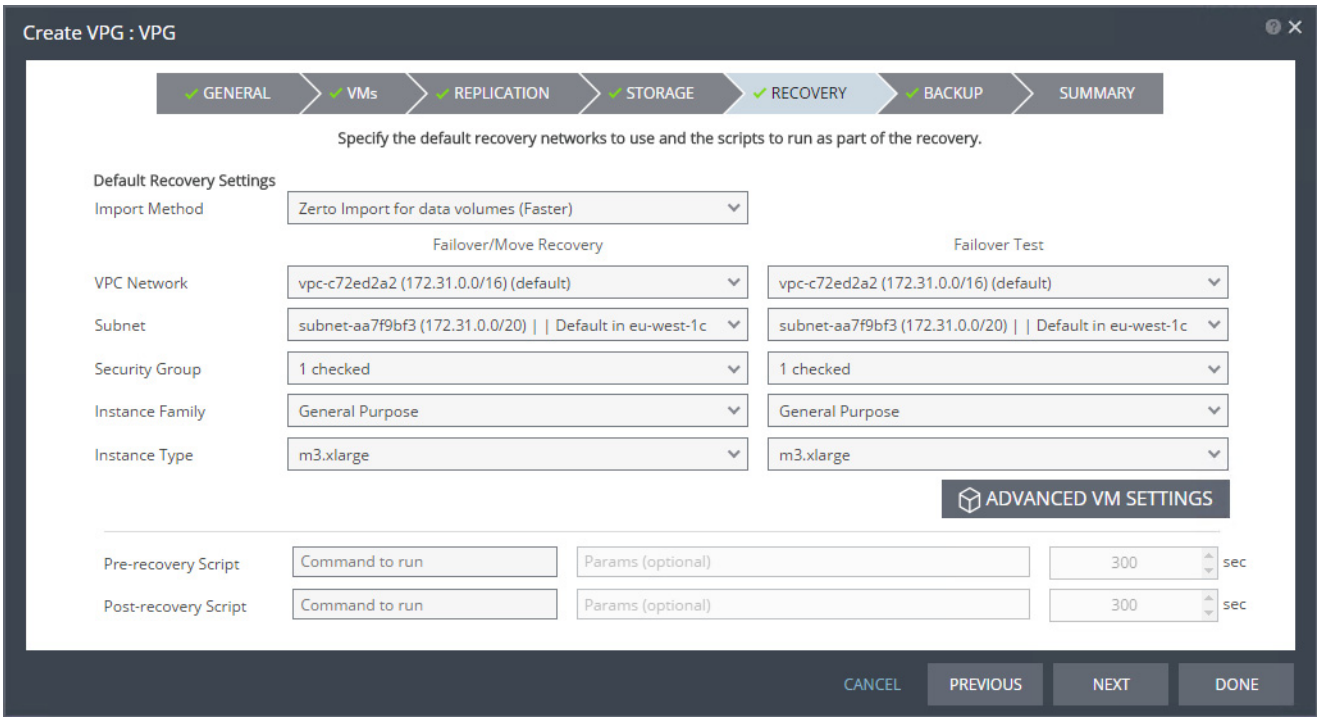

Command to run: The full path of the script to run. The script must be located on the same machine as the Zerto Virtual Manager for the recovery site.

Params: The values of any parameters to pass to the script. Separate parameters with a space.

Timeout (sec): The time-out in seconds for the script to run. If the script runs before executing a failover, move, or test failover and the script fails or a timeout value is reached, an alert is generated and the failover, move, or test failover is not performed. If the script runs after executing a failover, move, or test failover and the timeout value is reached, an alert is generated. The default timeout value is specified in the Site Configuration Advanced Settings dialog.

Creating a Script

There are many ways to create scripts to run before or after recovering a VPG. The following procedure uses a Windows PowerShell file (`.ps1`) or a batch (.bat) file.

To create a script:

1. Create a file on the machine where the Zerto Virtual Manager that manages the recovery is installed.
2. Enter the script that you want to run in the file.
3. Save the file as a Windows PowerShell file (`.ps1`) or batch (`.bat`) file.

When writing a PowerShell script, you can include the environment variables in the script. For example, the following code snippet shows the use of the `%ZertoOperation%` and `%ZertoVPGName%` environment variables:

```
$Operation = $env:ZertoOperation
$VPG = $env:ZertoVPGName
$time = Get-Date

if ($Operation -eq "Test") {
  "$time VPG: $VPG was tested." >> "C:\ZertoScripts\VPG_DR.txt"
}

if ($Operation -eq "FailoverBeforeCommit") {
  "$time Failover before commit was performed. VPG: $VPG" >> "C:\ZertoScripts\VPG_DR.txt"
}

if ($Operation -eq "MoveBeforeCommit"){
  "$time Move before commit was performed. VPG: $VPG" >> "C:\ZertoScripts\VPG_DR.txt"
}
```

Pre-recovery scripts must be saved on the protected site Zerto Virtual Manager machine. Post-recovery scripts must be saved on the recovery site Zerto Virtual Manager machine.

Note: Zerto recommends having both pre- and post-recovery scripts, available on both the protected and recovery Zerto Virtual Manager machines, so that they will work from the protected site.

4. Update `Command` to `run` and `Params` fields for all the VPG definitions that you want to run the script.

   Passing parameters is implemented differently for the two script types. For information about passing command line parameters, refer to the relevant PowerShell or batch file documentation.

Using a BAT File

Windows Batch (.bat) is an executable file that does not require anything in order to run. Update `Command` to `run` and `Params` fields for all the VPG definitions that you want to run the script.

Command to run – *<script_including_path>*

```
C:\ZertoScripts\PostScript.bat
```

Use quotes (") around the path if it includes spaces. The bat file is an executable file and is therefore included in the `Command` to `run` field.

Params – *<Zerto_Params>*, for example:

```
%ZertoOperation% %ZertoVPGName%
```

Using a PowerShell Script

Windows PowerShell scripts require Windows PowerShell (.exe) to execute. To specify a PowerShell script, update `Command` to `run` and `Params` fields for all the VPG definitions that you want to run the script.

Command to run – `powershell.exe`

Params – *<script_including_path> <Zerto_Params>*, for example:

```
C:\ZertoScripts\PostScript.ps1 %ZertoOperation% %ZertoVPGName%
```

Use quotes (") around the path if it includes spaces.

Note: You might have to set the remote signed execution policy. For example, using the following:

```
##PowerCLI requires remote signed execution policy - if this is not enabled,
##it may be enabled here by uncommenting the line below.

##Set-ExecutionPolicy -ExecutionPolicy RemoteSigned -Force
```

Note: Zerto recommends testing both PowerShell and batch scripts by running them from the command line, to ensure that they run correctly.

See also "Example Scripts", on page 61.

Example Scripts

The following script is an example of how to track failover tests.

The following script, `c:\ZertoScripts\TestedVPGs.bat`, writes the VPG name and date to the `ListOfTestedVPGs.txt` file every time a failover test is run:

```
SET isodt=%date:~10,4%-%date:~7,2%-%date:~4,2% %time:~0,2%-%time:~3,2%-%time:~6,2%
IF %1==Test ECHO %2 %isodt% >> c:\ZertoScripts\Results\TestedVPGs.txt
```

Where `%1` is the first parameter in the list of parameters, `%ZertoOperation%`, and `%2` is the second parameter in the list of parameters, `%ZertoVPGName%`.

Note: If the file `TestedVPGs.txt` does not exist it is created, as long as the folder, `c:\ZertoScripts\Results`, exists.

Exporting and Importing VPG Definitions

You can save VPG definitions to an external file and import these definitions back to Zerto Virtual Replication, for example, exporting the settings before uninstalling a version of Zerto Virtual Replication and importing the settings after reinstalling Zerto Virtual Replication.

Note: Zerto Virtual Replication regularly exports settings to the *Zerto_Installation_Folder*`\Zerto Virtual Replication\ExportedSettings` folder. You can use one of these exported files instead of creating a new export file. The default location of *Zerto_Installation_Folder* is `C:\Program Files\Zerto`.

To export VPG settings:

1. Open the *Zerto Diagnostics* application. For example, via *Start* > *Programs* > *Zerto Virtual Replication* > *Zerto Diagnostics*.

   The `Zerto Virtual Replication Diagnostics` menu dialog is displayed.

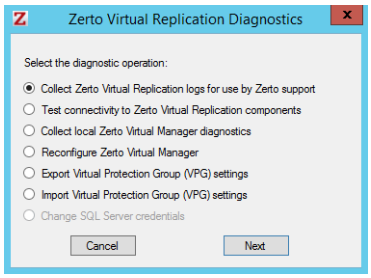

2. Select the *Export Protection Group Settings* option and click *Next*.

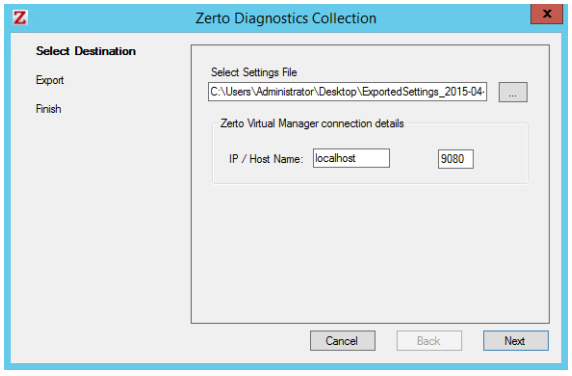

3. Select the destination for the file to contain exported settings and specify the Zerto Virtual Manager IP address and port where the VPGs are protecting virtual machines.
4. Click *Next*.

   The list of exported VPGs is displayed.

5. Click *Done*.

Note: If you are uninstalling Zerto Virtual Replication, the VPGs are deleted.

To import VPG settings:

1. Click *Start* > *Programs* > *Zerto Virtual Replication* > *Zerto Diagnostics*.

   The `Zerto Virtual Replication Diagnostics` menu dialog is displayed.
2. Select the *Import Protection Group Settings* option.
3. Click *Next*.

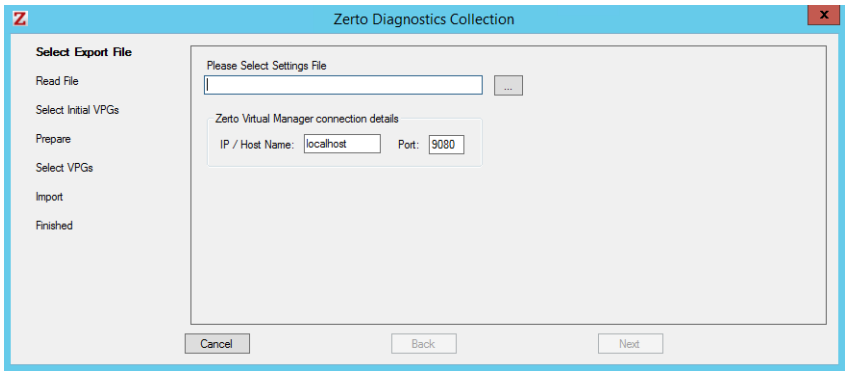

4. Select the file previously exported and enter the Zerto Virtual Manager IP address and port specified when exporting the VPGs.
5. Click *Next*.

   The list of exported VPGs is displayed.

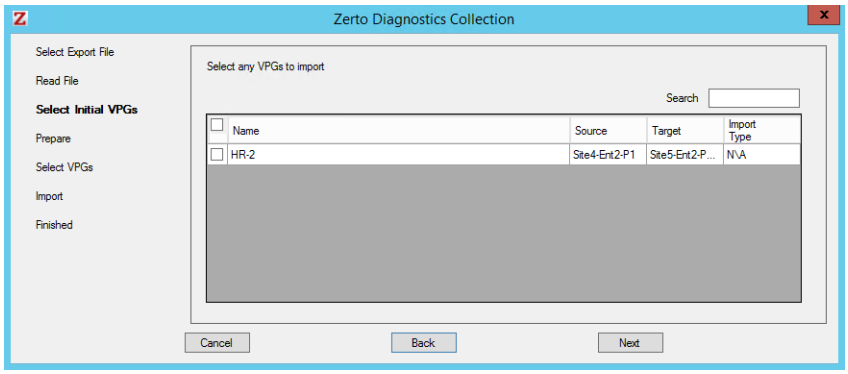

6. Select the VPGs to import. Only VPGs with names that are not already defined can be imported. VPGs in the import files with the same name as an existing VPG are disabled.
7. Click *Next*.

   The list of imported VPGs is displayed. If the VPG could not be imported, the reason for the failure is specified.

   Note: If a host was removed from and then re-added to the environment it is advisable to wait approximately 5 minutes from when the host was re-added before performing the import of the VPGs.
8. Click *Done*.

VPG Statuses and Synchronization Triggers

During normal operations the VPG status can change. For example, a change can be made to the VPG definition, or an operation such as move or failover is performed on the VPG, or an external event impacts the system such as the WAN going down. When the status changes, resulting in the VPG being synchronized, for example with a `Delta Sync`, the estimated time

to complete the synchronization is displayed under the VPG status, and if relevant, the synchronization trigger, such as `Network Congestion`.

See also:

- "VPG Statuses", on page 63
- "VPG Synchronization Triggers", on page 67

VPG Statuses

The following statuses are displayed:

STATUS	SUBSTATUS	COMMENT
Deleting	Deleting the VPG	
	VPG waiting to be removed	
Failing Over	Committing Failover	The VPG is being failed over.
	Failing over – Before commit	A VPG being failed over is in the initial stage, before committing the failover.
	Rolling back Failover	The failover is being rolled back to prior to the failover.
History Not Meeting SLA	See *Not Meeting SLA*, below.	The VPG is meeting the RPO SLA setting.
Initializing	Creating VPG	
	Initial Sync	
	Syncing	
	Volume Initial Sync	
Meeting SLA	—	
	Bitmap Syncing	
	Delta Syncing (When Force Sync is applied)	
	Recovery is Possible	After a rollback.
Moving	Committing Move	
	Moving – Before commit	
	Promoting	
	Rolling back Move	
Not Meeting SLA	Delta Sync (When Force Sync is not applied)	This status means that the VPG is not meeting the journal history nor RPO SLA settings.
	Delta Syncing a volume	
	Error	
	Needs configuration	
	Site disconnection	
	Site disconnection. No checkpoints	
	VM not protected	
	VPG has no VMs	
Recovered	—	The VPG has been recovered.
RPO Not Meeting SLA	See *Not Meeting SLA*, above.	The VPG is meeting the journal history SLA setting.

The following provides a full description of the sub-statuses are displayed:

SUBSTATUS	DESCRIPTION
Backing Up	An offsite backup is running.
Bitmap Syncing[1]	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. The bitmap is small and scales dynamically, containing references to the areas of the protected disk that have changed but not the actual I/O. The bitmap is stored locally on the VRA within the available resources. For example, when a VRA goes down and is then rebooted. When required, Zerto Virtual Replication starts to maintain a smart bitmap in memory, to track and record storage areas that change. When the issue that caused the bitmap sync is resolved, the bitmap is used to check updates to the protected disks and send any updates to the recovery site. A bitmap sync occurs when any of the following conditions occur: ■ Synchronization after WAN failure or when the load over the WAN is too great for the WAN to handle, in which case the VPGs with the lower priorities will be the first to enter a bitmap sync. ■ When there is storage congestion at the recovery site, for example when the VRA at the recovery site cannot handle all the writes received from the protected site in a timely fashion. ■ When the VRA service at the recovery site goes down and is then restarted, for example during a Zerto Virtual Replication upgrade During the synchronization, new checkpoints are not added to the journal but recovery operations are still possible, assuming there are valid checkpoints in the journal. If a disaster occurs requiring a failover during a bitmap synchronization, the VPG status changes to `Recovery Possible` and you can recover to the last checkpoint written to the journal. For synchronization to work, the protected virtual machines must be powered on so that the VRA has an active IO stack, which is only available when the virtual machine is powered on. **Note:** If the synchronization takes longer than the configured history, all the checkpoints in the journal can be lost, preventing a failover from being performed. For the resolution of this situation, see "To configure disaster recovery policies:", on page 78.
Committing Failover	Failing over the VPG.
Committing Move	Completing the move, including removing the protected virtual machines.
Creating VPG	The VPG is being created based on the saved definition.
Deleting the VPG	Deleting the VPG.

SUBSTATUS	DESCRIPTION
Delta Syncing[1]	The *Delta Sync* uses a checksum comparison to minimize the use of network resources. A Delta Sync is used when the protected virtual machine disks and the recovery disks should already be synchronized, except for a possible few changes to the protected disks, for example: ■ After a source VRA upgrade of a major release on the protected site: Depending on the nature of the upgrade, a VRA upgrade on the protected side may trigger either a Delta Sync or a Bitmap Sync. See the version release notes to determine if a sync will be triggered with a VRA upgrade. ■ A Force Sync operation was manually initiated on the VPG. ■ A host protecting virtual machines was restarted and the protected virtual machines on the host had not been moved to other hosts in the cluster or a protected virtual machine was moved to another host without a VRA, and then moved back to the original host. For synchronization to work, the protected virtual machines must be powered on so that the VRA has an active IO stack, which is only available when the virtual machine is powered on. During the synchronization, new checkpoints are not added to the journal but recovery operations are still possible, assuming there are valid checkpoints in the journal. If a disaster occurs requiring a failover during a delta synchronization, you can recover to the last checkpoint written to the journal.
Delta syncing a volume[1]	Synchronization when only delta changes for a volume needs synchronizing. For synchronization to work, the protected virtual machines must be powered on so that the VRA has an active IO stack, which is only available when the virtual machine is powered on. During the synchronization, new checkpoints are not added to the journal but recovery operations are still possible, assuming there are valid checkpoints in the journal. If a disaster occurs requiring a failover during a delta volume synchronization, you can recover to the last checkpoint written to the journal.
Error	Problem situation, for example, when a ZVM is disconnected from a VRA used to protect virtual machines. The VPG cannot be recovered until the problem is resolved,
Failing over - Before commit	Preparing and checking the VPG virtual machines in the recovery site.
Full Syncing[1]	Full synchronization to ensure that the protected disks and recovery disks are the same after some change to the system. This type of sync is the same as an `Initial Sync` but occurs after protection started. In general, this type of sync should not happen. For synchronization to work, the protected virtual machines must be powered on so that the VRA has an active IO stack, which is only available when the virtual machine is powered on. During the synchronization, new checkpoints are not added to the journal. Also, recovery operations are not possible.
Full syncing a volume[1]	Synchronization when a full synchronization is required on a single volume. For synchronization to work, the protected virtual machines must be powered on so that the VRA has an active IO stack, which is only available when the virtual machine is powered on. During the synchronization, new checkpoints are not added to the journal. Also, recovery operations are not possible.

SUBSTATUS	DESCRIPTION
Initial Sync[1]	Synchronization performed after creating the VPG to ensure that the protected disks and recovery disks are the same. Recovery operations cannot occur until after the initial synchronization has completed. For synchronization to work, the protected virtual machines must be powered on so that the VRA has an active IO stack, which is only available when the virtual machine is powered on. Adding a virtual machine to a VPG is equivalent to creating a new VPG and an initial synchronization is performed. In this case, any checkpoints in the journal become unusable and only new checkpoints added after the initial synchronization completes can be used in a recovery. The data in the journal however remains and is promoted to the recovered virtual machine as part of a recovery procedure.
Journal storage error	There was an I/O error to the journal. For example, if the journal was full and the size was increased. Once the problem is resolved a synchronization is required.
Moving - Before commit	Preparing and checking the VPG virtual machines in the recovery site.
Needs Configuration	One or more configuration settings are missing.
Promoting	Updating recovered virtual machines in the VPG with data from the journal.
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**).
Recovery storage error	There was an I/O error to the recovery storage. For example, the storage is almost full or the virtual machines are turned off and the recovery disks are inaccessible.
Recovery storage profile error	The storage profile in the recovery site specified to be used by the VPG cannot be found.
Rolling back	Rolling back to an initial status, for example, after canceling a cloning operation on the VPG.
Rolling back Failover	Rolling back a Failover operation before committing it.
Rolling back Move	Rolling back a Move operation before committing it.
Site disconnection	Communication with the Zerto Virtual Manager at the remote, recovery, site is down so continuing protection is halted (compare with **Recovery is possible**).
Site disconnection. No checkpoints	Communication with the Zerto Virtual Manager at the remote, recovery, site is down and there are no checkpoints to use to recover the VPG at the recovery site.
Syncing	Status while type of synchronization is being evaluated.
User paused protection	Protection is paused to enable solving a Journal disk space problem, for example, by increasing the disk size or cloning the VPG.
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.
Volume Initial Sync[1]	Synchronization when a full synchronization is required on a single volume, for example, when changing the target storage or adding a virtual machine to the VPG. For synchronization to work, the protected virtual machines must be powered on so that the VRA has an active IO stack, which is only available when the virtual machine is powered on. During synchronization, new checkpoints are not added to the journal. Also, recovery operations are not possible during a `Volume Initial Sync`.
VPG has no VMs	A configured VPG where the virtual machines have been removed from it, for example when changing both the storage and host for the virtual machines in the VPG, causes the VPG to be recreated.

SUBSTATUS	DESCRIPTION
VPG waiting to be removed	An attempt to remove the VPG failed and it must be forcibly removed. For details, see "Deleting a VPG When the Status is Deleting", on page 47.
Zerto Virtual Manager paused protection	Protection is paused to enable solving a Journal disk space problem, for example, by increasing the disk size or cloning the VPG.

1. Synchronization after a recovery starts after the promotion of data from the journal to the virtual machine disks ends. Thus, synchronization of virtual machines can start at different times, dependent on when the promotion for the virtual machine ends. All synchronizations are done in parallel, whether a delta sync or full sync, etc.

VPG Synchronization Triggers

The following synchronization triggers can be applied:

TRIGGER	DESCRIPTION
Force Sync	The user requested to synchronize the VPG, as described in "Forcing the Synchronization of a VPG", on page 46.
Network Congestion	The network bandwidth is not wide enough to handle all the data, causing some of the data to be backed up.
Protected Storage Error	An I/O error occurred to a protected virtual machine, after the data was sent to the recovery side.
Protected VRA Congestion	The host where the VRA is installed is highly loaded: many updates are made to the protected machines at the same time, causing a time lapse before the updates are passed to the recovery site.
Recovery or Journal Storage Error	There was an I/O error either to the recovery storage or journal, for example if the journal was full and the size was increased. Once the problem is resolved a synchronization is required.
Recovery Storage Congestion	The recovery storage is being written to a lot, causing a delay for some of the data passed from the protected site to be written to disk.
Recovery VRA Communication Problem	A network error, such as the network being down for a period, requires a synchronization of the VPG between the two sites, for example a bitmap sync.
VPG Configuration Changed	The configuration of the VPG changed resulting in a synchronization being required. For example, the size of the journal was changed.

CHAPTER 7: MANAGING A ZERTO VIRTUAL MANAGER

The Zerto Virtual Manager runs as a Windows service and connects to Zerto Virtual Replication components, such as the VRA.

The following topics are described in this chapter:

Check Connectivity Between Zerto Virtual Replication Components

If you think that there are connectivity problems to or from a Zerto Virtual Manager, you can use the Zerto diagnostics utility to check the connectivity.

To check connectivity between Zerto Virtual Manager components:

1. Open the *Zerto Diagnostics* application. For example, via *Start > Programs > Zerto Virtual Replication > Zerto Diagnostics*.

   The *Zerto Virtual Replication Diagnostics* menu dialog is displayed.

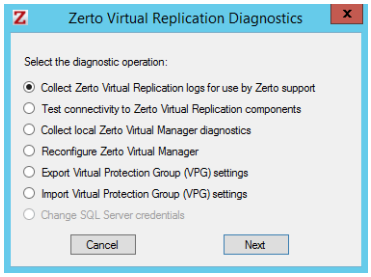

2. Select the *Test Connectivity to Zerto Virtual Replication components* option and click *Next*.

   The *IP Connectivity* dialog is displayed.

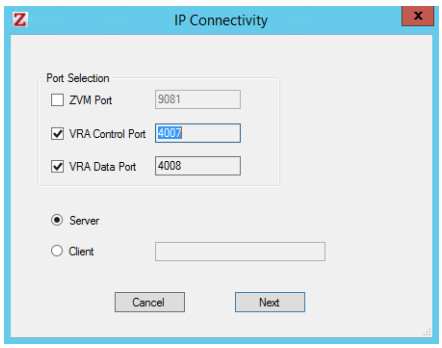

   You can use this dialog to check the following:
   - TCP communication between the Zerto Virtual Managers (ZVMs) on the protected and recovery sites. The default port, specified during installation, is 9081.
   - Communication between VRAs on the protected and recovery sites, via the control port and the data port.
3. Select the connectivity you want to test and in the case of the Zerto Virtual Manager (ZVM), specify the TCP communication port specified during the installation, if the default port, 9081, was changed.
4. Specify the type of test to perform:

   Server – Test for incoming communication.

Client – Test for outgoing communication. Specify the IP address of the receiving Zerto Virtual Manager.

5. Click *Next* to test the specified connectivity.

   The Server option listens for communication from a paired VRA. Stop listening by clicking *Stop*.

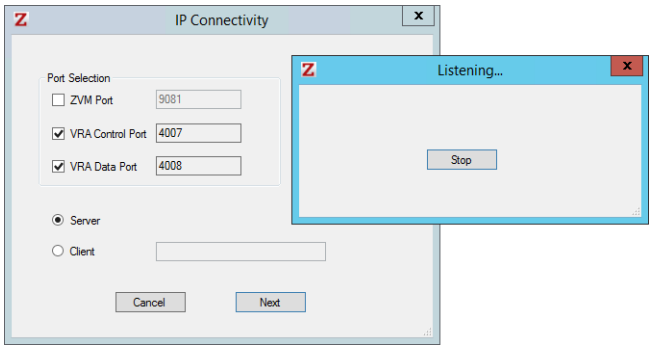

   The Client options tests the client; on completion a result dialog is displayed.

6. Click *Stop* (server test) or *OK* (client test) to return to the `Zerto Virtual Replication Diagnostics` dialog.

Reconfiguring the Zerto Virtual Manager Setup

When installing Zerto Cloud Appliance, you provide the IP address of the machine on which you are installing it. This is where the Zerto Virtual Manager runs and displays the Zerto User Interface.

You can change this IP address if necessary, using the Zerto Virtual Replication Diagnostics utility.

To reconfigure the Zerto Virtual Manager:

1. Click *Start* > *Programs* > *Zerto Virtual Replication* > *Zerto Diagnostics*.

   The *Zerto Virtual Replication Diagnostics* menu dialog is displayed.

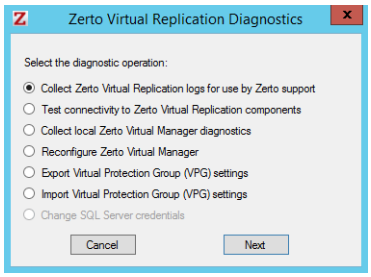

2. Select the *Reconfigure Zerto Virtual Manager* option and click *Next*.

   The settings for the connection to the Zerto Cloud Appliance are displayed.

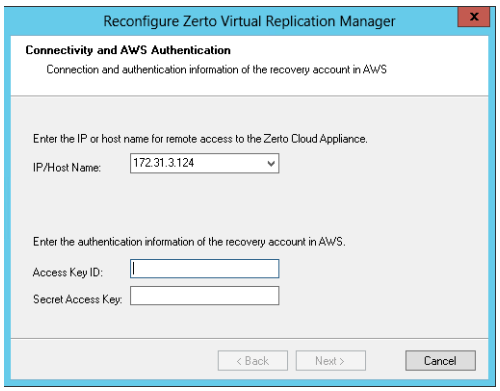

3. Change the IP/host name, and access key ID and secret access key, if necessary.

IP / Host Name – The IP address or host name of the machine on which the Zerto Cloud Appliance is installed.

Access Key ID – A unique identifier that is associated with a secret access key.

Secret Access Key – A key that is used with the access key ID.

4. Click *Next*.

   The the Zerto Virtual Manager is reconfigured to use the new information.

5. Click *Finish*.

If you changed the IP address of the Zerto Virtual Manager, you must unpair the AWS and protected sites, and then pair the sites again.

Reconfiguring the Microsoft SQL Server Database Used by the Zerto Virtual Manager

When installing Zerto Virtual Replication, you can specify a Microsoft SQL Server database to use by the Zerto Virtual Manager. If the access to this database changes, you can change the access in the Zerto Virtual Manager.

To reconfigure the access to the Zerto Virtual Manager database:

1. Click *Start* > *Programs* > *Zerto Virtual Replication* > *Zerto Diagnostics*.

   The *Zerto Virtual Replication Diagnostics* menu dialog is displayed.

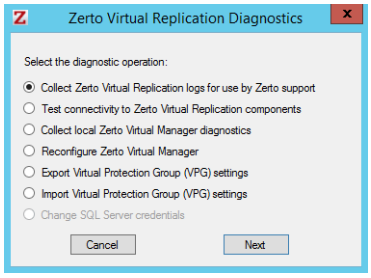

2. Select the *Change SQL Server Credentials* option and click *Next*.

   The installation settings for the SQL Server are displayed. Change the IP and username and password if necessary.

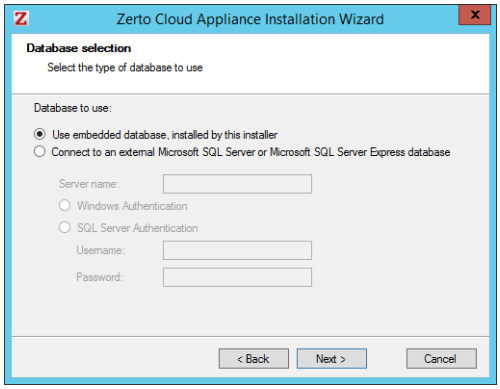

Server Name – The domain name and server instance to connect to, with the format `<server_name>\<instance_name>` or `<Server_IP>\<instance_name>`.

Specify either of the following authentication options:

Windows Authentication – Use Windows authentication. This option is only enabled if a specific service user account was specified in the previous *Service User* dialog, in which case the service account name and password are used.

SQL Server Authentication – Use SQL Server authentication.

User Name – The user name for SQL Server database.

Password – A valid password for the given user name.

3. Click *Next* to the end of the wizard and then click *Finish*.

The Zerto Virtual Manager service is restarted using the new credentials.

Replacing the SSL Certificate

The communication between the Zerto Virtual Manager and the user interface uses HTTPS. On the first login to the Zerto User Interface you must install a security certificate in order to be able to continue working without each login requiring acceptance of the security.

If you want to replace the SSL certificate, perform the procedure described in "To reconfigure the Zerto Virtual Manager:", on page 69 and select a new SSL certificate when the dialog for Zerto Virtual Manager setup is displayed:

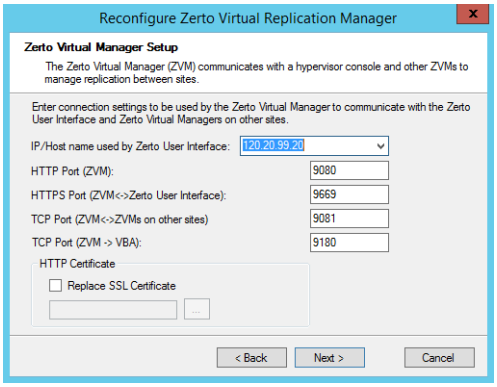

HTTP Certificate – Check `Replace SSL Certificate` and browse for a replacement certificate.

Pair to Another Site and Unpairing Sites

See the following sections:

- "Pair to Another Site", below
- "Unpairing Sites", on page 72

Pair to Another Site

You can pair to any site where Zerto Virtual Replication is installed.

To pair to a site:

1. In the Zerto User Interface, in the *SITES* tab click *PAIR*.
   The *Add Site* dialog is displayed.

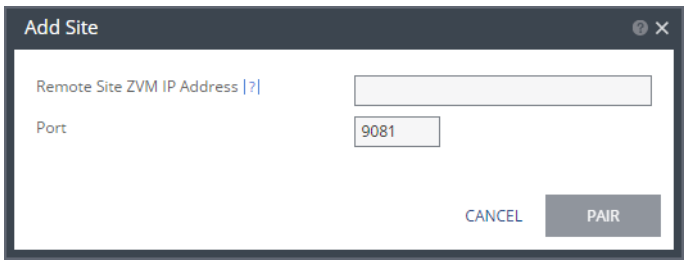

2. Specify the following:
   Remote Site ZVM IP Address: IP address or host name of the remote site Zerto Virtual Manager to pair to.

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

3. Click *PAIR*.

   The sites are paired, meaning that the Zerto Virtual Manager for the local hypervisor site is connected to the Zerto Virtual Manager at the remote hypervisor site.

Unpairing Sites

You can unpair any two sites that are paired to each other.

IMPORTANT: if there is a VPG on either of the sites you are unpairing, the VPGs will be **deleted**.

To unpair two sites:

1. In the Zerto User Interface, in the *SITES* tab, select the site which you want to unpair.
2. Click *UNPAIR*.

   A message appears warning the user that the sites are about to unpair.

   If there are either protected or recovered VPGs on the paired sites, a message appears warning the user that the VPGs will be deleted.
3. To unpair, click CONTINUE.

   The sites are no longer paired. If there are VPGs on either site, they are deleted.

CHAPTER 8: OVERVIEW OF DISASTER RECOVERY OPERATIONS

Zerto Virtual Replication provides a number of operations to recover virtual machines at the remote site. This chapter describes these operations. The following topics are described in this chapter:

The Failover Test Operation

Use the *Failover Test* operation to test that during recovery the virtual machines are correctly replicated in AWS.

The Failover Test operation creates test virtual machines in a sandbox, using the test network specified in the VPG definition as opposed to a production network, to a specified point-in-time, using the buckets in S3 managed by the VRA. For details, see "Testing Recovery to AWS", on page 83.

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

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

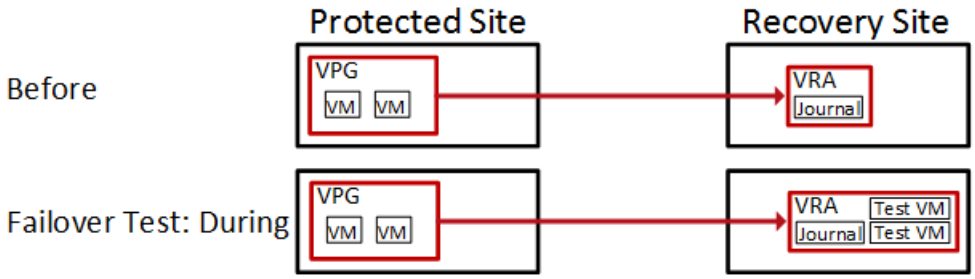

The Move Operation

Use the *Move* operation to transfer protected virtual machines from the protected site to AWS in a planned migration.

When you perform a planned migration of the virtual machines to AWS, Zerto Virtual Replication assumes that both sites are healthy and that you planned to relocate the virtual machines in an orderly fashion. For details, see "Migrating a VPG to AWS", on page 91.

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

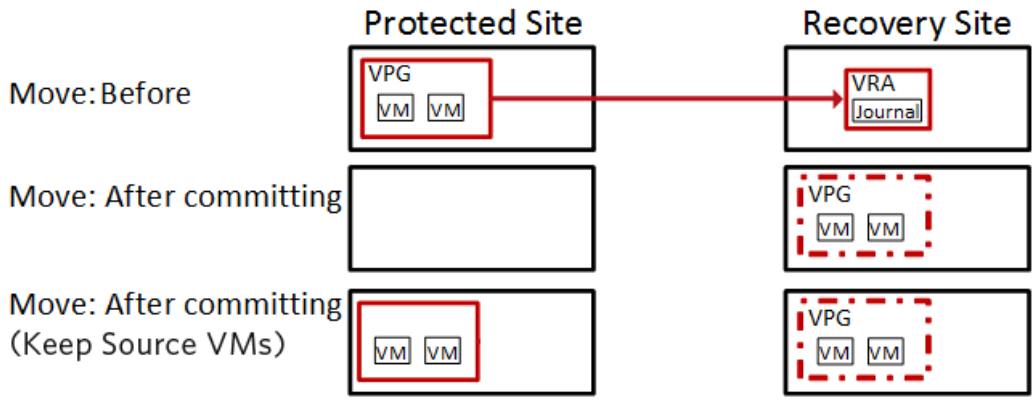

Note: The Move operation leaves the VPG in a *Recovered* state.

The Failover Operation

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

When you set up a failover you always specify a checkpoint to which you want to recover the virtual machines. When you select a checkpoint – either the last auto-generated checkpoint, an earlier checkpoint, or a user-defined checkpoint – Zerto Virtual Replication makes sure that virtual machines in AWS are recovered to this specified point-in-time. For details, see "Managing Failover to AWS", on page 97.

Note: To identify the checkpoint to use, you can perform a number of test failovers, each to a different checkpoint.

Failback after the Original Site is Operational

After completing a failover, when the original site is back up and running you can failback the recovered virtual machines back again.

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

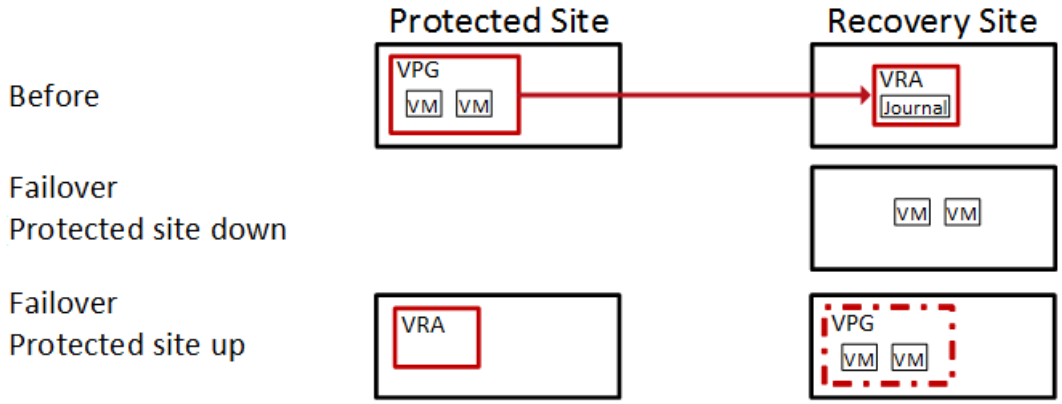

Note: The Failover operation leaves the VPG in a *Recovered* state.

The Restore File Operation

Use the *Restore File* operation to recover individual files and folders from the recovery site.

You can recover specific files and folders from the recovery site for virtual machines that are being protected by Zerto Virtual Replication and running Windows operating systems. You can recover the files and folders from a specific point-in-time. For details, see "Recovering Files and Folders", on page 107.

CHAPTER 9: ADVANCED SITE CONFIGURATION

There are a number of configuration tasks that you can perform, some of which should be done as part of the initial site configuration.

The following topics are described in this chapter:

Site Settings

The *Site Settings* dialog enables configuring various site settings. These include the maximum bandwidth that Zerto Virtual Replication uses between the protected and recovery sites, default script timeout, and protection policies such as the commit policy for a failover or move operation.

To specify site settings:

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

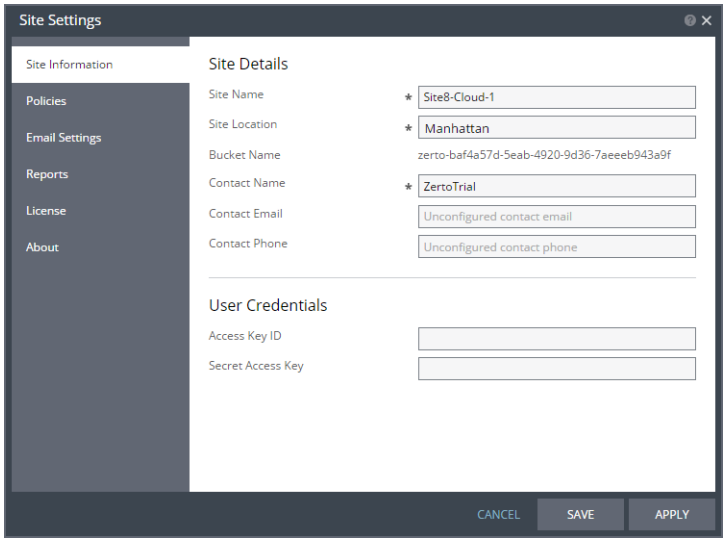

2. Make any required changes to the settings, click *SAVE* and then *APPLY*. The following settings can be defined:
   - "Editing Information About a Site", below
   - "Defining Site Policies", on page 78
   - "Configuring Email Settings", on page 79
   - "Defining the Resource Report Sampling Period", on page 80

Licensing is described in "Seeing What is Licensed", on page 80.

Editing Information About a Site

You provide information about the site during installation, to make it easier to identify the site in the in the user interface and to identify the contact person at the site. After installation you can updated these settings.

In the Zerto User Interface, site information is displayed at the top of the display.

To update information about the local site:

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.

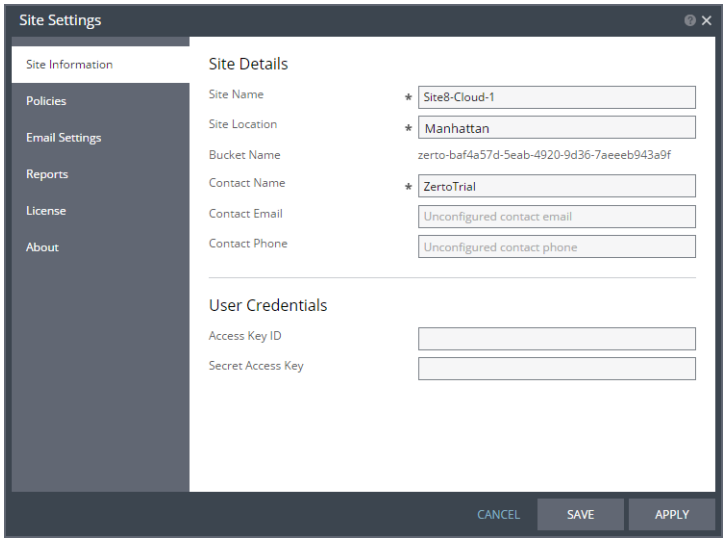

2. Define general information about the site.
   Site Name – The name used to identify the site. Mandatory.
   Site Location – Information such as the address of the site or a significant name to identify it. Mandatory.
   Bucket Name – The name of the bucket that was created when Zerto Virtual Replication was installed. This cannot be changed.
   Contact Name – The name of the person to contact if a need arises. Mandatory.
   Contact Email – An email address to use if a need arises.
   Contact Phone – A phone number to use if a need arises.
3. If the credentials to access AWS from the Zerto Virtual Manager change, specify the new credentials:
   Access Key ID – A unique identifier that is associated with a secret access key.
   Secret Access Key – A key that is used with the access key ID.
4. Click *APPLY* or *SAVE*.

Defining Site Policies

You can set default recovery and replication policies.

Configuring Disaster Recovery Policies

To configure disaster recovery policies:

1. Click *Policies*.

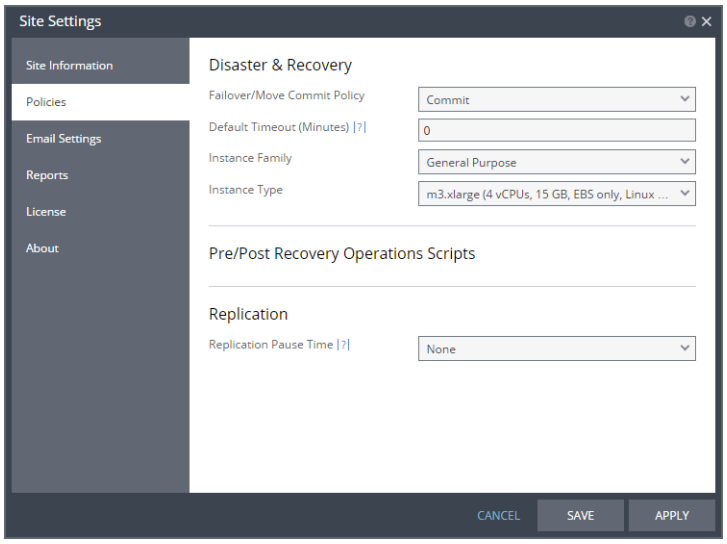

2. Choose the `Failover/Move Commit Policy` to use during a failover or move operation, described in "Initiating a Failover", on page 98 and "Moving Protected Virtual Machines to a Remote Site", on page 93 respectively. The following options are available:

   None – The failover or move operation must be manually committed or rolled back by the user.

   Commit – After the time specified in the `Default Timeout` field the failover or move operation is committed, unless manually committed or rolled back by the user before the time-out value is reached. During the specified time you can check the recovered VPG virtual machines.

   Rollback – After the time specified in the `Default Timeout` field the failover or move operation is rolled back, unless manually committed or rolled back by the user before the time-out value is reached. During the specified time you can check the recovered VPG virtual machines.

   The value set here applies as the default for all failover or move operations from this point on but can be changed when defining a failover or move operation.

3. Specify the `Default Timeout` after which a `Commit` or `Rollback` commit policy is performed. A value of zero indicates that the system will automatically perform the commit policy, without waiting for any user interaction.
4. Choose the `Instance Family` from which to select the type. AWS instance families are optimized for different types of applications.
5. Choose the `Instance Type` within the instance family, to assign to recovered instances. Different types within an instance family vary primarily in vCPU, ECU, RAM, and local storage size. The price per instance is directly related to the instance size.
6. Choose the `Replication Pause Time`, which is the time to pause when the journal might have problems, resulting in the loss of all checkpoints, for example, when the datastore for the journal is near to being full.

   The replication pause time is the amount of time that the transfer of data from the protected site to the journal on the recovery site is paused. This time can then be used by the administrator to resolve the issue, for example by cloning the virtual machines in the VPG, described in "Cloning Protected Virtual Machines to the Remote Site", on page 103. The value set here is applied to existing and new VPGs.
7. Click *APPLY* or *SAVE*.

Configuring Email Settings

You can configure Zerto Virtual Replication alerts to be sent to an email address, so as to be better informed when an alert occurs and backups are run.

Email Settings

To configure email settings:

1. Click *Email Settings*.

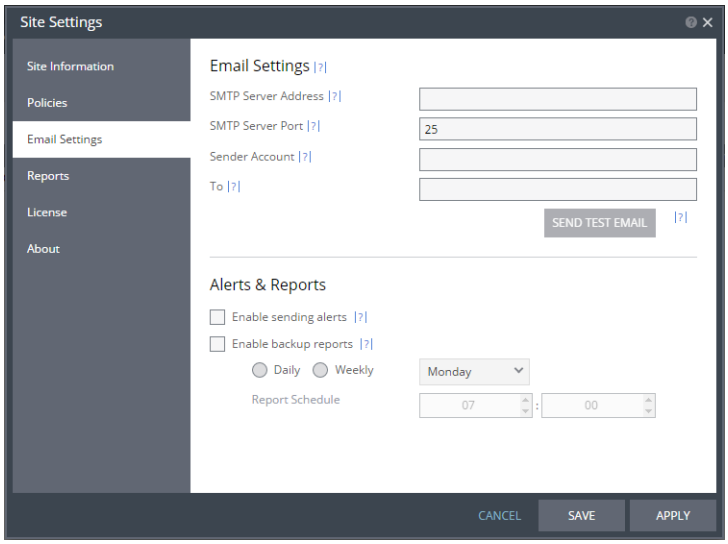

2. Specify the `SMTP server Address`. The Zerto Virtual Manager must be able to reach this address.
3. If the `SMTP Server Port` was changed from the default, 25, specify the port number.
4. Specify a valid email address for the email sender name in the `Sender Account` field.
5. Specify a valid email address where you want to send the email in the `To` field.

   You can test that the email notification is set up correctly by clicking *SEND TEST EMAIL*. A test email is sent to the email address specified in the `To` field.
6. Click *APPLY* or *SAVE*.

Alerts and Reports

You can configure when to send alerts and backup reports.

To configure when to send emails about alerts and backups:

1. To send an email when an alert is issued, check `Enable sending alerts`.
2. To send an email with a backup report, check `Enable backup reports`.
3. Specify whether you want a backup report sent daily or weekly.

   Daily – Send a daily backup report

   Weekly – Send a weekly backup report. Select the day of the week from the dropdown list.
4. Specify day of the week and the time of day to send the backup report.
5. Click *APPLY* or *SAVE*.

Defining the Resource Report Sampling Period

Specify when you want to take resource samples to identify resource usage, either daily at a specific hour and minute or hourly at a specific minute within each hour.

1. Click *Reports*.

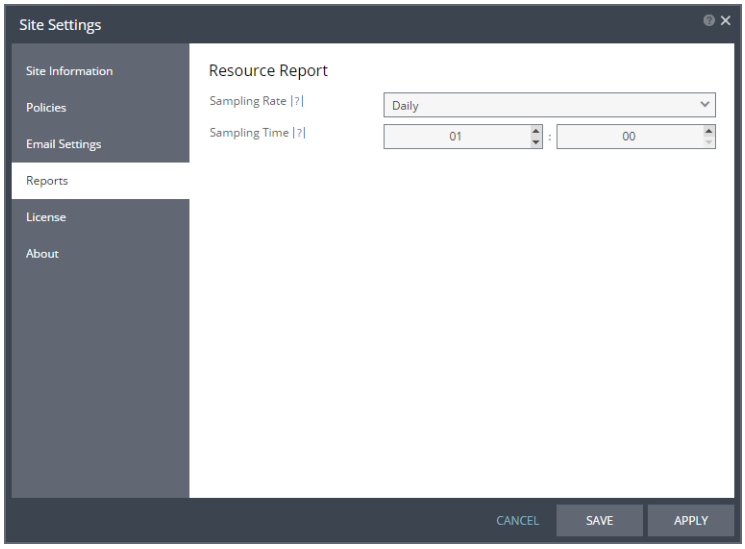

2. Choose the `Sampling Rate`.
3. Choose the `Sampling Time`.

   If you set the daily time to be 12:00, you will get a sample taken at noon every day. Collecting a sample hourly provides a higher resolution picture of replication traffic than if collected daily.
4. Click *APPLY* or *SAVE*.

Information is saved for 90 days when the sampling period is hourly and for one year when the sampling period is daily.

These samples are used to generate resource reports as described in "Zerto Virtual Replication Reports", on page 127.

Seeing What is Licensed

The Zerto license includes information such as the number of virtual machines that can be protected and the license expiry date. You can see these details in the *Site Settings* > *License* dialog.

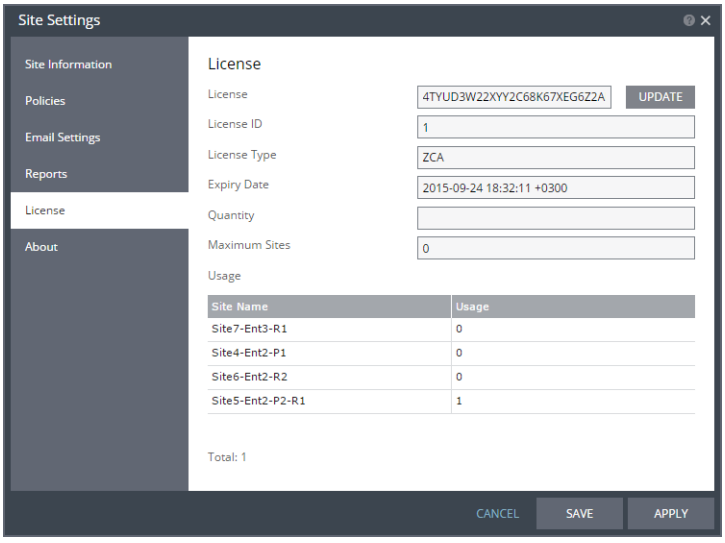

The Zerto license includes the following details:

License – The license key itself.

License ID – An identifier for the license.

License Type – What is licensed: whether the license restricts the number of virtual machines that can be protected or the number of sockets used.

Expiry Date – The license expiry date.

Quantity – The maximum amount licensed, either virtual machines or sockets, based on the license type. If blank, the quantity is unlimited.

Maximum Sites – The maximum number of sites allowed.

Usage – The sites using the license and number of protected virtual machines in each site.

A warning is generated when either the license expires or more than the licensed number of virtual machines are being protected. Protection continues but the license should be updated. After getting a new license key you can update Zerto Virtual Replication with this key.

To update a license key:

1. In the Zerto User Interface, in the top right of the header click *SETTING* ( ) and select *Site Settings*.
   The *Site Settings* dialog is displayed.
2. Click *License*.
3. Enter a valid license key and click *APPLY* or *SAVE*.

The license is updated on the local site and the paired remote sites.

About the Zerto Virtual Replication Version

You can see details about the version of Zerto Virtual Replication being run and specify whether the version can be automatically updated when new VMware vSphere versions are released, without the need to upgrade to a later version of Zerto Virtual Replication. This functionality is the Zerto CALLHOME feature. You can also enable or disable the Zerto Virtual Manager to send data to the SaaS platform for monitoring purposes.

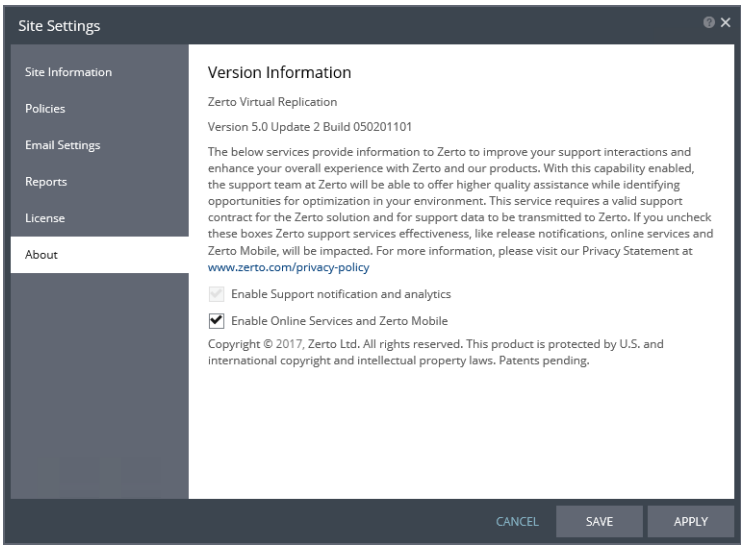

Enable Support notification and analytics – When selected, the CALLHOME feature is enabled, whereby analytics that are sent to Zerto that are used to improve Zerto Virtual Replication. The CALLHOME also enables to automatically update Zerto Virtual Replication when a new version of a hypervisor is released that is supported by Zerto.

Enable Online Services and Zerto Mobile – Allows licensed Zerto Virtual Manager users to enable or disable data being sent from the Zerto Virtual Manager to the SaaS platform thereby enabling site monitoring using the Zerto Mobile App.

To see version information, prepare to send analytics to Zerto, or send data to the Cloud:

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

   The *Site Settings* dialog is displayed.
2. Click *About*.

   The version and build of Zerto Virtual Replication that are installed in the site are displayed.
3. If you want to send analytics to Zerto automatically, check Enable Support notification and analytics. This initiates the CALLHOME feature. This information is used solely to improve Zerto Virtual Replication and to automatically update Zerto Virtual Replication when a new version of a hypervisor is released that is supported by Zerto.
4. If you want Zerto Virtual Replication to send information to our Online Services and Zerto Mobile App, check Enable Online Services and Zerto Mobile. This allows licensed Zerto Virtual Manager users to enable or disable data being sent from the Zerto Virtual Manager to the SaaS platform, thereby enabling site monitoring using the Zerto Mobile App.

   Note: The `Enable Online Services and Zerto Mobile` option is enabled by default.

   If the `Enable Online Services` and `Zerto Mobile` option is enabled and the user un-checks the checkbox, the user will receive a warning that unchecking `Enable Online Services and Zerto Mobile Application` will stop Zerto Virtual Replication from sending information to Online Services and to the Zerto Mobile Application, rendering these services inoperable for the entire installation.

CHAPTER 10: TESTING RECOVERY TO AWS

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. This chapter describes how to test VPG recovery.

The following topics are described in this chapter:

Recovering a protected virtual machine to AWS requires importing the machine and its associated volumes into EC2. Each machine and volume requires a separate import process. By default, Amazon limits accounts to a specific number of parallel import processes. If you have more machines and volumes than this limit, the process takes longer. The additional machines and volumes are queued and are imported only after an import process is available. If you intend to protect more machines and volumes than can be imported at one time by default, Zerto recommends that you contact AWS Support to increase the *ec2-import-instance/volume* limit.

Note: You cannot perform a failover test while a backup job is running.

The Test Failover Process

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 – instances – in a sandbox, using the test network specified in the VPG definition.

During the test, any changes to the protected virtual machines at the protected site are sent to the recovery site and new checkpoints continue to be generated, since replication of the protected machines continues throughout the test. You can also add your own checkpoints during the test period. You can initiate a failover during a test, as described in "Initiating a Failover During a Test", on page 102.

The Failover Test operation has the following basic steps:

- Starting the test.
  - The test virtual machine instances are created in AWS and configured to the checkpoint specified for the recovery.
  - The new instances 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 instances in a sandbox and powered on for testing.
- Stopping the test.
  - The instances in AWS are powered off and removed from the inventory.
  - The following tag is added to the checkpoint specified for the test: `Tested at startDateAndTimeOfTest` The updated 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.

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 and Stopping Failover Tests

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.

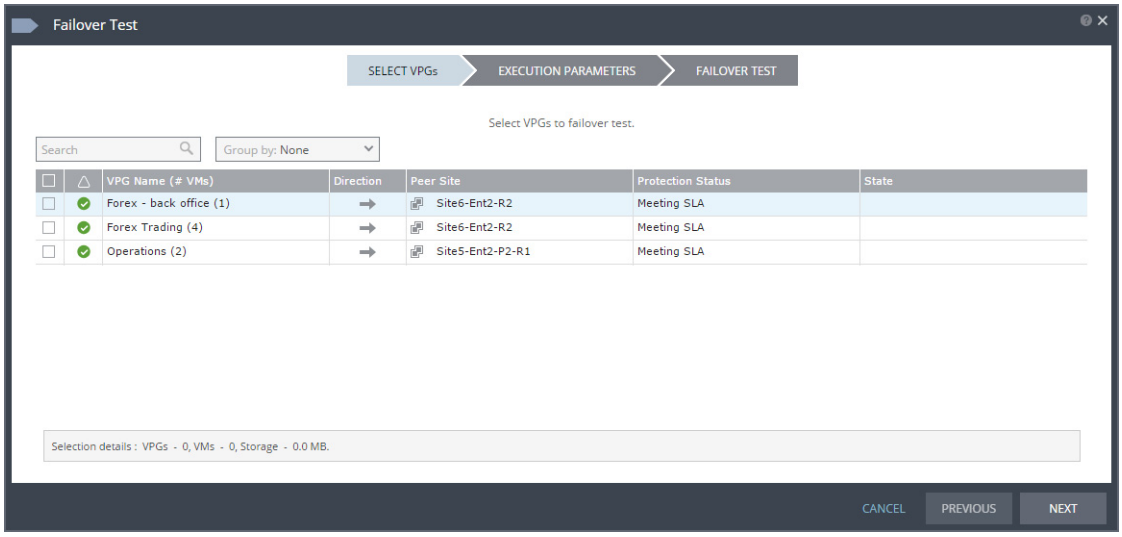

2. Select the VPGs to test. By default, all VPGs are listed.
   At the bottom, the selection details show the amount of data and the total number of virtual machines selected.
   The `Direction` arrow shows the direction of the process: from the protected site to the peer, recovery, site.
3. Click **NEXT**.
   The EXECUTION PARAMETERS step is displayed.

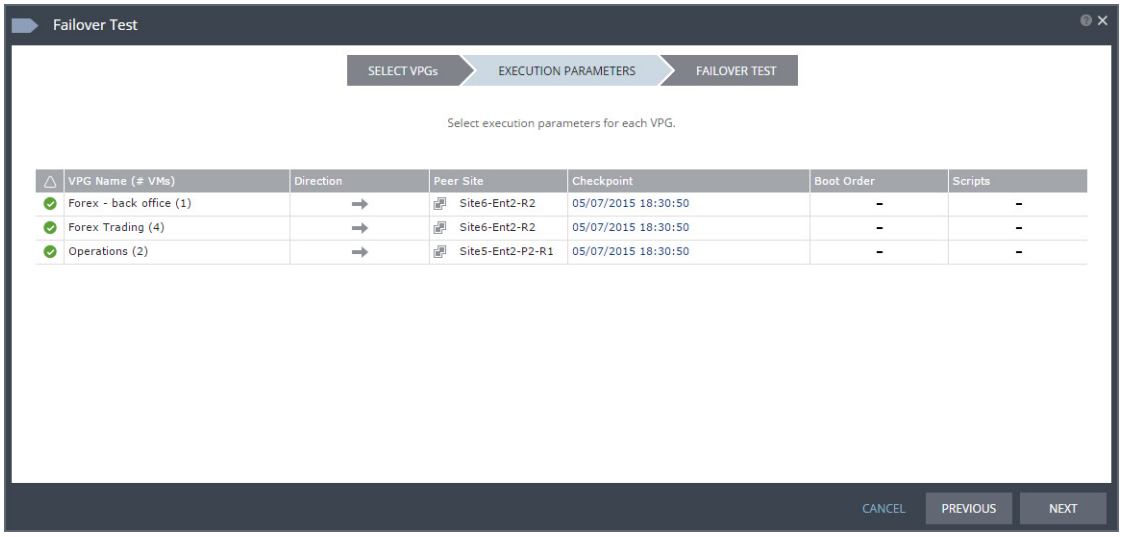

   You can select the checkpoint to use for the recovery and see if a boot order and scripts are defined for the VPG.
4. By default, the last checkpoint added to the journal is displayed in the Checkpoint column
   - To use this checkpoint, proceed to the next step.
   - To change the checkpoint, click the link that appears as the checkpoint.

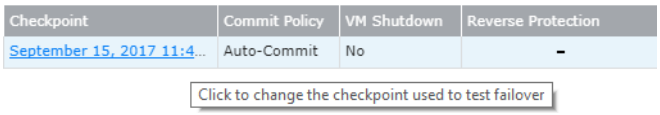

A window appears, displaying a list of the VPGs' checkpoints.

Latest: Recovery is to the **latest checkpoint**. This ensures that the data is crash-consistent for the recovery.

When selecting the latest checkpoint, the checkpoint used is the latest at this point.

If a **checkpoint is added** between this point and **starting the failover**, this **later** checkpoint is **not used**.

Latest Tagged Checkpoint: The recovery operation is to the latest checkpoint added in one of the following situations:

- By a user.
- When a failover test was previously performed on the VPG that includes the virtual machine.
- When the virtual machine was added to an existing VPG after the added virtual machine was synchronized.

5. **Latest VSS** – When VSS is used, the clone is to the latest VSS snapshot, ensuring that the data is both crash-consistent and application consistent to this point. The frequency of VSS snapshots determines how much data can be recovered. For details about VSS checkpoints, see "Ensuring Application Consistency – Checkpoints", on page 48.
6. To use a checkpoint which is **not** the latest checkpoint, or the latest tagged checkpoint, choose **Select from all available checkpoints**. By default, this option displays all checkpoints in the system. You can choose to display only automatic, or tagged checkpoints, or any combination of these types.
7. Click *OK*.
8. Click **NEXT**.

   The FAILOVER TEST step is displayed. The topology shows the number of VPGs and virtual machines being tested to failover to each recovery site. In the following example, 2 VPGs will be failed over to Site6-Ent2-R2, and they contain 5 virtual machines; and 1 VPG will be failed over to Site5-Ent2-P2-R2 and it contains 2 virtual machines.

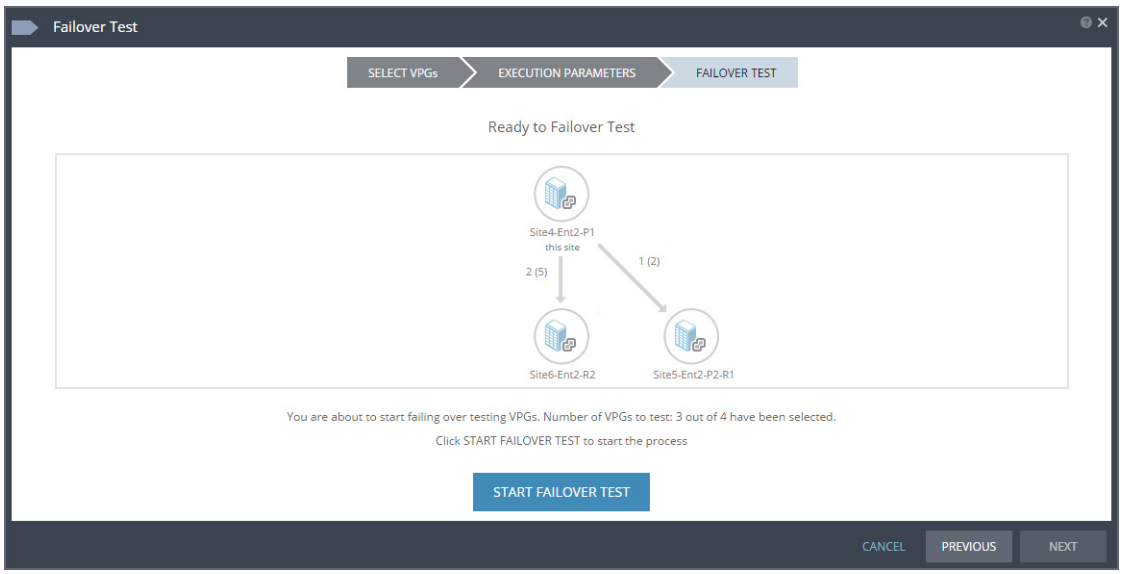

9. To start the test, click **START FAILOVER TEST**.

The test starts for the selected VPGs. The test begins with an initialization period during which the new instances are created in AWS. The protected virtual machines are created as new instances in EC2.

The default value for new instances in Zerto Virtual Replication is m3.xlarge except in the Asia Pacific (Seoul) region where they are defined as m4.xlarge instances. If these instances do not meet your needs, you can change this value in the Policies tab of the Site Settings dialog, see "Configuring Disaster Recovery Policies", on page 78. You can also change the instance type of new instances when you create or edit a VPG.

If you did not define a private IP for a virtual machine in the VPG definition, during recovery AWS sets the private IP from the defined subnet range.

After Starting a Test, What Happens?

The virtual machines in the virtual protection group are created in AWS. In the AWS console, the new virtual machines appear with their original names and the suffix *testing recovery*.

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, as described in "Initiating a Failover During a Test", on page 102.
- 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 `Testing Failover` when a failover test is being performed.

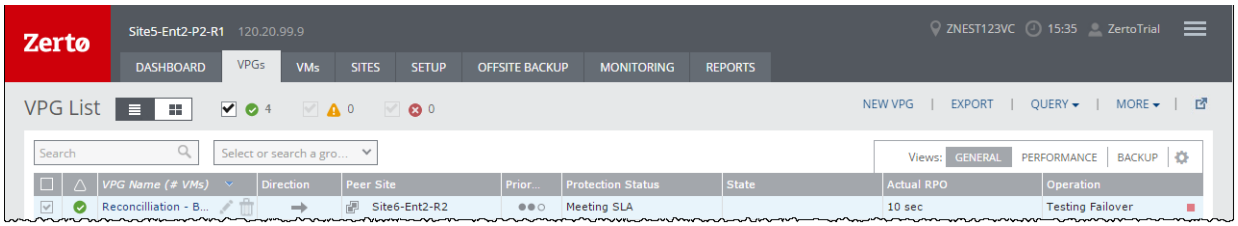

- In the Zerto User Interface, click the *VPGs* tab, and then click the name of a VPG you are testing. A dynamic tab is created displaying the specific VPG details including the status of the failover test.

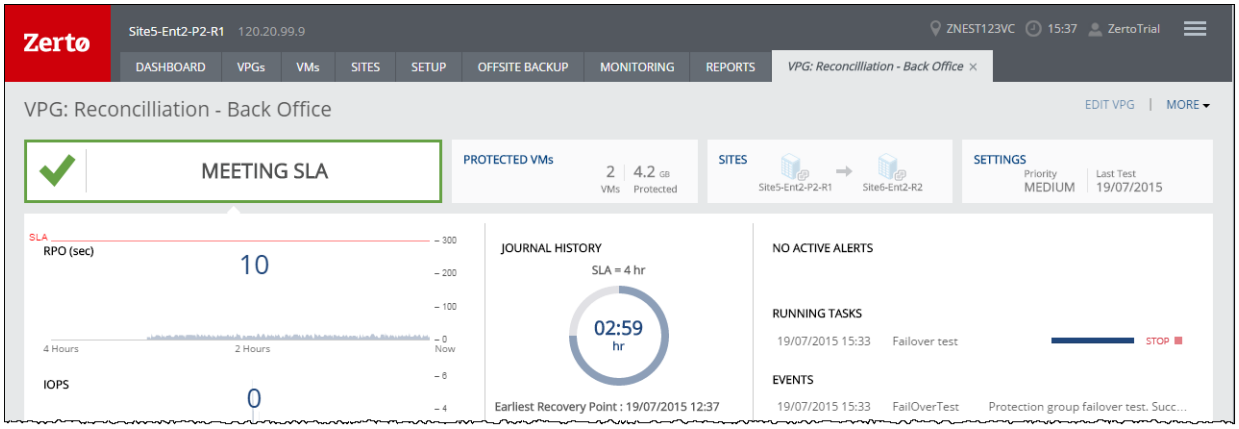

To stop a failover test:

1. Click the *Stop* icon, in either the Dashboard or the dynamic tab, to stop the test in the specific VPG tab.

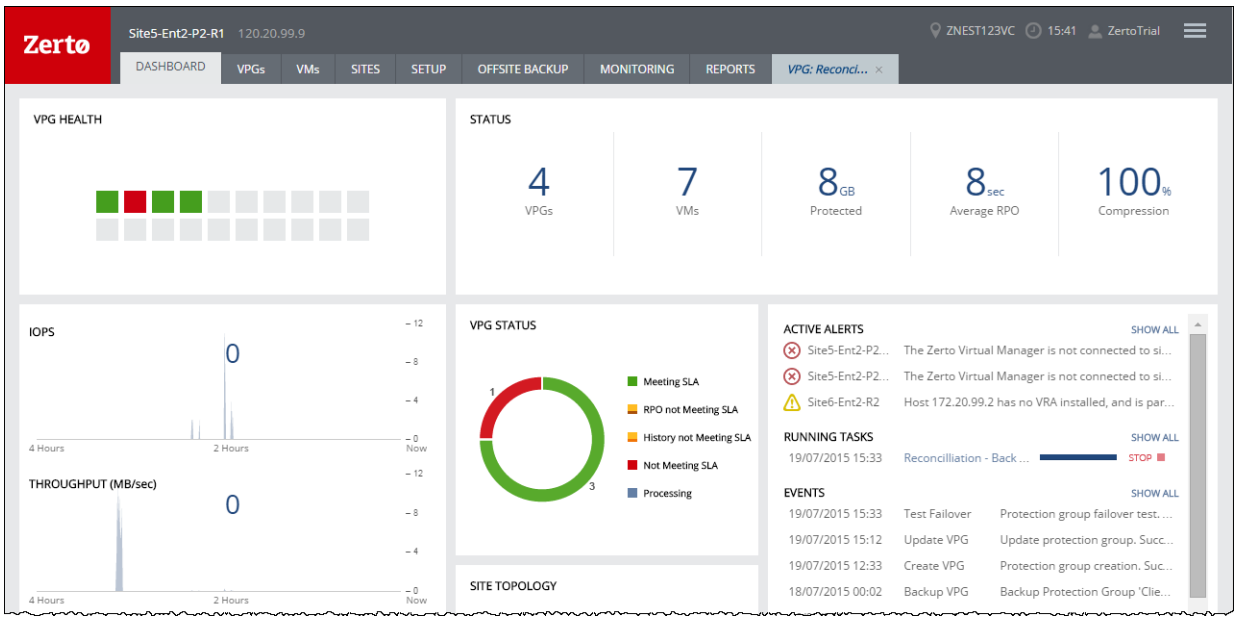

You can also stop the test via the *TASKS* popup dialog in the status bar, or by selecting *MONITORING* > *TASKS*.

The *Stop Test* dialog is displayed.

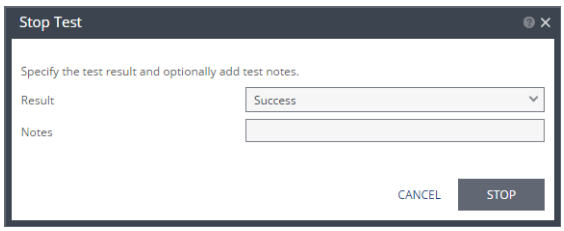

2. In the `Result` field specify whether the test succeeded or failed.
3. Optionally, in the `Notes` field, add a description of the test. For example, specify where external files that describe the tests performed are saved. Notes are limited to 255 characters.
4. Click *STOP*.

After stopping a test, the following occurs:

- Virtual machines in the recovery site are powered off and removed.
- The resource group created for the operation is deleted.
- 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.

Viewing Test Results

After stopping a test, you can see the test results in Zerto Virtual Replication reports. For more information, see "Zerto Virtual Replication Reports", on page 127.

Live Disaster Recovery Testing

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

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

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

The following flowchart shows the testing decision flow:

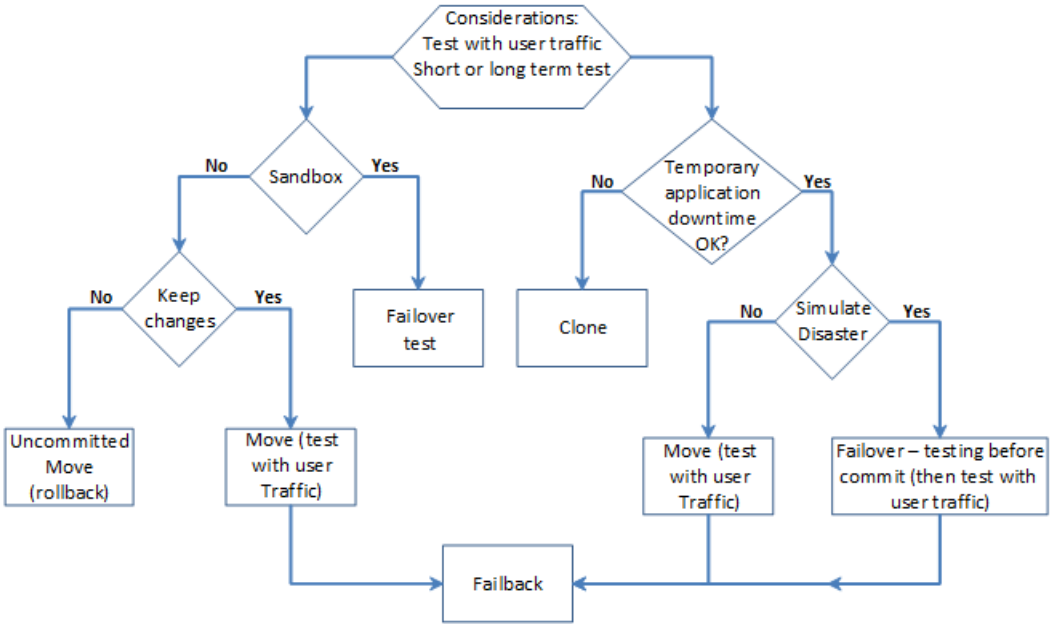

During any live test, Zerto recommends that you only maintain one working version of the same virtual machine. Thus, the first step in any test, except for a Failover Test, is to make sure that the protected virtual machines are shut down before starting to test recovered machines. During a Zerto Virtual Replication Move operation the first step Zerto Virtual Replication performs is to shut down the protected machines, to ensure data integrity. However, a Zerto Virtual Replication Failover operation assumes that the protected virtual machines are no longer accessible (the total site disaster scenario) and does not attempt to shut them down at the beginning of the operation. In a live test using a Failover operation you have to manually shut down the virtual machines to be tested at the beginning of the test in order to prevent potential *split-brain* situations where two instances of the same applications are live at the same time.

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

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

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

Using a Failover Test Operation

Use a Failover Test operation if recovering the virtual machines in a sandbox. Using the test network specified in the VPG definition for network isolation, is sufficient for a test.

Procedure

The Failover Test operation is described in detail in "Starting and Stopping Failover Tests" on page 84. The following highlights specific steps to enable using the Failover Test when recovering the virtual machines in a sandbox.

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

Failover Test Considerations

- You do not have to shut down the protected virtual machines, and changes from the test phase are not kept or applied to the protected applications.
- You can recover to a specific point-in-time.
- You can use an isolated network to enable testing in a sandbox environment and not a live DR environment. This is the recommended practice.
- At the end of the test, you can power on the virtual machines in the protected site and continue to work without the need to save or replicate back any data changed during the test.

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

Using an Uncommitted Move Operation

Use a Move operation with the commit/rollback policy set to rollback after the test period, if you need to test the recovery of virtual machines in the recovery site production environment.

Procedure

The Move operation is described in detail in "Moving Protected Virtual Machines to a Remote Site", on page 93. The following procedure highlights specific steps to enable using the Move functionality for a DR test.

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

Move Considerations

- Changes from the pre-commit phase are not kept or applied to the protected applications.
- The new instances are connected to the network for a full test of the environment.
- The protected machines are turned off until the end of the test, ensuring that there are no conflicts between the protected site and recovery site.
- You can only recover to the last checkpoint written to the journal, at the start of the Move operation.

CHAPTER 11: MIGRATING A VPG TO AWS

This chapter describes a planned migration of a virtual protection group - VPG - to a remote site. The following topics are described in this chapter:

Recovering a protected virtual machine to AWS requires importing the machine and its associated volumes into EC2. Each machine and volume requires a separate import process. By default, Amazon limits accounts to a specific number of parallel import processes. If you have more machines and volumes than this limit, the process takes longer. The additional machines and volumes are queued and are imported only after an import process is available. If you intend to protect more machines and volumes than can be imported at one time by default, Zerto recommends that you contact AWS Support to increase the *ec2-import-instance/volume* limit.

Note: You cannot perform a move while a backup job is running.

The Move Process

Use the *Move* operation to move groups of protected virtual machines from a protected site to a recovery site in a planned migration.

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

Note: To recover virtual machines on the recovery site during disaster recovery, see "Managing Failover to AWS", on page 97.

The MOVE operation has the following basic steps:

- Shutting down the protected virtual machines gracefully. This ensures data integrity.

  If the machines cannot be gracefully shut down, for example, when VMware Tools or Microsoft Integration Services is not available, you must manually shut down the machines before starting the Move operation or forcibly power off the virtual machines as part of the Move operation. If the machines cannot be gracefully shut down automatically and are not shut down manually and the Move operation does not forcibly power them off, the Move operation stops and Zerto Virtual Replication rolls back the virtual machines to their original status.
- Inserting a clean checkpoint. This avoids potential data loss since the virtual machines are not on and the new checkpoint is created after all I/Os have been written to disk.
- Transferring all the latest changes that are still in the queue to the recovery site, including the new checkpoint.
- Creating the virtual machines in the recovery site and attaching each virtual machine to its relevant virtual disks, based on the last checkpoint.

  Note: The new instances are created without CD-ROM or DVD drives, even if the protected virtual machines had CD-ROM or DVD drives. Also, as long as the virtual machines are created, the operation is considered successful, even if the virtual machines are not created with their complete definition, for example setting a private IP cannot be performed.
- Powering on the new instances, making them available to the user. If applicable, the boot order defined in the VPG settings is used to power on the machines.

  If the new instances do not power on, the process continues and the new instances must be powered on manually.
- By default, automatically committing the Move operation without testing. However, you can also run basic tests on the new instances to ensure their validity. Depending of the commit/rollback policy that you specified for the operation, the operation is committed, finalizing the move, or rolled back, aborting the operation.
- If `Keep Source VMs` is not selected, the protected virtual machines are removed from the inventory.

  Note: If *Keep Source VMs* is not selected, and **the virtual machines or vCD vApp** are already protected in other VPGs, continuing with the operation will cause the **virtual machines or vCD vApp to be deleted** from **other VPGs** that are

protecting them and to the journals of these VPGs to be reset. In the event of vCD vApp or if no other virtual machines are left to protect, the **entire VPG will be removed**.

- Copying data from the S3 buckets to the new EBS disks of the new instances. During this process, the new instances cannot be used. The protected virtual machines are created as new instances in EC2. The default value for new instances in Zerto Virtual Replication is m3.xlarge except in the Asia Pacific (Seoul) region where they are defined as m4.xlarge instances. If these instances do not meet your needs, you can change this value in the Policies tab of the Site Settings dialog, see "Configuring Disaster Recovery Policies", on page 78. You can also change the instance type of new instances when you create or edit a VPG.

A move differs from a failover in that with a move you cannot select a checkpoint to restore the virtual machine to. Also, to ensure data integrity, the protected virtual machines are powered off completely and a final checkpoint created so that there is no data loss before the move is implemented.

You can initiate the Move operation from either the protected site or recovery site.

Moving Protected Virtual Machines to a Remote Site

You can move the virtual machines in a virtual protection group to AWS, where the virtual machines are replicated.

To initiate a move:

1. In the Zerto User Interface select ACTIONS > MOVE VPG.
   The Move wizard is displayed.

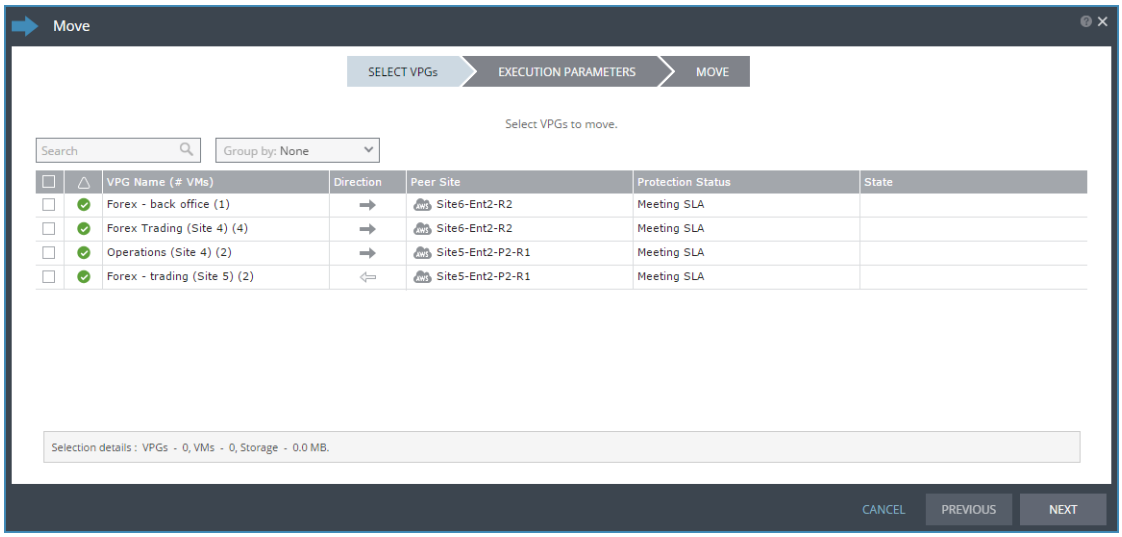

2. Select the VPGs to move.
   At the bottom, the selection details show the amount of data and the total number of virtual machines selected.
   The Direction arrow shows the direction of the process: from the protected site to the peer, recovery, site.
3. Click **NEXT**. The EXECUTION PARAMETERS step is displayed.

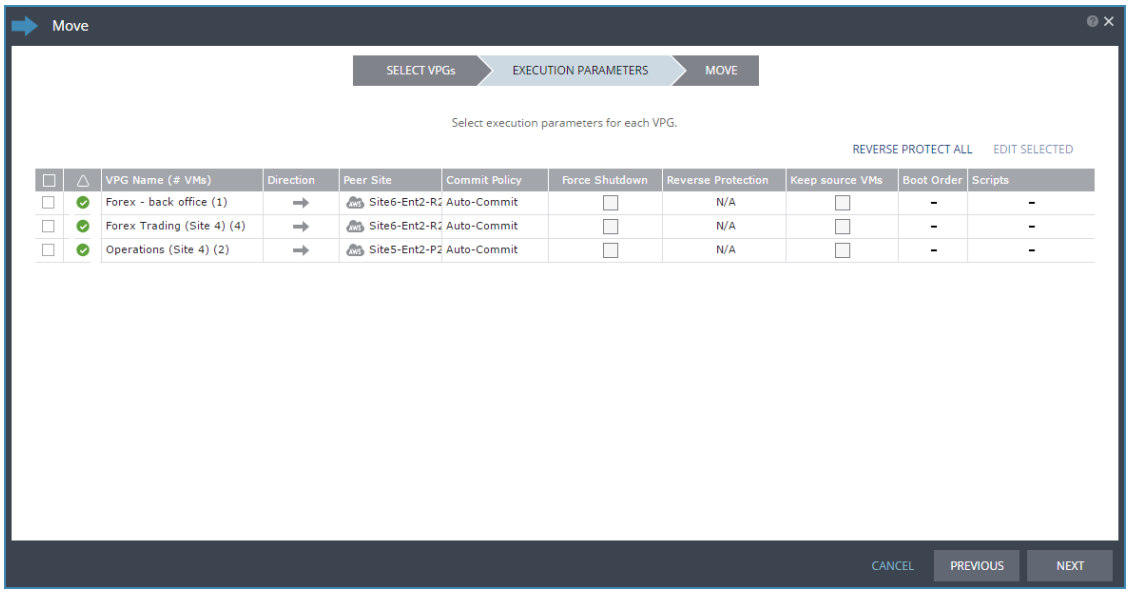

   You can change the following values to use for the recovery:
   - **Commit Policy**
   - **Force Shutdown** policy
   - **Keep Source VMs** settings

   You can also see if a boot order and scripts are defined for the VPG.

4. To change the commit policy, click on the field or select the VPG and click *EDIT SELECTED*.
   a) To commit the recovery operation automatically, without any checking, select `Auto-Commit` and 0 minutes.
   b) If you do not want an automatic commit or rollback, select `None`. You must manually commit or roll back.

   To allow checking before committing or rolling back, specify an amount of time to check the recovered machines, in minutes, before the automatic commit or rollback action is performed. During this time period, check that the new virtual machines are OK and then commit the operation or roll it back. The maximum amount of time you can delay the commit or rollback operation is 1440 minutes, which is 24 hours.
5. To specify the shutdown policy, double-click the `VM Shutdown` field. If the virtual machines cannot be gracefully shut down, for example if a utility such as VMware Tools or Microsoft Integration Services is not installed on one of the virtual machines in the VPG, the Move operation fails unless you specify that you want to force the shutdown. If a utility is installed on the protected virtual machines, the procedure waits five minutes for the virtual machines to be gracefully shut down before forcibly powering them off.
6. To **prevent** the protected virtual machines or vCD vApp from being **deleted** in the protected site, click the *Keep source VMs* checkbox.

   IMPORTANT:
   - The **virtual machines or vCD vApp** will be **removed** from the **other VPGs** that are protecting them if the following conditions apply:
     - The virtual machines or vCD vApp are already protected in other VPGs
     - Reverse protection is specified
     - *Keep Source VMs* is not checked
   - If your VPG has a vCD vApp, or if there are no other virtual machines left to protect, the **entire VPG will be removed.**

- Protecting virtual machines or vCD vApps in several VPGs is enabled only if both the protected site and the recovery site, as well as the VRAs installed on these sites, are of version 5.0 and higher.

7. Click **NEXT**.
   - When reverse protection is specified for a VPG residing on a vCD site that is replicating to either a vSphere or Hyper-V site, the boot order settings will not reserve the start delay vCD vApp settings for virtual machines with the same order number.

   The MOVE step is displayed. The topology shows the number of VPGs and virtual machines being moved to each peer site. In the following example, 2 VPGs will be moved to Site6-Ent2-R2, and they contain 5 virtual machines; and 1 VPG will be moved to Site5-Ent2-P2-R2 and it contains 2 virtual machines.

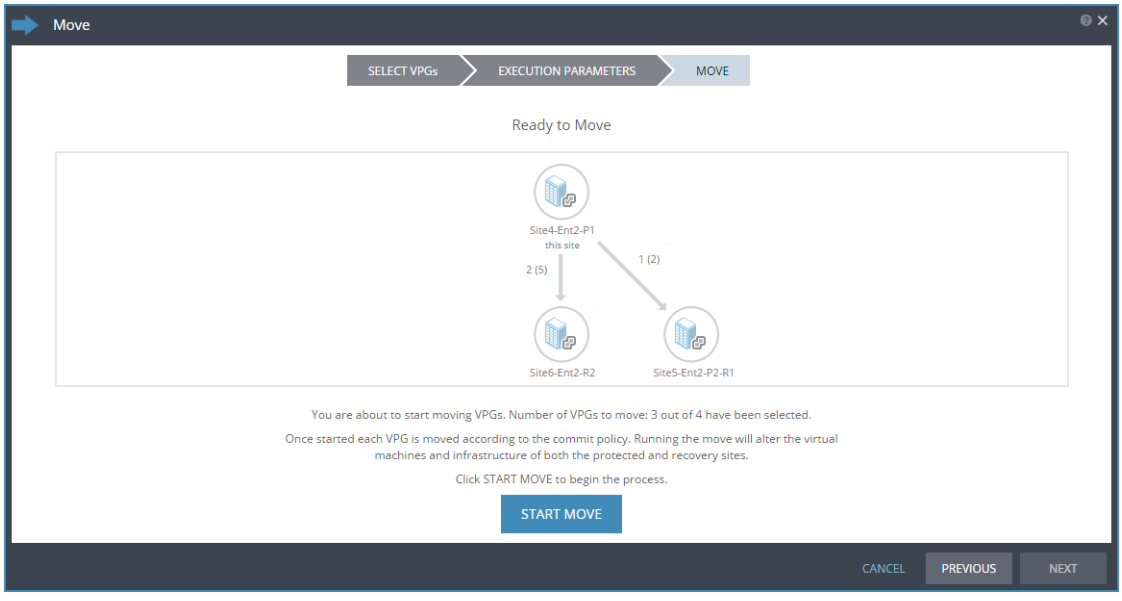

8. Click **START MOVE** to start the migration.
9. If a commit policy was set with a timeout greater than zero, as described in step 4, you can check the new instances on AWS before they are removed from the protected site.

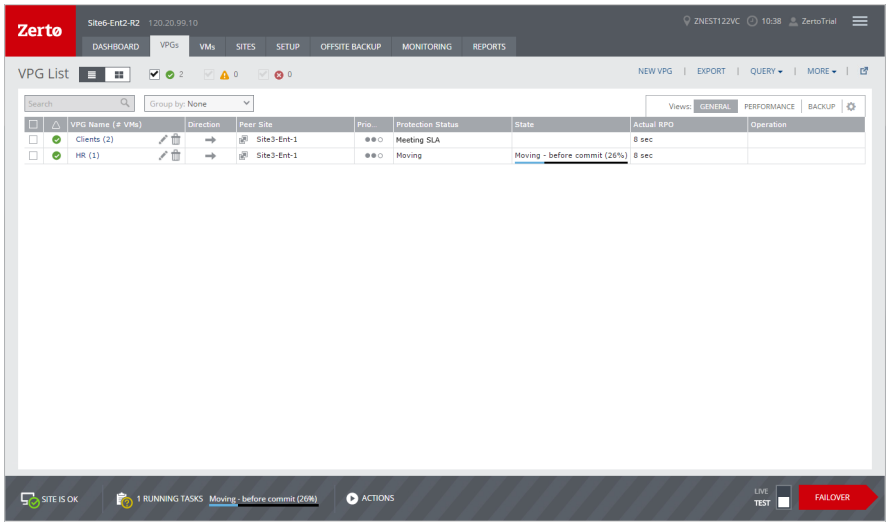

Note: If an instance exists on the recovery site with the same name as a virtual machine being migrated, the machine is moved and named in the recovery site with a number added as a suffix to the name, starting with the number 1.

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

All testing done during this period, before committing or rolling back the Move operation, is written to EBS disks. These virtual disks are automatically defined by AWS when the new instances are created.

Note: You cannot take a snapshot of a virtual machine before the Move operation is committed and the data from the journal promoted to EBS disks, since the virtual machine volumes are still managed by the VRA and not directly by the virtual machine. Taking a snapshot of a machine that is in the process of being moved will corrupt that machine.

10. Check the virtual machines on the recovery site, then either:
    - Wait for the specified **Commit Policy** time to elapse, and the specified operation, either **Commit** or **Rollback**, is performed automatically.
    - Or, in the specific VPG tab, click the Commit or Rollback icon ( ).
      - Click *Commit* to confirm the commit.
      - Click *Rollback* to roll back the operation, removing the virtual machines that were created on the recovery site and rebooting the machines on the protected site. The *Rollback* dialog is displayed to confirm the rollback.

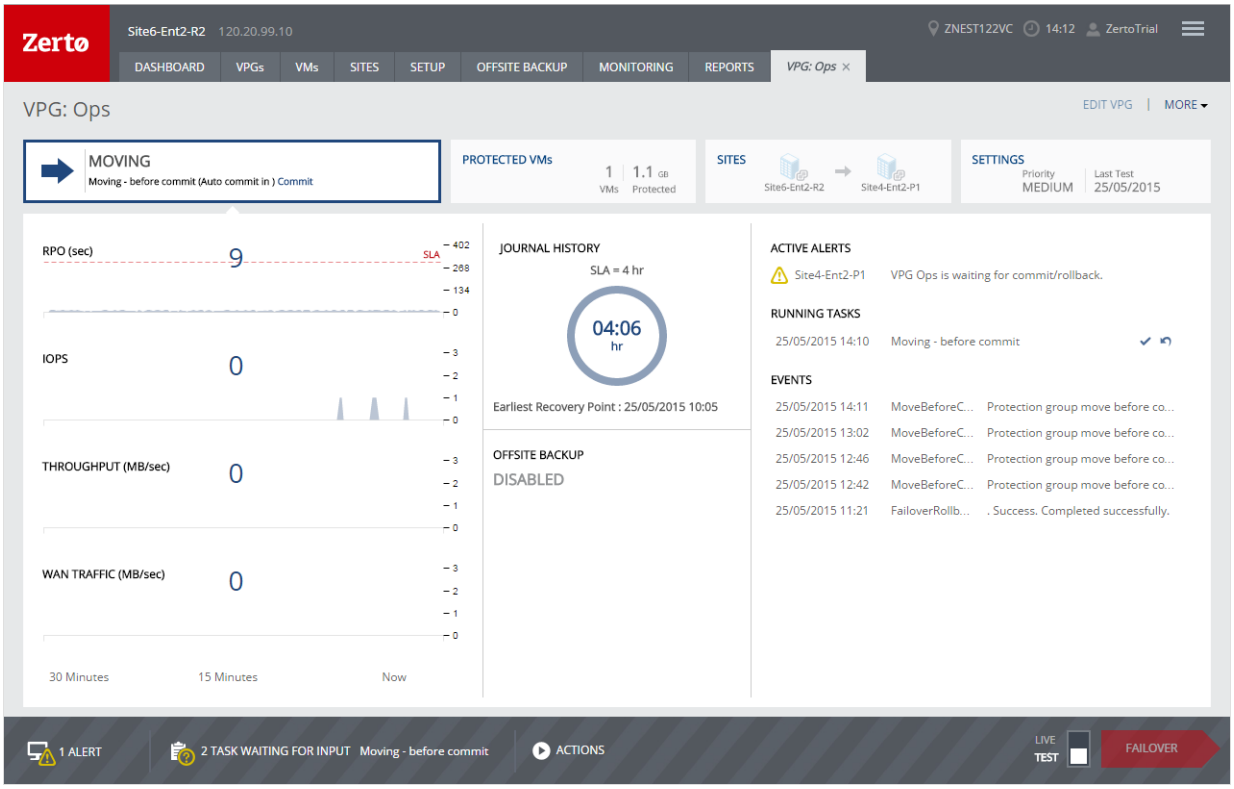

You can also commit or roll back the operation via the *TASKS* popup dialog in the status bar, or by selecting *MONITORING* > *TASKS*.

- After the new instances are up and running and committed in the recovery site, the powered off virtual machines in the protected site are removed from the protected site.
- Finally, data is copied from the S3 buckets to the EBS disks attached to the new instances in AWS.
- While data is being copied, the new instances are not available.

Notes:

- If virtual machines or vCD vApp are already protected **in several VPGs**, and reverse protection is configured, the **virtual machines or vCD vApp are deleted** from the protected site. This will result in the **removal of these virtual machines from other VPGs** that are protecting them and to the journals of these VPGs to be reset. In the event of vCD vApp or if no other virtual machines are left to protect, the **entire VPG will be removed**.

  Protecting virtual machines in several VPGs is enabled only if both the protected site and the recovery site, as well as the VRAs installed on these sites, are of version 5.0 and higher.
- If `Keep Source VMs` is selected, the protected virtual machines are not removed from the protected site.

Protecting virtual machines in several VPGs is enabled only if both the protected site and the recovery site, as well as the VRAs installed on these sites, are of version 5.0 and higher.

The protected virtual machines are created as new instances in EC2. The default value for new instances in Zerto Virtual Replication is m3.xlarge except in the Asia Pacific (Seoul) region where they are defined as m4.xlarge instances. If these instances do not meet your needs, you can change this value in the Policies tab of the Site Settings dialog, see "Configuring Disaster Recovery Policies", on page 78. You can also change the instance type of new instances when you create or edit a VPG.

If you did not define a private IP for a virtual machine in the VPG definition, during recovery AWS sets the private IP from the defined subnet range.

Note: If the new instances do not power on, the process continues and the new instances must be manually powered on.

CHAPTER 12: MANAGING FAILOVER TO AWS

This chapter describes how to perform a failover to AWS after an unforeseen disaster. The following topics are described in this chapter:

Recovering a protected virtual machine to AWS requires importing the machine and its associated volumes into EC2. Each machine and volume requires a separate import process. By default, Amazon limits accounts to a specific number of parallel import processes. If you have more machines and volumes than this limit, the process takes longer. The additional machines and volumes are queued and are imported only after an import process is available. If you intend to protect more machines and volumes than can be imported at one time by default, Zerto recommends that you contact AWS Support to increase the *ec2-import-instance/volume* limit.

Note: If you need to perform a failover while a backup job is running, the backup job is aborted to enable the failover to run.

The Failover Process

Use the *Failover* operation following a disaster to recover protected virtual machines to AWS.

Note: You can also move virtual machines from the protected site to AWS in a planned migration. For details, see "Migrating a VPG to AWS", on page 91.

When you set up a failover you always specify a checkpoint to which you want to recover the virtual machines. When you select a checkpoint – either the last auto-generated checkpoint, an earlier checkpoint, or a tagged checkpoint – Zerto Virtual Replication makes sure that the new instances on AWS are recovered to this specified point-in-time. By setting a commit policy that enables checking the recovered machines before committing the failover, you can check the integrity of the recovered machines. If the machines are OK, you can commit the failover. Otherwise, you can roll back the operation and then repeat the procedure using a different checkpoint.

The Failover operation has the following basic steps:

- If the protected site or Zerto Virtual Manager is down, the process continues with the next step.

  If the protected site or Zerto Virtual Manager is still running, the failover requirements are determined:
  - If the default is requested, doing nothing to the protected virtual machines, the Failover operation continues with the next step.
  - If shutting down the protected virtual machines is requested and the protected virtual machines do not have a utility such as VMware Tools or Microsoft Integration Services available, the Failover operation fails.
  - If forcibly shutting down the protected virtual machines is requested, the protected virtual machines are shut down and the Failover operation continues with the next step.
- Creating the new instances on AWS in the production network and attaching each instance to its relevant EBS disks, configured to the checkpoint specified for the recovery. The new instances are created without CD-ROM or DVD drives, even if the protected virtual machines had CD-ROM or DVD drives. Also, as long as the virtual machines are created, the operation is considered successful, even if the virtual machines are not created with their complete definition, for example setting a private IP cannot be performed. The protected virtual machines are created as new instances in EC2. The default value for new instances in Zerto Virtual Replication is m3.xlarge except in the Asia Pacific (Seoul) region where they are defined as m4.xlarge instances. If these instances do not meet your needs, you can change this value in the Policies tab of the Site Settings dialog, see "Configuring Disaster Recovery Policies", on page 78. You can also change the instance type of new instances when you create or edit a VPG.

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

- Powering on the new instances, making them available to the user. If applicable, the boot order defined in the VPG settings is used to power on the machines.

  If the instances do not power on, the process continues and the instances must be manually powered on.
- The default is to automatically commit the Failover operation without testing. However, you can also run basic tests on the machines to ensure their validity to the specified checkpoint. Depending on the commit/rollback policy that you specified for the operation after testing either the operation is committed, finalizing the failover, or rolled back, aborting the operation.
- If the protected site is still available, for example, after a partial disaster, the original protected site virtual machines are not powered off and are not removed.

Initiating a Failover

You can initiate a failover, whereby the virtual machines in the virtual protection group are replicated to a set checkpoint in AWS.

You can initiate a failover to the last checkpoint recorded in the journal, even if the protected site is no longer up. You can initiate a failover during a test, as described in "Initiating a Failover During a Test", on page 102.

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

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

To initiate a failover:

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

   The *Failover* wizard is displayed.

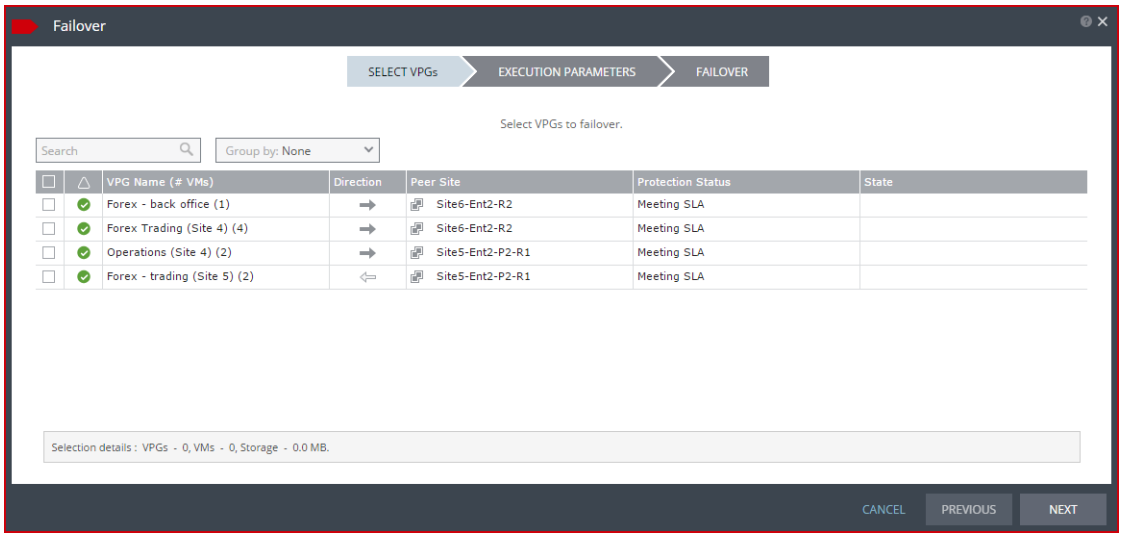

2. Select the VPGs to failover. By default, all VPGs are listed.

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

   The `Direction` arrow shows the direction of the process: from the protected site to the peer, recovery, site.
3. Click *NEXT*.

The *EXECUTION PARAMETERS* step is displayed.

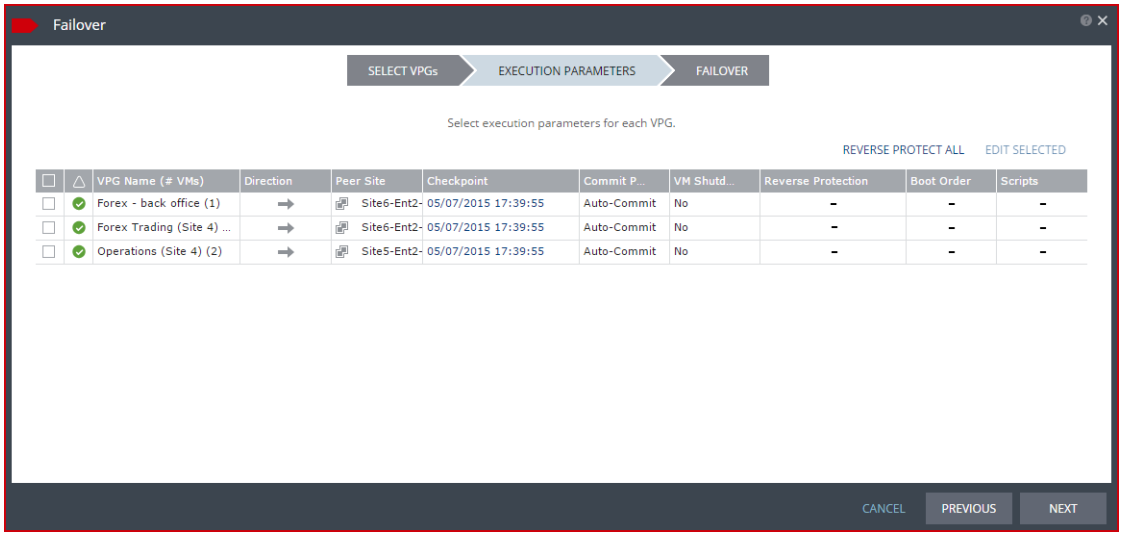

You can change the following values to use for the recovery:

- The commit policy
- The checkpoint to use
- The shutdown policy

You can also see if a boot order and scripts are defined for the VPG.

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

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

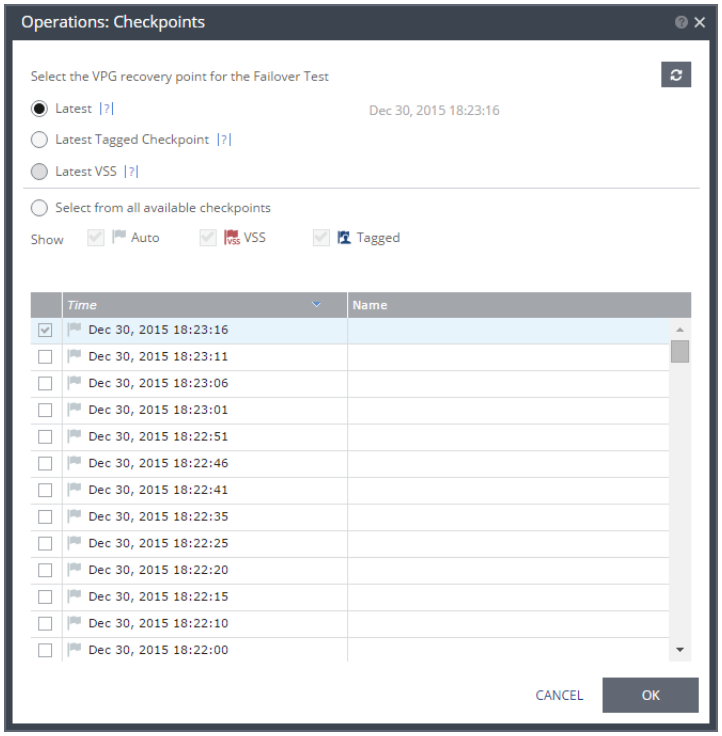

5. Select the checkpoint to use. Click the refresh button to refresh the list. You can choose from one of the following checkpoints:

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

Latest Tagged Checkpoint – The recovery operation is to the latest checkpoint added in one of the following situations:

- By a user.
- When a failover test was previously performed on the VPG that includes the virtual machine.
- When the virtual machine was added to an existing VPG after the added virtual machine was synchronized.

Latest VSS – When VSS is used, the clone is to the latest VSS snapshot, ensuring that the data is both crash-consistent and application consistent to this point. The frequency of VSS snapshots determines how much data can be recovered. For details about VSS checkpoints, see "Ensuring Application Consistency – Checkpoints", on page 48.

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

6. Click *OK*.
7. To change the commit policy, click on the field or select the VPG and click *EDIT SELECTED*.
   a) To commit the recovery operation automatically, without any checking, select `Auto-Commit` and 0 minutes.
   b) If you do not want an automatic commit or rollback, select `None`. You must manually commit or roll back.

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

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

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

   Force Shutdown – The protected virtual machines are forcibly shut down before starting the failover. This is similar to performing a Move operation to a specified checkpoint. If the protected virtual machines have Microsoft Integration Services available, the procedure waits five minutes for the virtual machines to be gracefully shut down before forcibly powering them off.
9. Click *NEXT*.
10. Click **OK**. If a virtual machine is deleted from other VPGs, the journals of these VPGs are reset.

The *FAILOVER* step is displayed. The topology shows the number of VPGs and virtual machines being failed over to each recovery site. In the following example, 2 VPGs will be failed over to Site6-Ent2-R2, and they contain 5 virtual machines; and 1 VPG will be failed over to Site5-Ent2-P2-R2 and it contains 2 virtual machines.

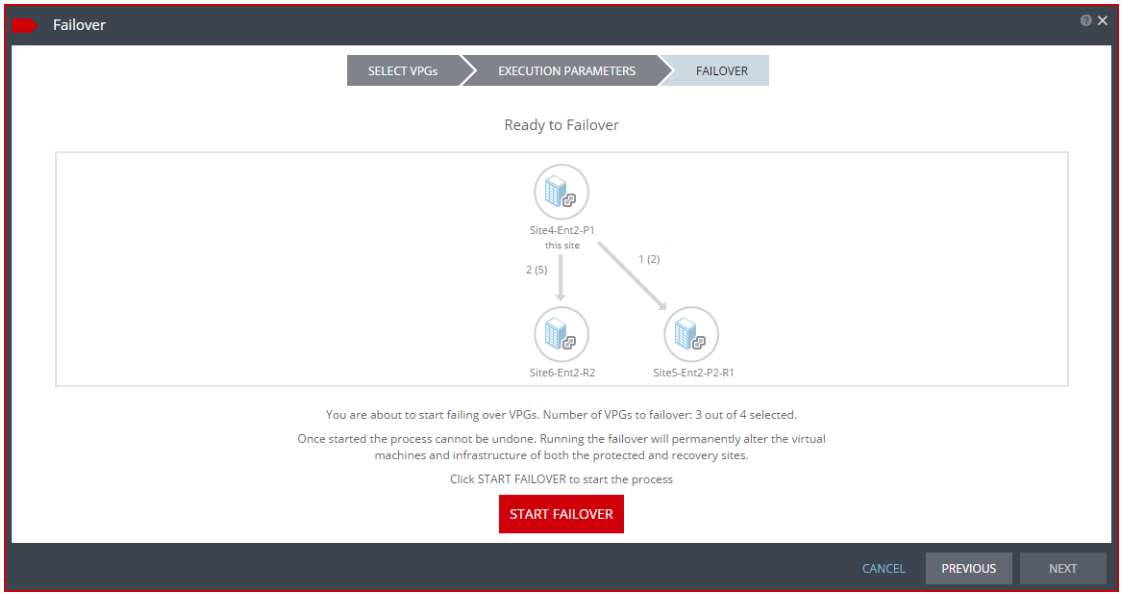

11. Click *START FAILOVER* to start the failover.

    If a commit policy was set with a timeout greater than zero, you can check the new instances on AWS before committing the failover operation.

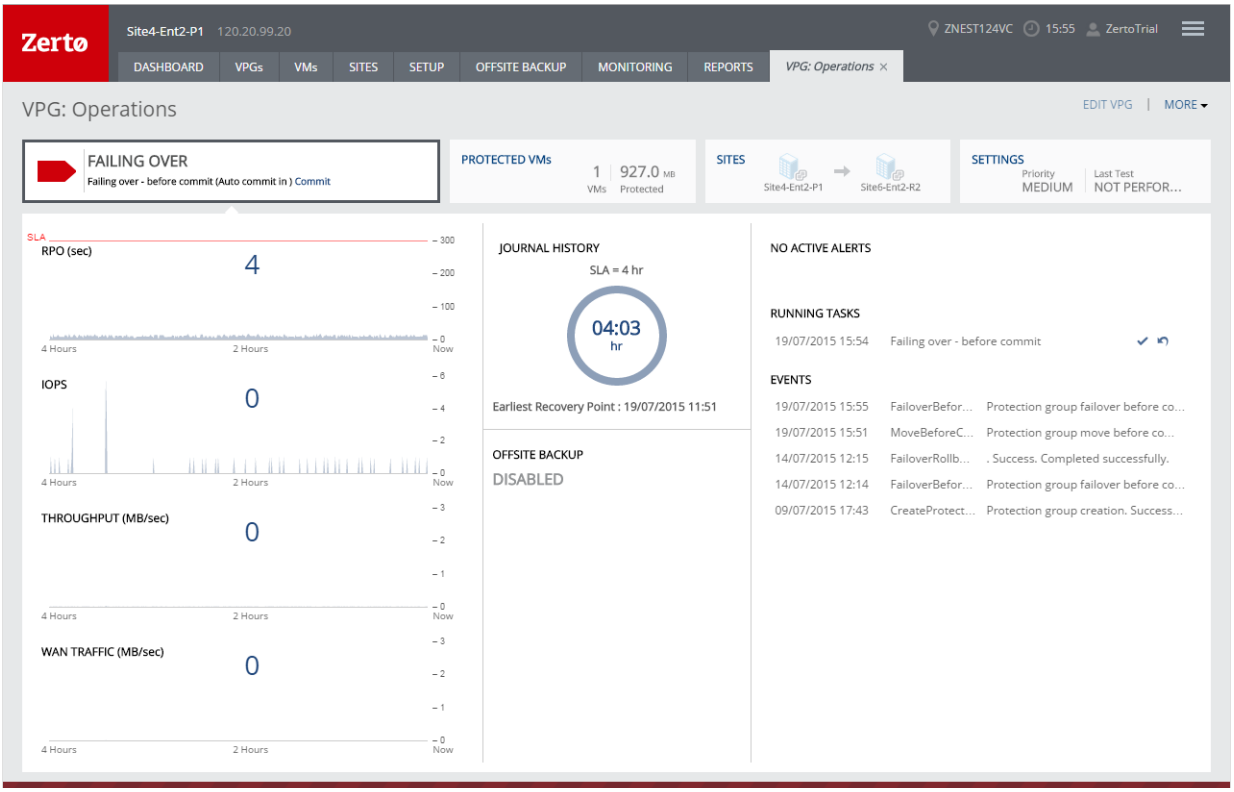

The failover starts by creating the new instances on AWS to the point-in-time specified: either the last data transferred from the protected site or to one of the checkpoints written in the journal.

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

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

All testing done during this period, before committing or rolling back the failover operation, is written to EBS virtual disks. These virtual disks are automatically defined when the instances are created on AWS for testing.

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

12. After checking the virtual machine instances in AWS, choose one of the following:
    - Wait for the specified `Commit Policy` time to elapse, and the specified operation, either `Commit` or `Rollback`, is performed automatically.
    - Click the *Commit* or *Rollback* icon ( ) in the specific VPG tab.
      Click *Commit*. The *Commit* dialog is displayed to confirm the commit.
      Click *Rollback* to roll back the operation, removing the virtual machines that were created on the recovery site and rebooting the machines on the protected site. The *Rollback* dialog is displayed to confirm the rollback.

    You can also commit or roll back the operation via the *TASKS* popup dialog in the status bar, or by selecting *MONITORING* > *TASKS*.

The protected virtual machines are created as new instances in EC2. The default value for new instances in Zerto Virtual Replication is m3.xlarge except in the Asia Pacific (Seoul) region where they are defined as m4.xlarge instances. If these instances do not meet your needs, you can change this value in the Policies tab of the Site Settings dialog, see "Configuring Disaster Recovery Policies", on page 78. You can also change the instance type of new instances when you create or edit a VPG.

If you did not define a private IP for a virtual machine in the VPG definition, during recovery AWS sets the private IP from the defined subnet range.

Note: If the new instances do not power on, the process continues and the instances must be manually powered on.

What Happens When the Protected Site is Down

If the protected site is down, you can initiate the failover from AWS, as described above in "To initiate a failover:", on page 98.

The tab for a specific VPG tab for a VPG shows that recovery is possible.

If the Zerto Virtual Manager service is down the actual machines that are being protected can still be up, but they are only recoverable to the last checkpoint written before the Zerto Virtual Manager service went down. If the hypervisor management tool, such as vCenter Server or Microsoft SCVMM, is down, some of the protected virtual machines might not be protected.

When there is no connection with the protected site, the status for recovered VPGs is red with an `Error` status and green while recovery is being performed. If the protected site restarts, the status changes to orange.

Initiating a Failover During a Test

Replication continues during a test. If you need to initiate a failover during a test, you initiate the failover. The test stops to enable the failover and then a normal failover is performed, as described in "Initiating a Failover", on page 98. Any changes made to test the failover are not replicated, as only changes to the protected machines in the VPG are replicated.

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

CHAPTER 13: CLONING A VPG TO AWS

You can create a clone of each virtual machine in a VPG. The clone is a copy of the protected virtual machine, located on the recovery site, while the virtual machine on the protected site remains protected and live.

The following topics are described in this chapter:

Note: You cannot clone virtual machines in a VPG test while a backup job is running.

The Clone Process

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

The Clone operation has the following basic steps:

- Creating the cloned disks at the recovery site with the data from the journal to the specified checkpoint.
- Creating the virtual machines at the recovery site in the move/failover network and attach each virtual machine to its relevant cloned disks, configured to the checkpoint specified for the clone.

  Note: The virtual machines are created without CD-ROM or DVD drives, even if the protected virtual machines have CD-ROM or DVD drives.

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

Cloning Protected Virtual Machines to the Remote Site

You might want to create a clone if you need to have a copy of the virtual machines saved to a specific point-in-time, for example, when a VPG enters a *Replication Paused* state, or when testing a VPG in a live DR test.

To clone a VPG:

1. In the Zerto User Interface, in the *VPGs* tab click the name of the VPG to be cloned.

   A new tab is added to the Zerto User Interface, with the name of the VPG that you clicked. The tab displays data about the VPG.

   Note: If the VPG was previously viewed, and the tab for this VPG is still displayed, you can access the details by selecting the tab.
2. Select the new tab and click *MORE* > *Offsite Clone*.

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

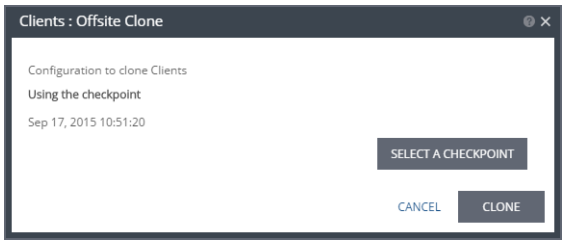

3. If you intend to use the last checkpoint, which is displayed in the dialog, go to step 6.

To select the checkpoint, click *SELECT A CHECKPOINT*.

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

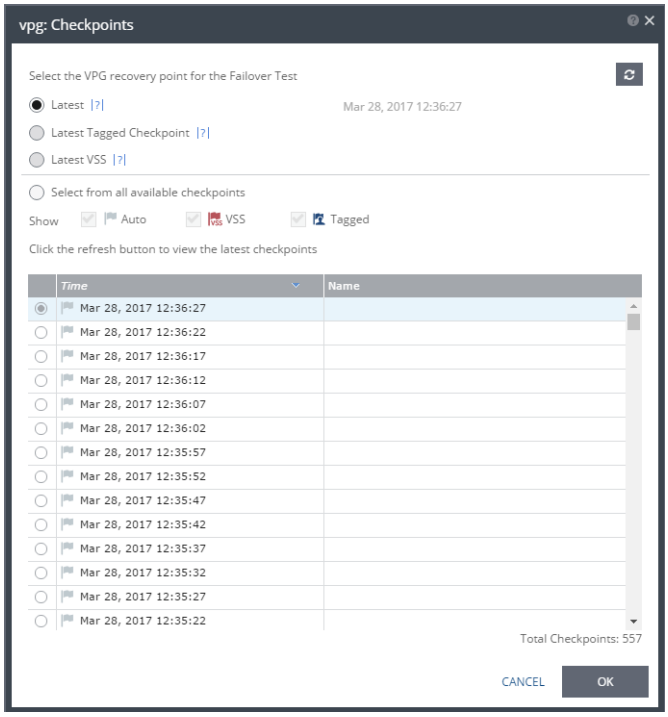

When selecting the point to recover to:

- The refresh button is initially grayed out and is enabled for clicking after 5 seconds. It is also grayed out for 5 seconds after being clicked, before being re-enabled.
- A ***Click the refresh button to view the latest checkpoints*** reminder is displayed 10 seconds after the refresh button is clicked to remind the user that there is a new *Latest Checkpoint*.
- If the user has scrolled to, and selected, a checkpoint anywhere in the checkpoints list, clicking the refresh button will automatically return the user to the selected checkpoint in the list.

4. Select the checkpoint to use:

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

   Latest Tagged Checkpoint – The recovery operation is to the latest checkpoint added by a user or when a failover test was previously performed on the VPG which includes the virtual machine or when the virtual machine was added to an existing VPG after the added virtual machine was synchronized. Checkpoints added to the virtual machine journals in the VPG ensure that the data is crash-consistent to this point. If a checkpoint is added between this point and starting the operation, this later checkpoint is not used.

   Latest VSS – When VSS is used, the clone is to the latest VSS snapshot, ensuring that the data is both crash-consistent and application consistent to this point. The frequency of VSS snapshots determines how much data can be recovered. For details about VSS checkpoints, see "Ensuring Application Consistency – Checkpoints", on page 48.

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

5. Click *OK*.
6. Click *CLONE*.

The cloning starts and the status is displayed in the VPG details tab.

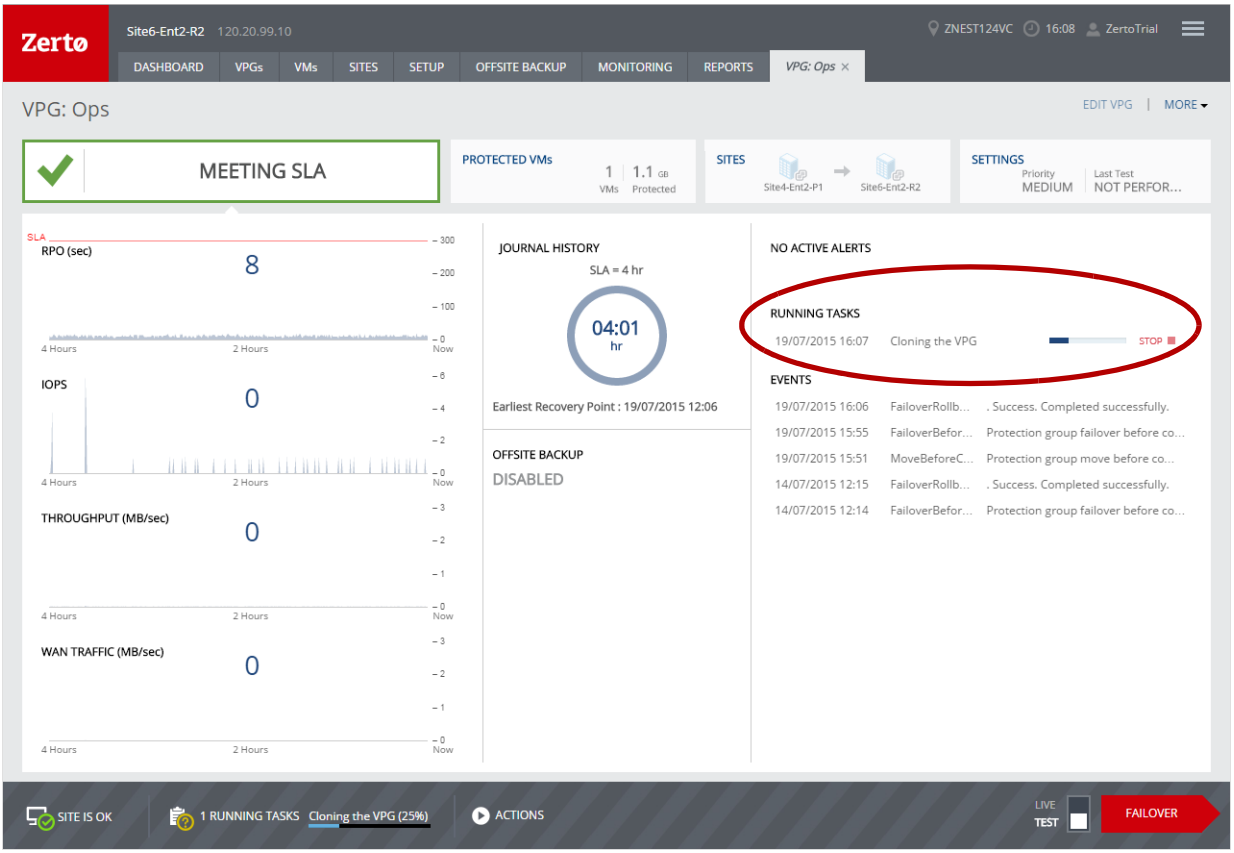

The cloned machines are named with names of the protected machines with the timestamp of the checkpoint used for the clone. The cloned virtual machines are not powered on.

Cloned volumes are recovered in EC2 as EBS disks with magnetic disk type. Virtual machines with disks that are less than 1GB are recovered with disks of 1GB. Additional volumes might be created in the cloned instance, depending on the instance type used for the clone. These volumes can be ignored.

CHAPTER 14: RECOVERING FILES AND FOLDERS

You can recover specific files and folders from the recovery site for virtual machines that are being protected by Zerto Virtual Replication and running Windows operating systems. You can recover the files and folders from a specific point-in-time. Thus, you can recover files and folders for a virtual machine for as far back as the journal history is configured.

This section describes how to recover files and folders. The following topics are described in this section:

The File and Folder Recovery Process

Use the *RESTORE FILE* operation to recover specific files and folders from the recovery site.

When you set up file and folder recovery, you always specify a checkpoint to which you want to recover the files and folders. When you select a checkpoint – either the last automatically generated checkpoint, an earlier automatically generated checkpoint, or a tagged checkpoint – Zerto Virtual Replication makes sure that the files and folders replicated at the remote site are recovered to this specified point-in-time.

The file and folder operation has the following basic steps:

- Selecting the virtual machine that is protected on which the files or folders to recover are located.
- Selecting the checkpoint at which the files and folders will be recovered.
- Selecting the disk which contains the files and folders to recover.

  Note: You can only recover files and folders from one disk at a time.
- Mounting the selected disk.
- Selecting the files and folders on the disk to recover.
- Downloading the selected files and folders. The files are downloaded to the machine where you run the Zerto User Interface. Make sure that this machine has enough space for the recovered files.
- Unmounting the selected disk.

You can only recover files and folders from one disk at a time. After the required disk is mounted, if you want to recover files or folders from another disk, you can begin the mount process for the second disk. Zerto Virtual Replication does not support mounting the same volume twice, for example if you want a file from two different checkpoints.

The operating system of the machine on which the recovery site Zerto Virtual Manager is installed determines the types of file systems from which files and folders can be recovered. When the operating system supports a file system, files and folders can be recovered from it. For example, if a protected virtual machine running Windows 2012 has files using the ReFS file system and requires one or more of these files to be recovered and the recovery site Zerto Virtual Manager is on a machine with Windows 2008, which does not support ReFS, the protected virtual machine files and folders cannot be recovered.

Note the following:

- You cannot recover files or folders from Linux machines.
- You cannot recover files or folders from a virtual machine when a test failover, live failover, move, clone, or backup is being performed on a VPG that contains the virtual machine.
- Zerto Virtual Replication does not support disks larger than 2TB.
- You cannot recover files or folders from the Zerto plugin.

Recovering Files and Folders

The procedure to recover files and folders involves the following steps:

1. "Mounting the Disk that Contains the Required Files and Folders", on page 107

   Note: While the disk is mounted:
   - If you start a live failover or move, Zerto Virtual Manager forcibly unmounts the disk so the live failover or move can be performed.
   - Manual backups fail.
   - You can perform a test failover or clone.
2. "Downloading the Files and Folders from the Disk", on page 110

Mounting the Disk that Contains the Required Files and Folders

Before you can recover files or folders, you must first select the checkpoint in time from which you will recover the files or folders. Then you must select and access the disk in which the files or folders are contained.

To mount a disk that includes files and folders to restore:

1. From either the protected or recovery site, select *ACTIONS > RESTORE FILE*.

   The *File and Folder Restore* wizard is displayed.

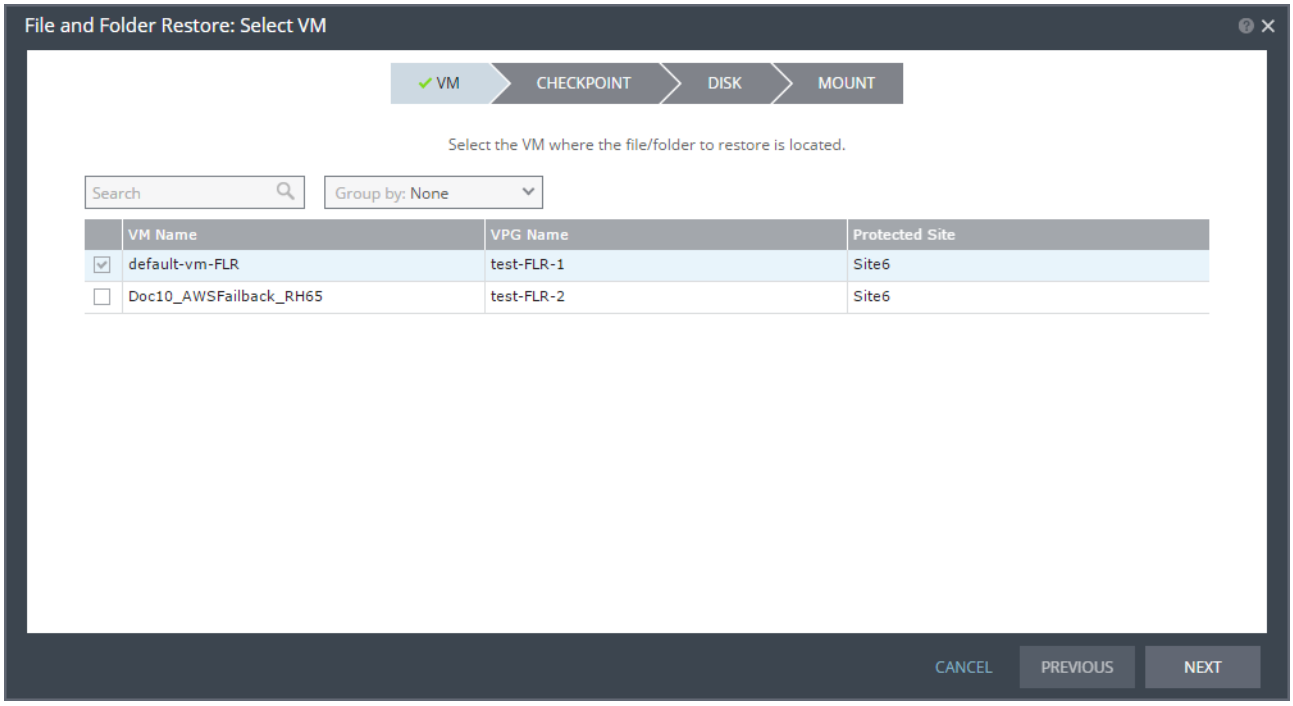

   The list of all protected virtual machines is displayed. You can only recover files or folders from one virtual machine at a time.
2. Select the virtual machine on which the file or folder is located and click *NEXT*.

The *CHECKPOINT* step is displayed. By default, all checkpoints are displayed.

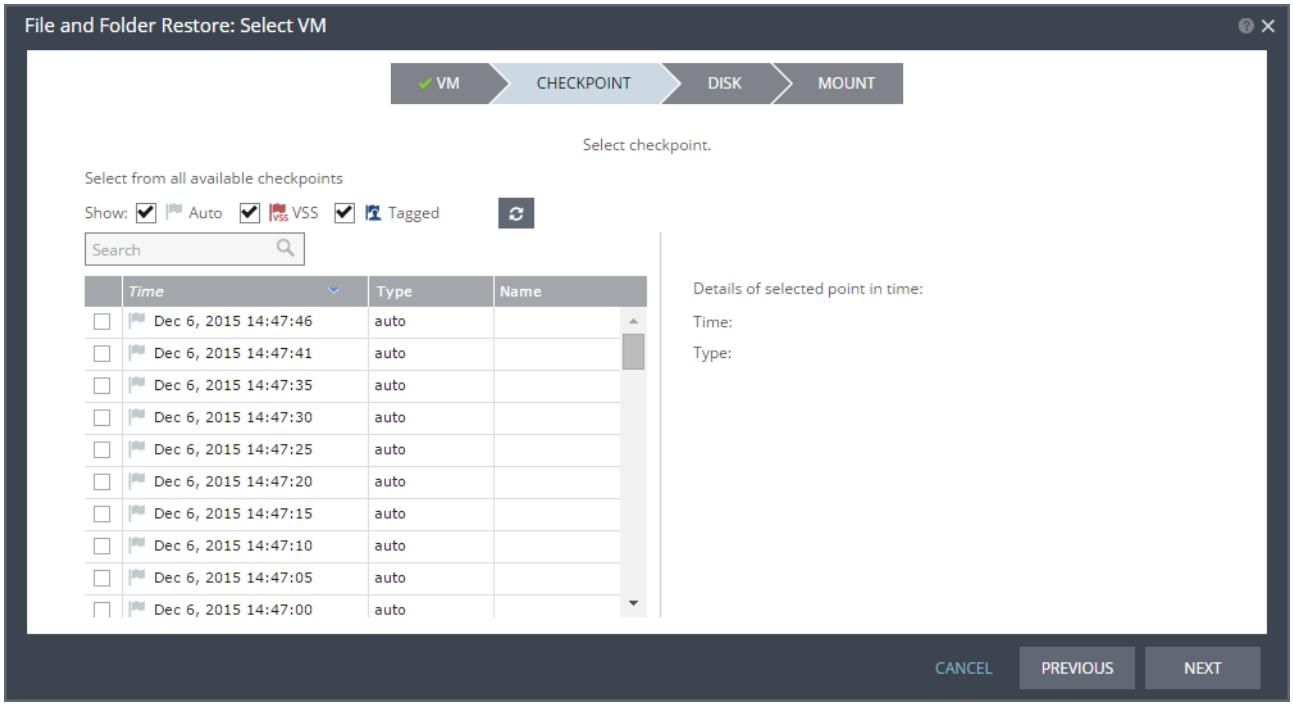

3. Select the checkpoint from which to recover the file or folder.
   - **Auto:** Checkpoints generated by the Zerto Virtual Manager are displayed.
   - **VSS:** Checkpoints that were synchronized with VSS snapshots are displayed. When VSS is used, recovery to the latest VSS snapshot ensures that the data is both crash-consistent and application consistent to this point. The frequency of VSS snapshots determines how much data can be recovered.
   - **Tagged:** Checkpoints that were added by a user, or were added by the Zerto Virtual Manager when a failover test was performed on the VPG that included the virtual machine, or when the virtual machine was added to an existing VPG after the virtual machine was synchronized.
4. Click *NEXT*.

The *DISK* step is displayed. All disks associated with the selected virtual machine are displayed.

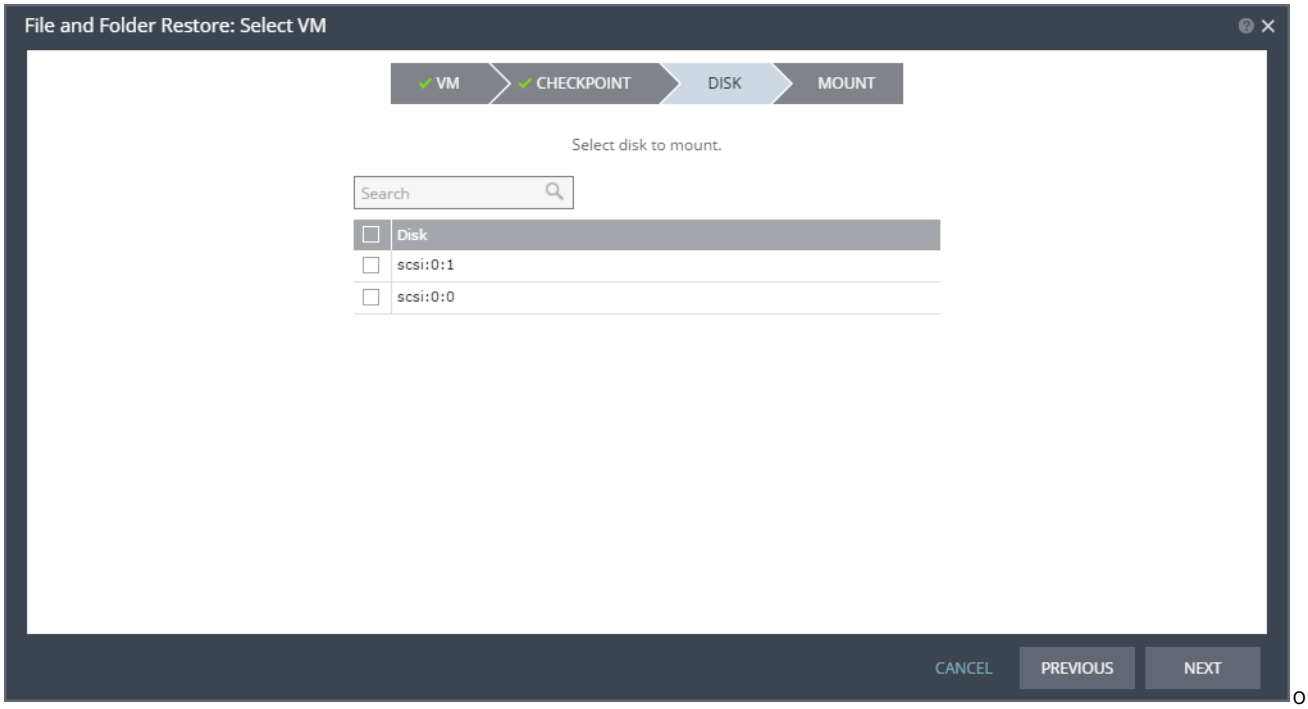

o

5. Select a disk to mount and click *NEXT*.

   The *MOUNT* step is displayed. The settings you selected are displayed.

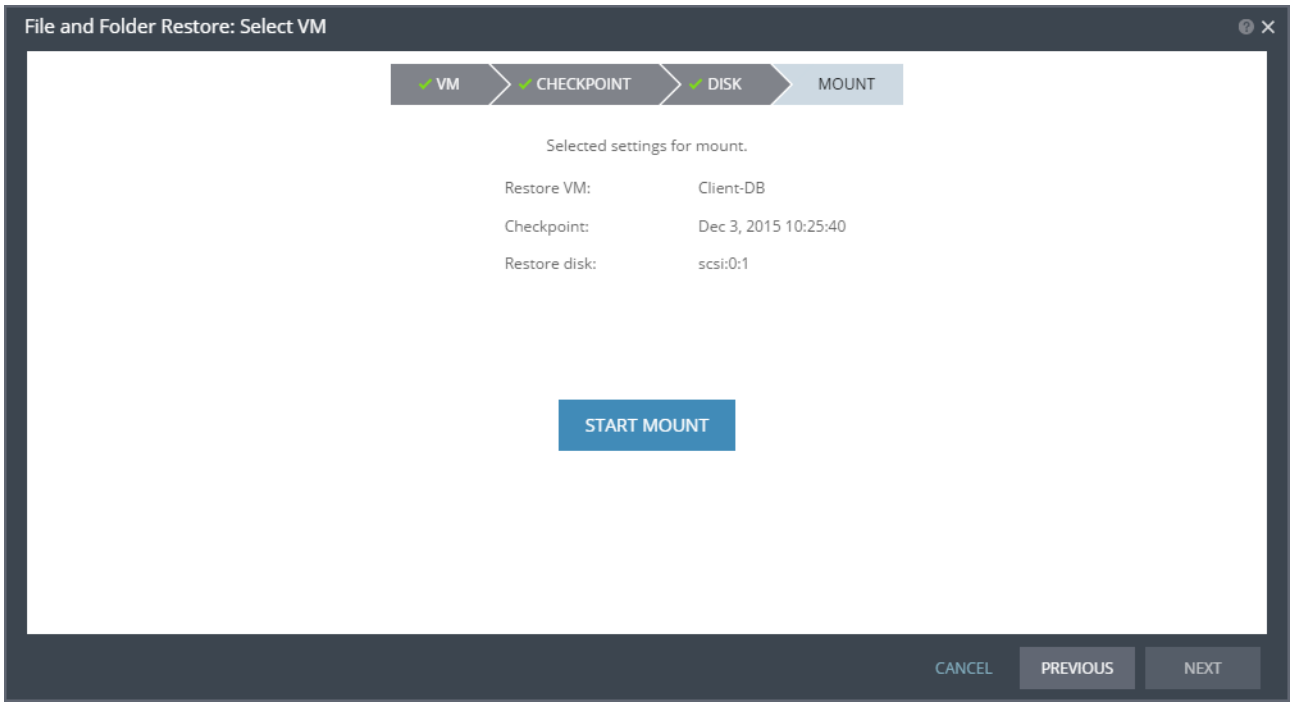

6. Click *START MOUNT* to mount the disk.

   Mounting the disk may take some time, depending on the selected checkpoint and the number of files and folders on the disk.

   - When the first part of the restore process is done, icons appear next to the completed task.
   - By clicking the folder icon ( ) you can browse the folders and files on the disk.
   - By clicking the unmount icon ( ) you can unmount the disk without restoring any files or folders.

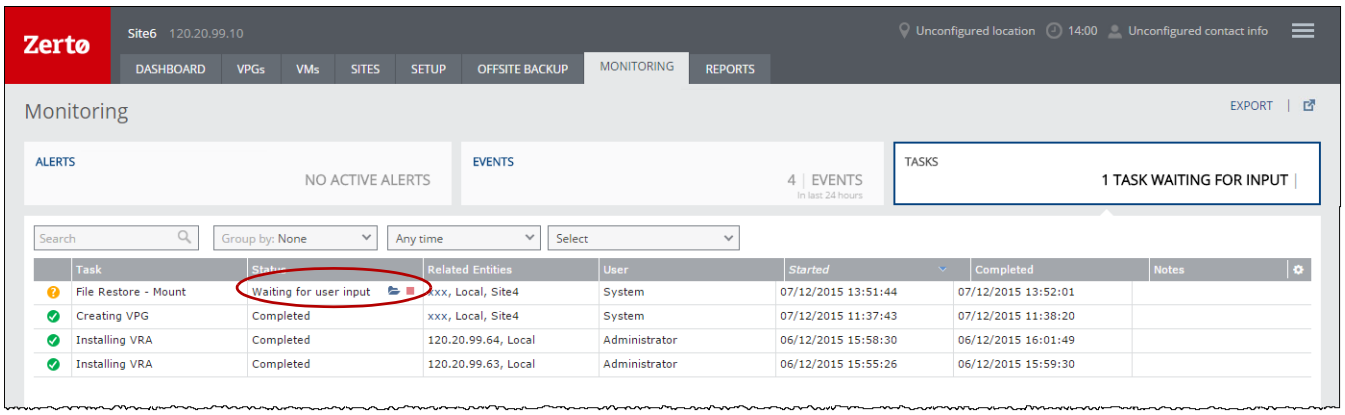

7. Continue with "Downloading the Files and Folders from the Disk" on page 110.

Downloading the Files and Folders from the Disk

In this procedure you select the files and folders. The files are downloaded to the machine where you run the Zerto User Interface. Make sure that this machine has enough space for the recovered files.

To download folders or files:

1. Click the folder icon ( ).

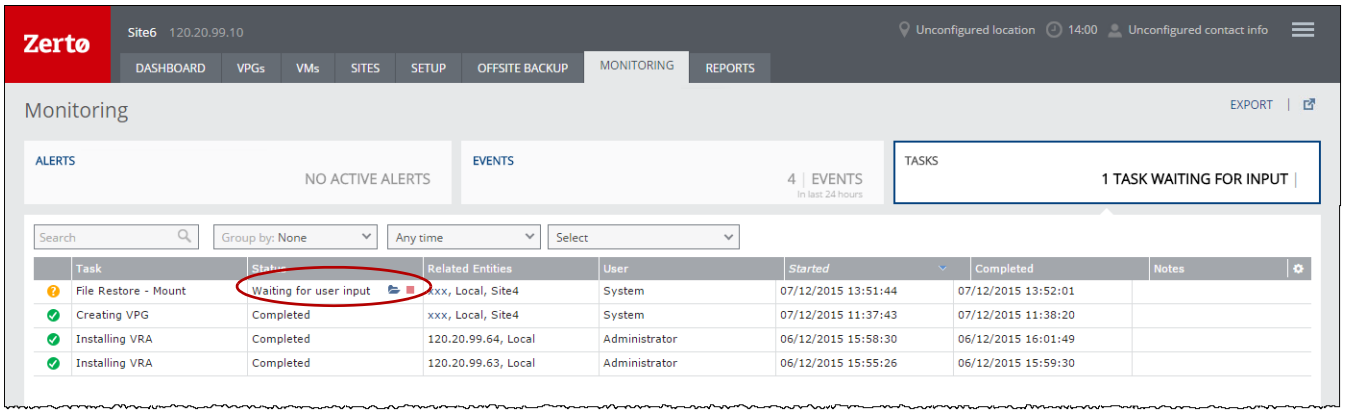

The *File and Folder Restore* dialog is displayed.

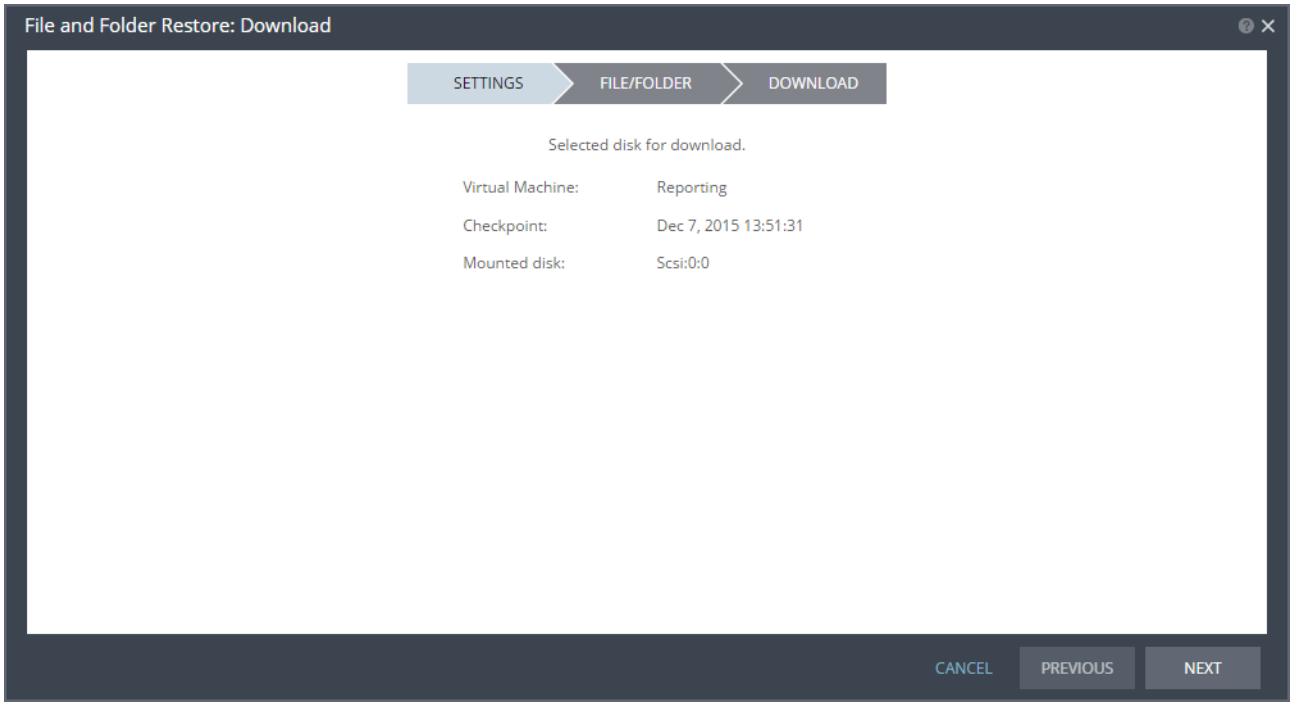

2. Click *NEXT*.
   The *FILE/FOLDER* step is displayed.

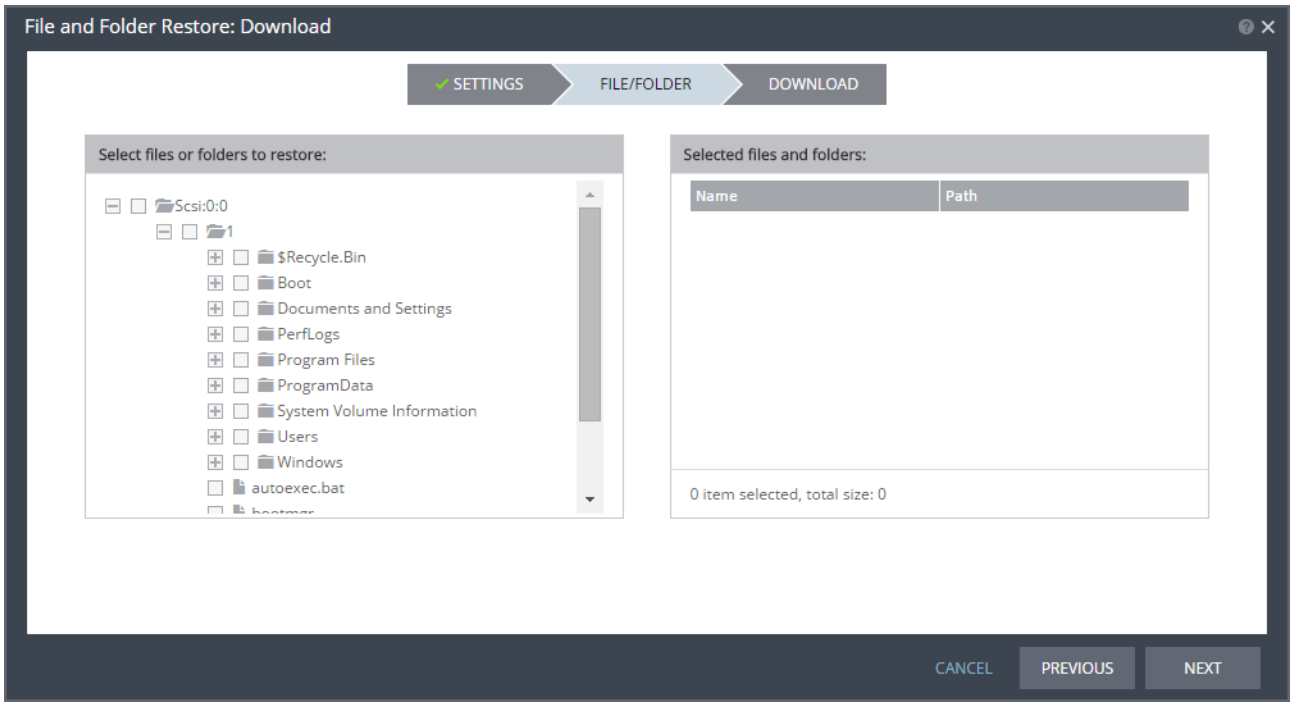

3. Select the files and folders you want to download.

The selected files or folders are displayed in the right pane of the dialog. The number of items selected is displayed and the size of the selected files is also displayed.

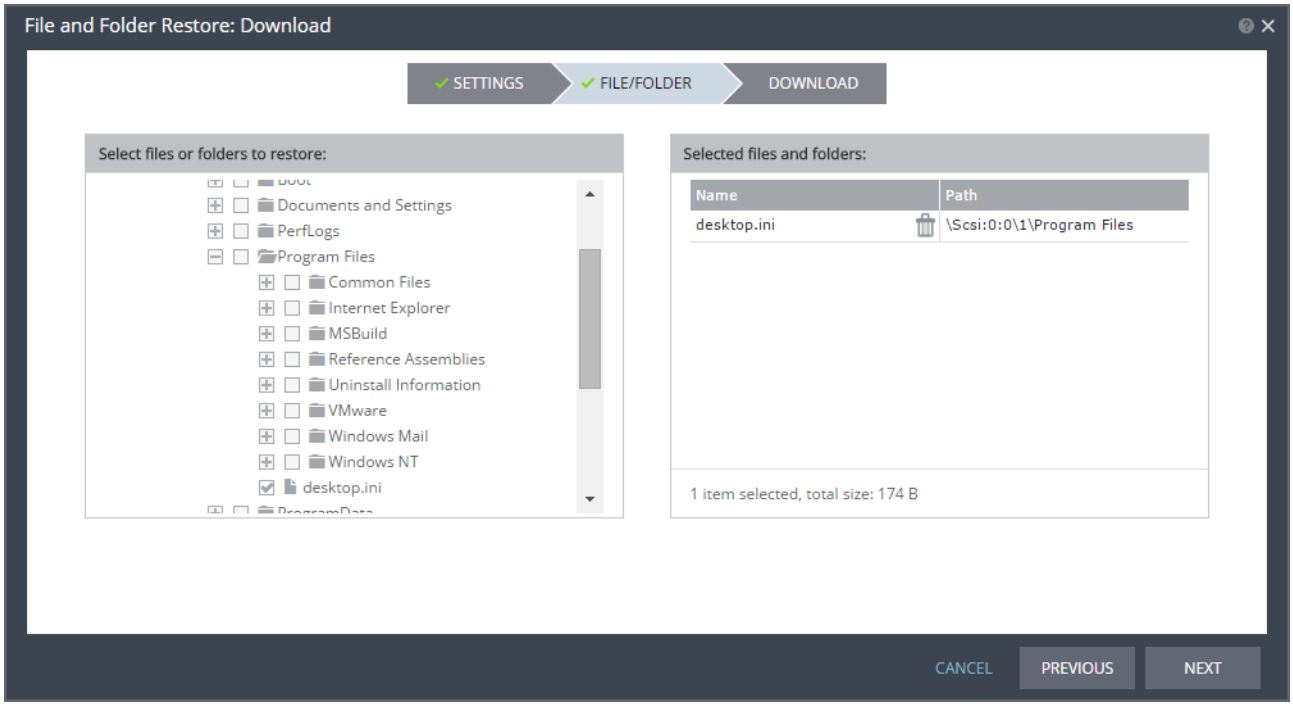

4. Click *NEXT*.

   The *DOWNLOAD* step is displayed. It shows the files and folders you selected for downloading. By default, when you select multiple files or one or more folders, the data is compressed before it is downloaded. If you select only one file, for download, you can choose whether or not the file is compressed.

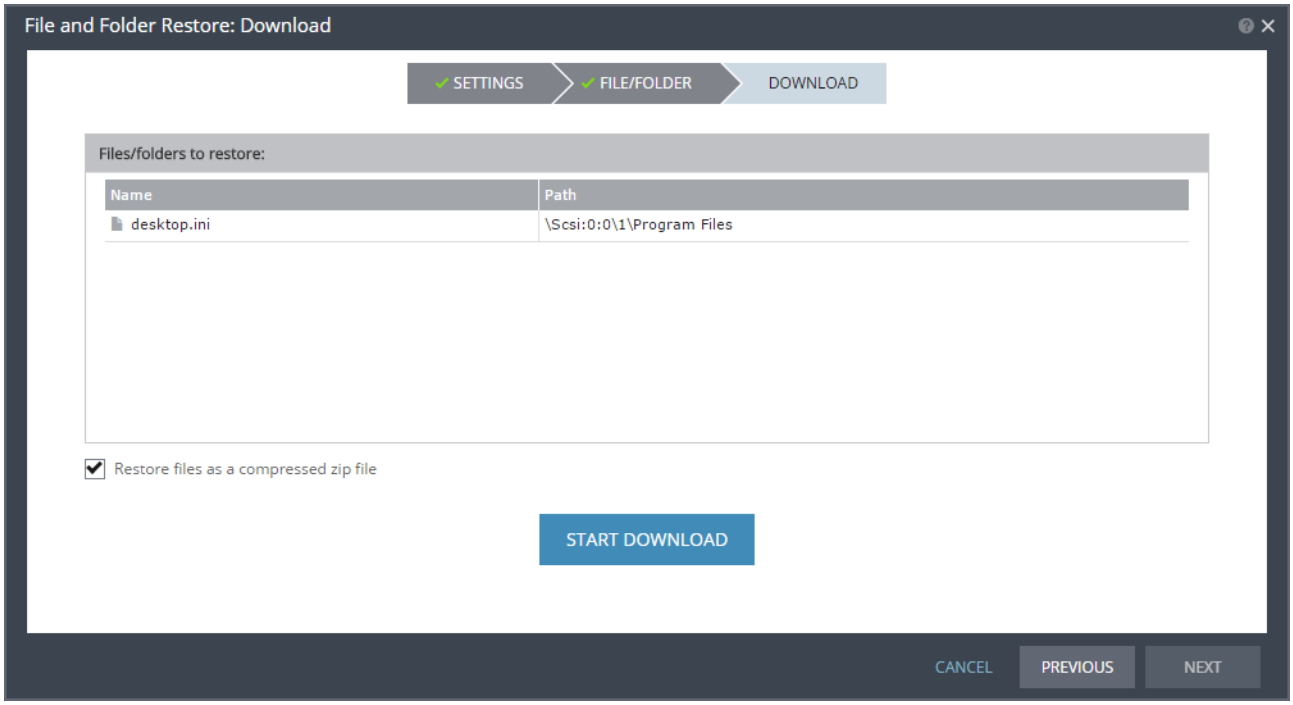

5. Click *START DOWNLOAD*.

   The files and folders are downloaded by default to the downloads folder on the computer where you run the Zerto User Interface.

Note: Saving the files and folders to a network share is dependent on the browser used to display the Zerto User Interface and the settings for this browser.

- When you select one file to download, and do not compress the file, the name of the downloaded file is the name of the file. For example, if you download a file called `Important-file.docx`, the name of the file on your computer will be `Important-file.docx`.
- When you choose one file and choose to compress it, or you select multiple files, the files are zipped into a file called `ZertoDownloads.zip`.

6. Zerto recommends that you **unmount the disk** after the files or folders are downloaded. To unmount the disk, click the unmount icon (■).

CHAPTER 15: FAILING BACK FROM AWS

If you have moved or failed over virtual machines to Amazon EC2, you can fail back those instances to VMware vSphere. This chapter describes the procedures for failing back both Windows machines and Linux machines.

The following topics are described in this chapter:

Failing Back a Windows Machine to a VMware ESXi Host

The procedure to fail back recovered Windows machines from AWS to a VMware ESXi host is done on a machine-by-machine basis; failing back is not performed for a VPG. You must use a VMware tool. The description in this section uses information derived in part from the VMware Knowledgebase article:
http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1018015

The procedure is divided into the following parts:

1. Prepare the instance for failback. See "Preparing the Instance for Failback", below.
2. Use the VMware converter to fail back the instance to VMware vSphere. See "Failing Back the Instance to vSphere", on page 118.

Preparing the Instance for Failback

1. Open the AWS EC2 Management Console and under *Instances* note the availability zone of the instance you want to fail back.
2. Select *ELASTIC BLOCK STORE > Volumes*.
3. Click *Create Volume*.

   The *Create Volume* dialog is displayed.

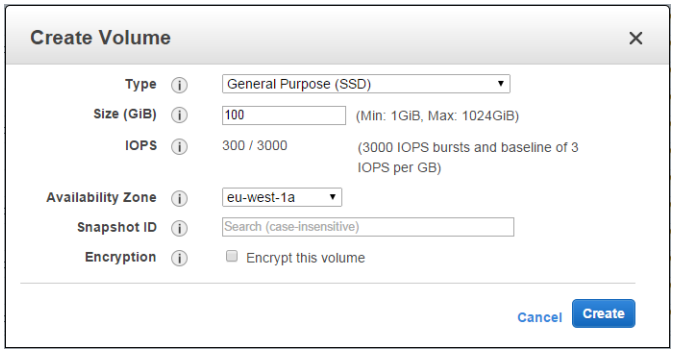

4. Enter the following:

   Type – General Purpose (SSD)

   Size – Specify a size that is at least the size of all the volumes for the virtual machine instance.

   Availability Zone – Select the availability zone of the instance you want to fail back.

   Leave both **Snapshot ID** and **Encryption** with the default values.
5. Click *Create*.

   It may take a few minutes for AWS to create the volume.
6. After the volume is created, attach it to your instance:
   a) Select the volume you just created.
   b) Select *Actions > Attach Volume*.

The *Attach Volume* dialog is displayed.

c) Choose the instance to fail back from the drop-down list.
d) Click *Attach*.

AWS attaches the new volume to the instance you selected.

7. Reboot the instance to see that the new volume is attached to the instance.

   In AWS, you can see that the volume is attached to the instance by selecting the volume and checking the *Attachment information* in the bottom section of the screen.

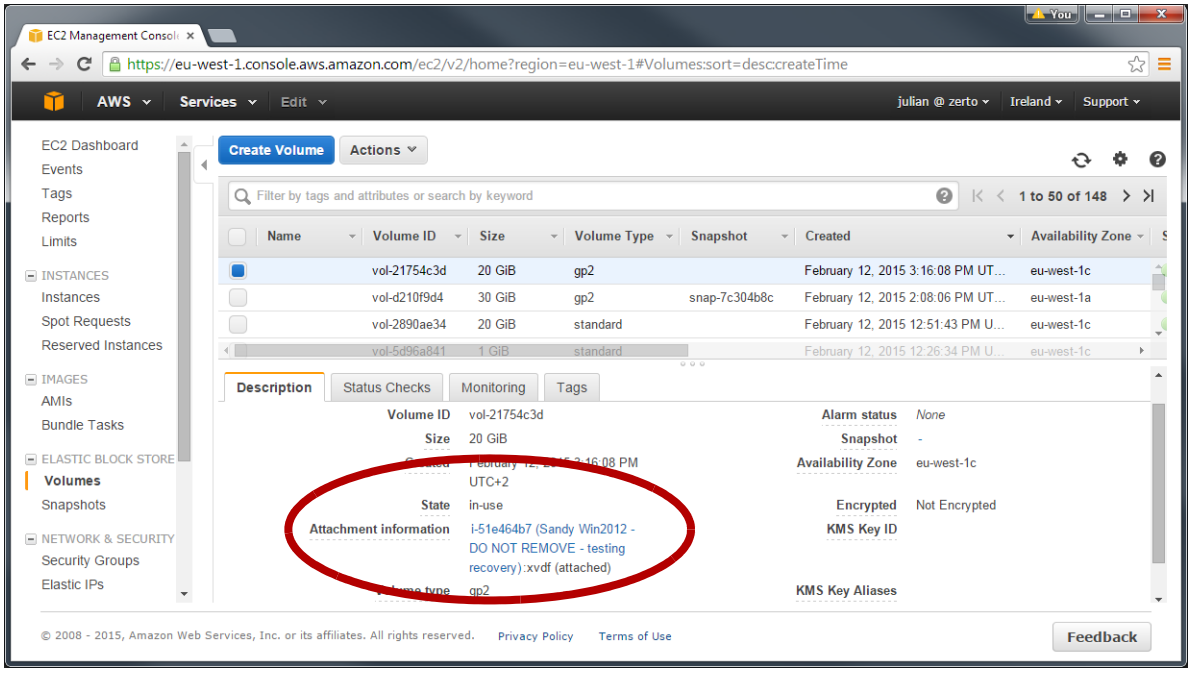

8. Download the VMware vCenter converter from https://my.vmware.com/web/vmware/evalcenter?p=converter and copy it to the new volume on the instance to fail back.

   You must open an account with VMware to download the software, which is free. The following documentation is based on VMware vCenter Converter Standalone version 5.5.3.

9. Run the VMware vCenter Converter Standalone and follow its instructions until the *Setup Type* step.
10. Select **Local installation** and click *Next*.
11. In the *Ready to Install* step, click *Install*.

    The converter is installed on the instance.

12. Select *Run Converter Standalone Client now* and click *Finish*.

The VMware vCenter Converter Standalone opens.

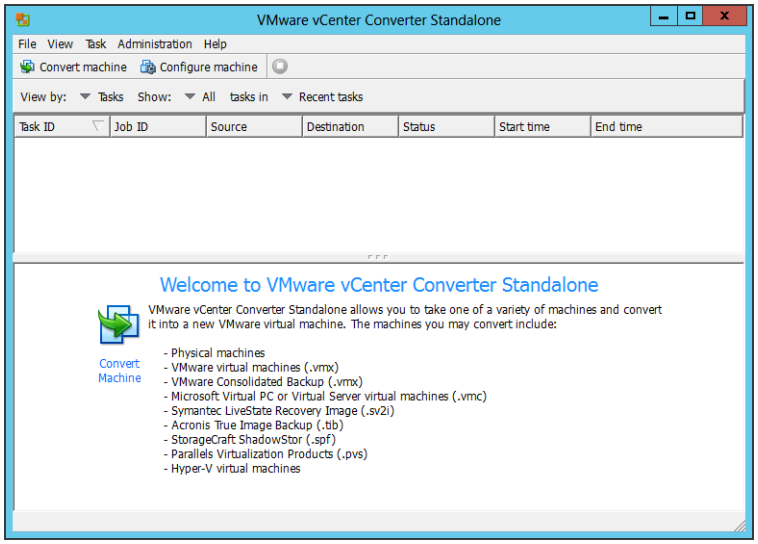

13. Click *Convert machine*.

    The *Source System* page is displayed.

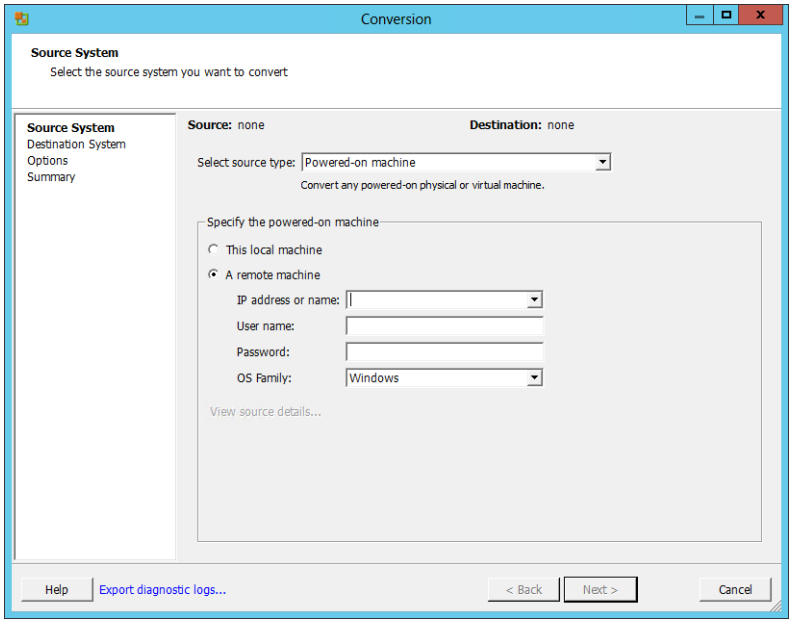

14. Select the following:

    Select source type – Select *Powered-on machine*

    This local machine – Select the local machine.

15. Click *Next*.

The *Destination System* page is displayed.

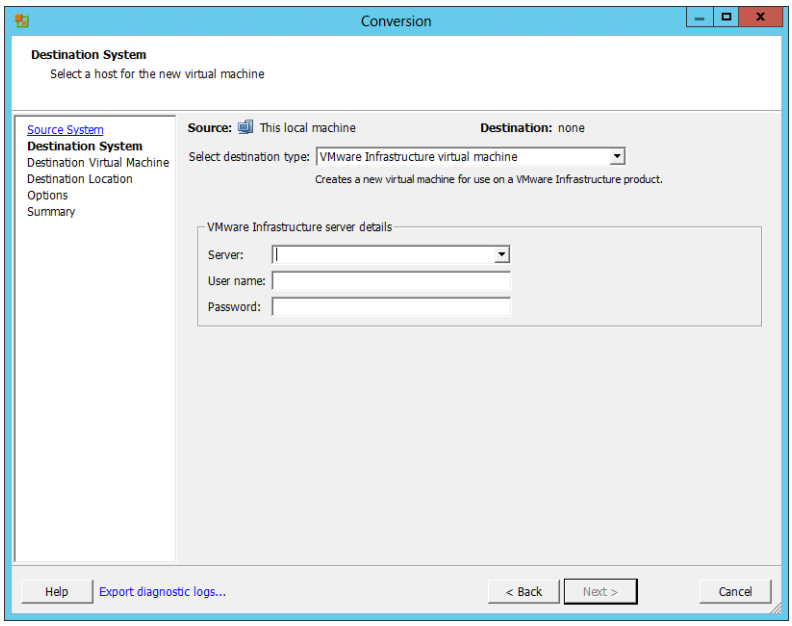

16. Select *VMware Workstation or other VMware virtual machine* for `Select destination type`.

    The *Destination System* page is redisplayed

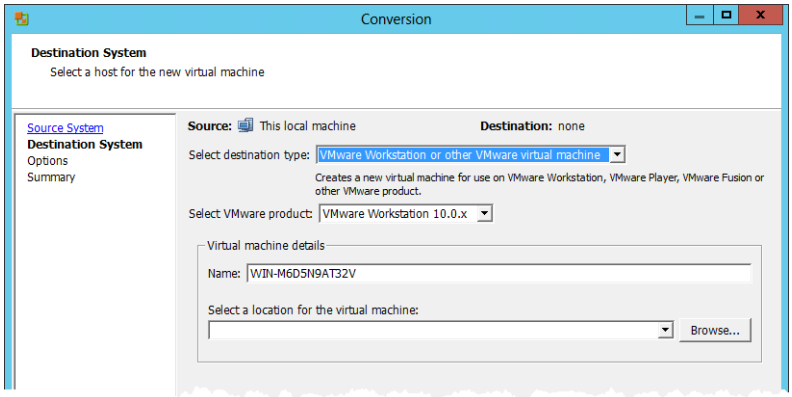

17. Set the following:

    Select VMware product – Select the VMware Workstation version you need. Zerto recommends using the latest version.

    Name – Enter a name to assign to the machine to be created in vCenter Server.

    Select a location for the virtual machine – Select the volume you created in steps 3-5 as the location of the new instance that will be created.

18. Click *Next*.

The *Options* page is displayed.

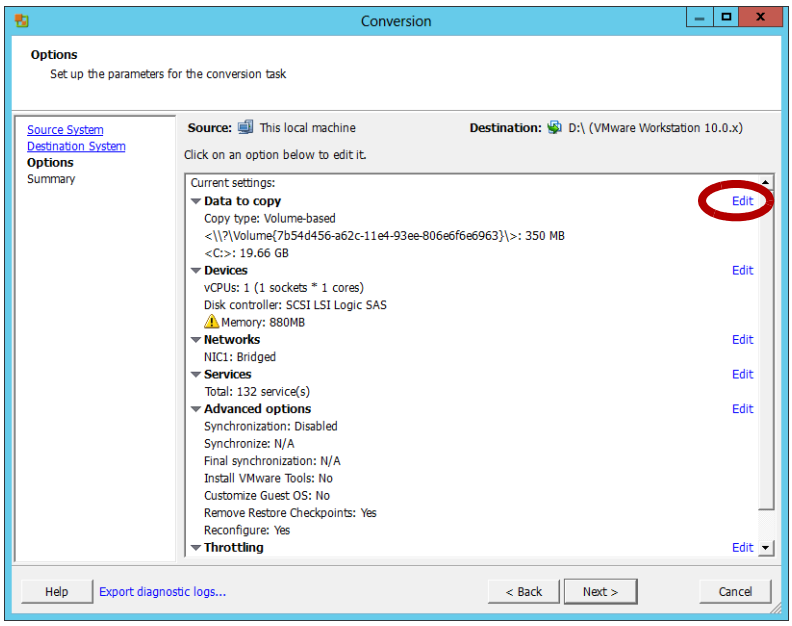

19. Check which volumes will be copied by clicking *Edit* for *Data to copy*. Make sure the checkbox next to the volume you created in steps 3-5 is **not** checked. This ensures that this volume will not be included in the conversion process.

    Make sure that all other check boxes **are** checked, to ensure that all other volumes will be included in the conversion process.

    Check the other settings and, if necessary, edit them.
20. Click *Next*.
21. In the *Summary* step, review the conversion parameters that will be used.
22. Click *Finish*.

The instance and its volumes are copied to the new volume.

Failing Back the Instance to vSphere

1. In VMware vCenter Converter Standalone, click *Convert machine*.
2. In the *Source System* page, specify the following:

   Select source type – Select *VMware Workstation or other VMware virtual machine*

   Virtual machine file – Select the `<name>.vmx` file that was created, where `<name>` is the name you gave in step 17 in the section Preparing the Instance for Failback.
3. Click *Next*.
4. In the *Destination System* page, specify the following:

   Select destination type – Select *VMware infrastructure virtual machine*.

   Server – Enter the IP address of the VMware vCenter Server to which the new machine will be copied.

   User name – Enter the user name of the VMware vCenter Server.

   Password – Enter the password of the VMware vCenter Server.
5. Click *Next*.

The converter connects to the VMware vCenter Server you specified in the previous step, step 4, and the *Destination Virtual Machine* page is displayed.

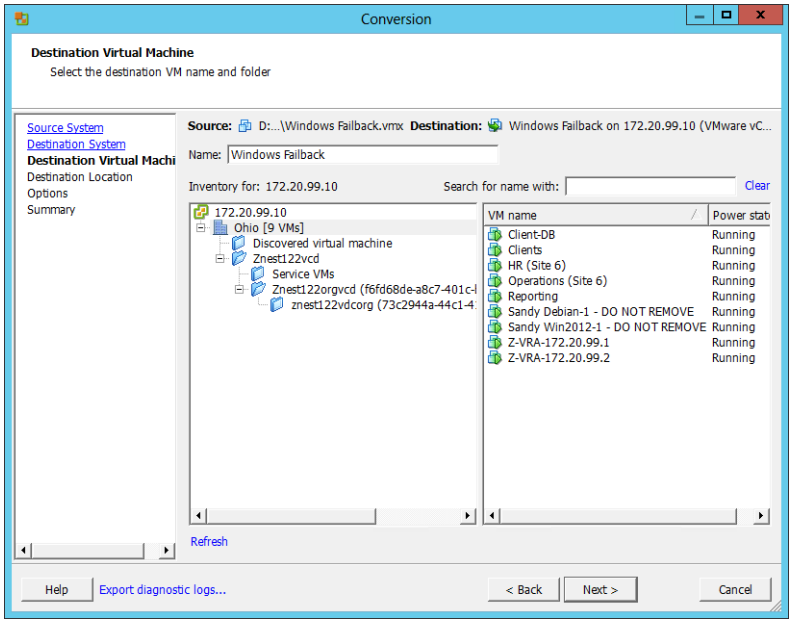

6. In *Name*, enter a name for the new machine.
7. Click *Next*.

   The *Destination location* page is displayed.

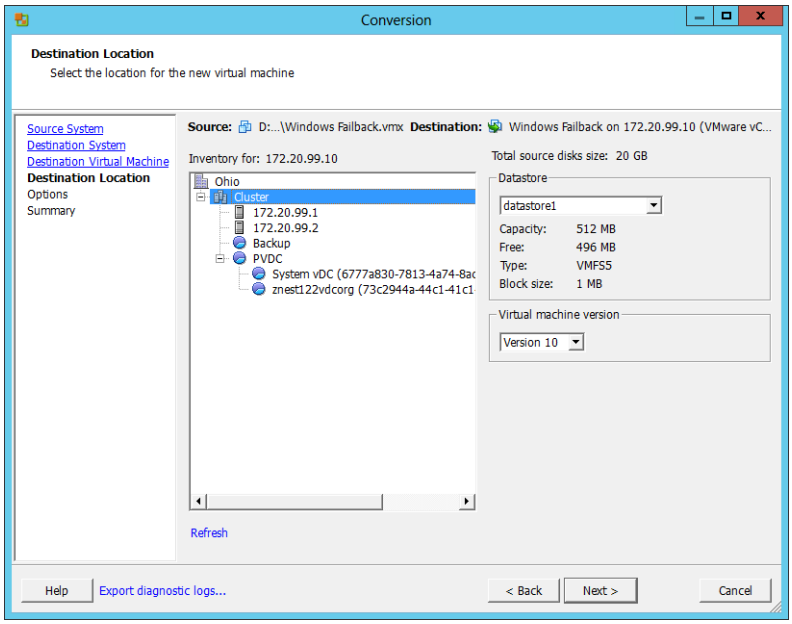

8. Choose the host and datastore and virtual machine version.
9. Click *Next*.
10. In the *Options* step, check the values of the settings and, if necessary, edit them. For example, in the *Data to copy* section, check whether the disk should be thick or thin provisioned.
11. Click *Next*.
12. In the *Summary* step, review the conversion parameters that will be used. Click *Finish*.

    The instance in AWS is converted to a virtual machine in your vCenter Server. The virtual machine is not powered on.

Failing Back a Linux Instance to a VMware ESXi or Microsoft Hyper-V Host

The procedure to failback a Linux instance from AWS to either a VMware ESXi host or a Microsoft Hyper-V host is done on a machine-by-machine basis; failing back is not performed for a VPG. The description in this section uses information derived from the following Amazon documentation:

- http://docs.aws.amazon.com/AWSEC2/latest/CommandLineReference/set-up-ec2-cli-windows.html
- http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ExportingEC2Instances.html

Prerequisites to Set Up a PC to Failback a Recovered Linux Machine from AWS

In order to failback a Linux machine you require the following software on the PC:

- The Amazon EC2 CLI tools that you use to export the Linux machine from EC2. These tools are available from Amazon EC2 CLI Tools.
- Java version 1.7 or higher. Either the runtime environment, JRE, or developer environment, JDK, must be available. Java is available from the Java web site.

You also require the following user variables definitions on the PC:

- Define the following user variables:
  - JAVA_HOME
  - EC2_HOME
  - AWS_ACCESS_KEY
  - AWS_SECRET_KEY
    The Amazon EC2 CLI tools use your access keys to identify you. There are two types of access keys: access key IDs, for example, `AKIAIOSFODNN7EXAMPLE`, and secret access keys, for example, `wJalrXUtn+7xRfiCYEXAMPLEKEY`.

  You define these variables as follows:
  a) Click *Start*, right-click *Computer*, and select *Properties*.
  b) Select *Advanced system settings*.
  c) Click *Environment Variables*.
  d) Under *User variables*, click *New*.
  e) In *Variable name*, type `JAVA_HOME`
  f) In *Variable value*, type the path to your Java home, for example, `C:\Program Files (x86)\Java\jre1.8.0_45`
  Don't include the bin folder in JAVA_HOME.
  g) Click *OK*.
  h) Repeat steps d to g for the other variables:
    - `EC2_HOME`, set the *Variable name* value to the path of the folder into which you saved the Amazon EC2 CLI tools, for example, `C:\Tools\AWS Tools\ec2-api-tools-1.7.4.0`
    - `EC2_URL`, set the *Variable name* value to the region in AWS where the instance to be failed back is located, for example, `https://ec2.eu-west-1.amazonaws.com`
    - `AWS_ACCESS_KEY`, set the *Variable name* value to the AWS access key ID.
    - `AWS_SECRET_KEY`, set the *Variable name* value to the AWS secret key.
  i) Select the *Path* variable and click *Edit* and add the following to the beginning of the *Variable value*:
  `;%EC2_HOME%\bin;%JAVA_HOME%\bin`
  Add the bin folder that contains the Java executable to the path before other versions of Java.
  j) Click *OK* and then click *OK* to exit the *Environment Variable* dialog.
  k) Verify that the environment is set up correctly: Open a new Command Prompt window and enter the following:
    - `C:\> java -version`
      Verifies the JAVA_HOME and path are set correctly. The output from the command should be similar to the following:
      ```
      java version "1.8.0_45"
      Java(TM) SE Runtime Environment (build 1.8.0_45-b15)
      Java HotSpot(TM) Client VM (build 24.45-b02, mixed mode, sharing)
      ```
    - `C:\> dir "%EC2_HOME%"`
      Verifies the EC2_HOME and path are set correctly. The output from the command is the folder listing where you saved the Amazon EC2 CLI tools.
    - `C:\> ec2-describe-regions`

Verifies the EC2 environment variables are set correctly. The output from the command should be similar to the following:

```
REGION  us-east-1        ec2.us-east-1.amazonaws.com
REGION  eu-west-1        ec2.eu-west-1.amazonaws.com
REGION  sa-east-1        ec2.sa-east-1.amazonaws.com
REGION  ap-northeast-1   ec2.ap-northeast-1.amazonaws.com
REGION  us-west-2        ec2.us-west-2.amazonaws.com
REGION  us-west-1        ec2.us-west-1.amazonaws.com
REGION  ap-southeast-1   ec2.ap-southeast-1.amazonaws.com
```

If you get an error that any of these commands are not recognized as an internal or external command, check the setting of Variable name and Variable value settings and the Path setting. Fix any errors, open a new Command Prompt window, and try the command again. If you get an error that the required option `-O` is missing, check the setting of AWS_ACCESS_KEY. If you get an error that the required option `-W` is missing, check the setting of AWS_SECRET_KEY. If you get a `Client.AuthFailure error`, check that you've entered your AWS_ACCESS_KEY and AWS_SECRET_KEY correctly, and check that the date and time are set correctly on your computer.

- Define an S3 bucket for the failed back instance to store the exported instances. To create a bucket:
  a) Sign into the AWS Management Console and open the Amazon S3 Management Console.
  b) Click *Create Bucket*.
  c) In the *Create Bucket* dialog box, enter a name for the bucket and select the region where you want the bucket to reside. The name of the bucket must be unique across all existing bucket names in Amazon S3. The bucket name can contain lowercase letters, numbers, periods (.), and hyphens (-), must start with a number or letter and be between 3 and 63 characters long. The name cannot include underscores (_), two, adjacent periods or dashes next to periods and must not end with a hyphen.
  d) Click *Create*.
  e) Expand *Permissions*, and click *Add more permissions*.
  f) In the new line, enter *vm-import-export@amazon.com as the Grantee* and check *Upload/Delete* and *View Permissions*.
  g) Click *Save*.

Failing Back a Recovered Linux Machine from AWS

The same Amazon EC2 CLI tool is used to fail back a Linux machine to either VMware or Hyper-V.

To failback a Linux machine to VMware:

1. In a Command Prompt window run the following EC2 CLI command:

```
ec2-create-instance-export-task instance_id -e target_environment -f disk_image_format
-c container_format -b bucket_name
```

PARAMETER	DESCRIPTION
instance_id	The ID of the instance to fail back.
-e *target_environment*	The target environment. Set this parameter to `VMware`
-f *disk_image_format*	The format used for disk images. Set this parameter to `VMDK`
-c *container_format*	The container format used to combine disk images with metadata. If absent, only the disk image is exported. Set this parameter to `OVA`
-b *bucket_name*	The name of the destination Amazon S3 bucket where the file will be exported.

The instance is exported to the S3 bucket as an OVA file.

Note: You can monitor the export of your instance, at the command prompt, by typing the following command:

```
ec2-describe-export-tasks
```

Or, for the specific task:

```
ec2-describe-export-tasks filename
```

where *filename* is the name of the file for the instance in the S3 bucket.

Monitor the export to determine when the export has completed.

2. In the S3 bucket, right-click the file and select *Download*.
3. Right-click the Download link and save the OVA file to your site, for example, in Internet Explorer via *Save target as* and in Google Chrome via *Save Link as*.
4. In the VMware Web Client or the VMware Client console, deploy the OVA file.
   a) Select *File > Deploy OVF Template*.
   b) In the *Deploy OVF Template* wizard, select the OVA file as the OVF package.
   c) Continue through the wizard.

The Linux machine is deployed in the vCenter Server. By default, the name of the machine is the name of the file for the instance in the S3 bucket.

To failback a Linux machine to Hyper-V:

1. In a Command Prompt window run the following EC2 CLI command:

   ```
   ec2-create-instance-export-task instance_id -e target_environment -f disk_image_format
   -b bucket_name
   ```

PARAMETER	DESCRIPTION
instance_id	The ID of the instance to fail back.
-e *target_environment*	The target environment. Set this parameter to `Microsoft`
-f *disk_image_format*	The format used for disk images. Set this parameter to `VHD`
-b *bucket_name*	The name of the destination Amazon S3 bucket where the file will be exported.

   The instance is exported to the S3 bucket as a VHD file.

   Note: You can monitor the export of your instance, at the command prompt, by typing the following command:

   ```
   ec2-describe-export-tasks
   ```

   Or, for the specific task:

   ```
   ec2-describe-export-tasks filename
   ```

   where *filename* is the name of the file for the failed back instance in the S3 bucket.
2. In the S3 bucket, right-click the file and select *Download*.
3. Right-click the Download link and save the VHD file to your site, for example, in Internet Explorer via *Save target as* and in Google Chrome via *Save Link as*.
4. In the Hyper-V Manager, right-click the Hyper-V host and select *New > Virtual Machine*.
   a) Specify the generation to use for the machine and the startup memory.
   b) Continue through the *New Virtual machine* wizard until the Connect Virtual Hard Disk step and in this step select to use an existing virtual hard disk and specify the file downloaded from the S3 bucket.
   c) Click *Finish*.

The Linux machine is deployed under the Hyper-V host.

CHAPTER 16: OFFSITE BACKUP CONFIGURATION

When using Zerto Virtual Replication for offsite backups, you must first define the repository that will be used to store the backups.

Disaster recovery using Zerto Virtual Replication enables recovering from a disaster to any point between the moment just before the disaster and a specified amount of time in the past up to 30 days. The recovery is done in real time at the recovery site with a minimal RTO.

If there is an additional requirement to extend the recovery ability to more than 30 days, Zerto Virtual Replication provides an offsite back up option that enables saving the protected virtual machines offsite for up to one year in a state where they can be easily deployed.

The virtual machine files are saved in a repository for the required period. Each virtual machine can have multiple offsite backups created according to a fixed schedule.

The offsite backups are managed by a Windows service, the Virtual Backup Appliance (VBA). The VBA is installed as part of the Zerto Virtual Replication installation. During an offsite backup, the VBA communicates with the VRAs on the recovery site to create the virtual machine files, such as the configuration and virtual disk files in a repository. The offsite backups are fixed points saved either weekly or monthly in the repository. Before you can create an offsite backup for virtual machines, you must first create one or more repositories for the offsite backup jobs.

The following offsite backup set up options are described in this chapter:

Creating an Offsite Backup Repository

You define the repositories where offsite backups are defined on the recovery site and can be stored, either locally at the recovery site, or on a network share that uses the SMB, Server Message Block, protocol. The repository where you want this offsite backup stored is specified when an offsite backup is defined.

To create an offsite backup repository:

1. In the Zerto User Interface, click *SETUP* > *REPOSITORIES*.
2. Click *NEW REPOSITORY*.

The *New Repository* dialog is displayed.

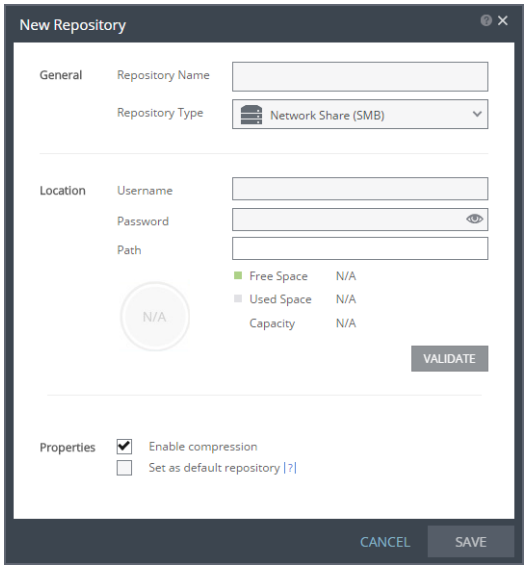

3. Specify the following settings:

   Repository Name – Specify a unique name for the repository.

   Repository Type – Specify the type of repository. The options are *Local* or *Network Share (SMB)*. If *Local* is specified, backups are stored on the local machine where the Zerto Virtual Manager is installed. If *SMB* is specified, the network share drive must be an SMB drive and if specified the username and password to access the drive must be provided. If the repository location is a network drive, this drive can be mounted to third party storage, such as Amazon Web Services (AWS). Using TntDrive, from Amazon, enables you to save your offsite backups to a cloud repository mounted disk as if you are using a LAN or locally mounted drive. You can mount one or more Amazon S3 buckets as network drives or as removable local drives, and to use them exactly as you would use any other drive folder on your computer.

   Username – Username to access the Network Share drive. The name can be entered using either of the following formats:
   - `username`
   - `domain\username`

   This field is not displayed when the type is `Local`.

   Password – Password to access the Network Share drive. This field is not displayed when the type is `Local`.

   Path – The path where the repository will reside. The path must be accessible from the Zerto Virtual Manager, so if the repository is on a different domain to the Zerto Virtual Manager, the domain must be included in the path.

   Enable Compression – Check this option to compress backups stored in the repository. Compression is done using zip compression, set to level six. If you want better compression, which requires more CPU, or less compression to reduce the CPU overhead, contact Zerto support.

   Note: Compression usually reduces the effectiveness of deduplication on stored data. If the backup repository resides on a deduplication-enabled storage appliance, it is recommended that the data be stored uncompressed.

   Note: Backup to TntDrive with compression enabled is not supported.

   Set as Default Repository– Check if you want the repository to be used as the default when specifying extended recovery in a VPG.
4. Click *VALIDATE*.You must validate the path specified. If the folder does not exist, you are asked if you want to create it.
5. Click *SAVE*.

   The repository is created.

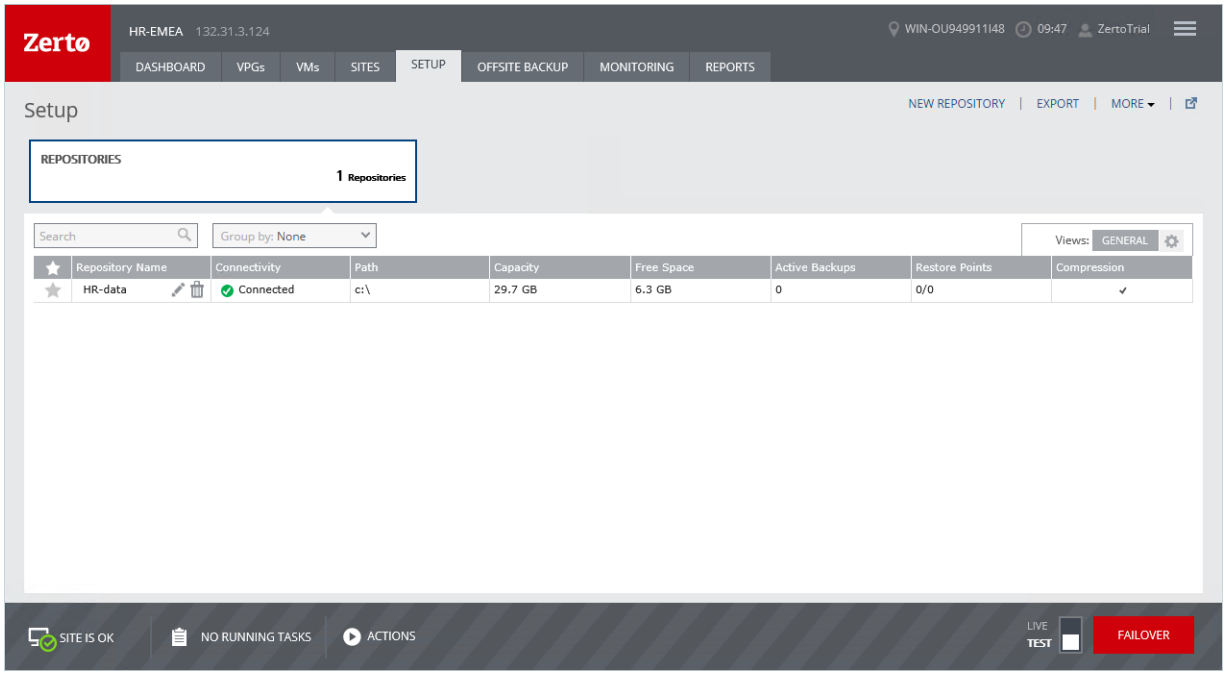

You can define more than one repository. When defining offsite backup, you specify which repository to use.

Editing an Offsite Backup Repository

You edit the repositories from the *Repositories* tab.

To edit an offsite backup repository:

1. In the Zerto User Interface, click *SETUP > REPOSITORIES*.
2. Select the repository to edit and click the edit, pencil, icon.

   The *Edit Repository* dialog is displayed.

   Edit any of the following settings:

   Repository Name – Specify a unique name for the repository.

   Repository Type – Either specify that the repository resides on a local or shared network disk, using the SMB protocol, accessible from the recovery site. If the repository location is a network drive, this drive can be mounted to third party storage, such as Amazon Web Services, AWS, or Microsoft Azure.

   Username – Username to access the Network Share drive. The name can be entered using either of the following formats:
   - `username`
   - `domain\username`

   This field is not displayed when the type is `Local`.

   Password – Password to access the Network Share drive. This field is not displayed when the type is `Local`.

   Path – The path from the recovery site where the repository will reside. The path must be accessible from the Zerto Virtual Manager, so if the repository is on a different domain to the Zerto Virtual Manager, the domain must be included in the path.

   Enable compression – Check this option to compress backups stored in the repository. Compression is done using zip compression, set to level six. If you want better compression, which requires more CPU, or less compression to reduce the CPU overhead, contact Zerto support.

   Note: Compression usually reduces the effectiveness of deduplication on stored data. If the backup repository resides on a deduplication-enabled storage appliance, it is recommended that the data be stored uncompressed.

   Note: Backup to TntDrive with compression enabled is not supported.

Set as default repository– Check if you want the repository to be used as the default when specifying extended recovery in a VPG.

3. Click *VALIDATE*.You must validate the path specified. If the folder does not exist, you are asked if you want to create it.
4. Click *SAVE*.

The updated definition of the repository is saved.

CHAPTER 17: ZERTO VIRTUAL REPLICATION REPORTS

Zerto Virtual Replication includes reporting for the following:

Recovery Reports

Information about recovery operations — failover tests, moves, and failovers — can be displayed in Recovery Reports under the REPORTS tab.

The information includes the name of the user who initiated the report, which recovery operation, the point in time , protected and the recovery sites involved, when the recovery operation was started, when it ended, the time it took to bring up the machines in the recovery site, the RTO, whether the operation succeeded or not, the VPG recovery settings, the virtual machine recovery settings, and detailed recovery steps, and any notes added during a failover test

Recovery Reports are always kept, and **never deleted**.

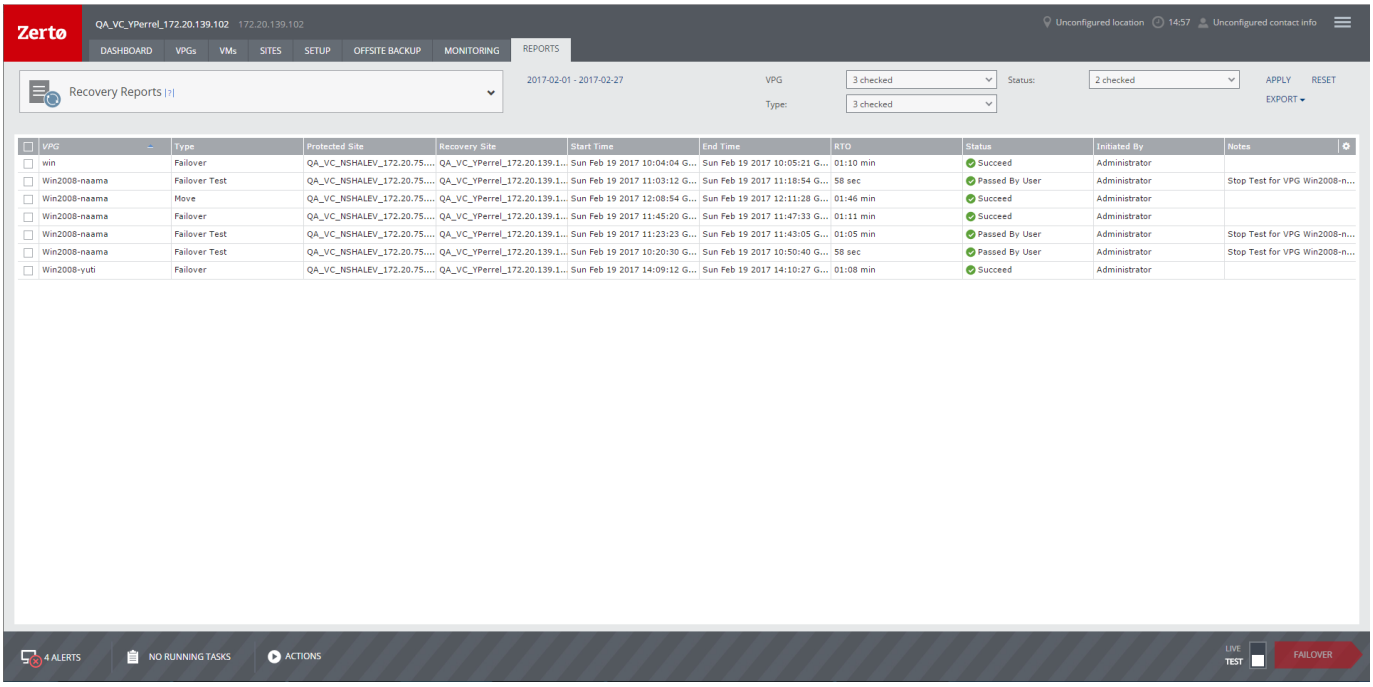

You can filter the tests by the following:

Dates: The dates for which you want information. Only operations performed between these dates are displayed.

VPG: Select the VPGs for which you want information. The number of VPGs you selected is displayed. If you select `All`, the total number of VPGs is shown.

Type: Select the recovery operations for which you want information: `Failover`, `Move`, `Failover Test`. If more than one operation is selected, the number of recovery operations you selected is displayed.

Status: Select the statuses for which you want information: `Success`, `Failed`. If more than one status is selected, the number of statuses you selected is displayed.

Click **APPLY** to apply the selected filtering.

Click **RESET** to reset the display to the default values.

Click **EXPORT** and choose PDF or ZIP to generate a report.

The report displays information by VPG and then by virtual machine within the VPG. The VPG information includes who initiated the operation, the type of operation, the start time and the end time of the operation, the recovery host, storage, network, any boot order information, etc. The information for each machine includes the steps taken during the operation, such as creating a machine and scratch volumes for testing, when each process began and ended, and whether the operation succeeded or not.

Note: When FOT is in still in progress, the **end time** in the Recovery Report appears as **NA**.

The Recovery operation start time and Recovery operation end time values are shown in UTC according to the Zerto Virtual Manager clock in the recovery site. The Point in time value takes the checkpoint UTC time, which was created in protected site, and converts it to the recovery site time zone.

Branding the Recovery Report

A branded logo can be placed in the report in the top left corner by adding the logo as a .png file to the <ZertoInstallFldr>\Zerto\Zerto Virtual Replication\gui\ folder with the name provider_logo.png.

The folder ZertoInstallFldr is the root folder where Zerto Virtual Replication in the recovery site is installed. For example, C:\Program Files\Zerto.

Resources Report

Information about the resources used by the virtual machines being recovered to a particular site is displayed in the *Resources* report under the *REPORTS* tab. The information is collected at fixed times that are defined in the *Reports* tab of the *Site Settings* dialog in the recovery site. Information for the report is saved for 90 days when the sampling period is hourly and for one year when the sampling period is daily.

The report collects the resource information for the virtual machines being recovered to the site where the report is run.

If no virtual machines are recovered to the site where the report is run, the report is empty.

You can filter the information by the following:

From and **To** – The dates for which you want information.

Click *EXPORT* to generate the report, which is produced as an Excel file.

The information presented in this report is divided into three tabs:

Details Tab – Shows information for each protected virtual machine.

Performance Tab – Shows bandwidth and throughput information for each virtual machine in a table and in a graph.

Target Host Tab – Shows information per host in the recovery site.

Using a REST API to Generate a Report

Zerto Virtual Replication exposes a REST API to produce resource data. The report is generated by passing a URL. For details about the ResourcesReport API (and all other Zerto Virtual Replication REST APIs), see the *Zerto Virtual Replication RESTful API Reference Guide*.

Details Tab

The **Details** tab includes the names and IDs of the virtual machines being protected and, for each virtual machine, the timestamp for the information, where it is protected, the CPU used, the memory used by the host and the guest, the storage used, and other information.

Interpreting the Details Tab

The **Details** tab provides a breakdown of every protected virtual machine, identified by its internal identifier and name in the hypervisor manager. The report also includes the name of the VPG that is protecting the virtual machine and information such as the protected and recovery sites, the protected and recovery vCD Org, cluster, etc.

The `Timestamp` column displays the time when the last sample, as defined in the *Reports* tab of the *Site Settings* dialog, was taken.

The VPG Type column is one of:

VC2VC – vCenter to vCenter replication

VC2VCD – vCenter to vCloud Director replication

VCD2VCD – vCloud Director to vCloud Director replication

VCD2VC – vCloud Director to vCenter replication

The `ZORG` column defines organizations set up in the Zerto Cloud Manager that use a cloud service provider for recovery.

The `Bandwidth (Bps)` and `Throughput (Bps)` columns display the average between two consecutive samples. With daily samples, these figures represent the average daily bandwidth and throughput. For hourly samples, the timestamp represents an average between the sample at the timestamp and the previous sample. A value of -1 means that the system failed to calculate the value, which can happen for several reasons, for example:

- Sites were disconnected when the sample was collected. Although the protected site measures the throughput and bandwidth, the recovery site logs the results.
- The bandwidth or throughput values at the time of the sample was lower than the bandwidth or throughput value in the previous sample. This can happen, for example, if the protected site VRA is rebooted since the sample values are not stored persistently by the VRA.
- If `valueInLastSample` does not exist, since `currentValue` is the first sample for the virtual machine, the data is not calculated.

Bandwidth is calculated as: `(currentValue – valueInLastSample)/elapsedTtime`

For example:

TIME	ACTION/DESCRIPTION
2:29:59.999	A virtual machine is placed in a VPG
2:30	A sample is generated. The total transmitted bytes is zero since the virtual machine was just placed in the VPG
2:30-2:59.999	The VM is writing data at 1MB/minute
3:00	The virtual machine lowers its write rate to 0.5MB/minute

3:30	A new sample is calculated. Current value of total data transmitted is 45MB:
	`1MB/minute)*(30 minutes) + (0.5MB/minute)*(30 minutes)`
	Last value of total data transmitted is 0, from the 2:30 sample.
	`Bandwidth = (45MB-0)/(60 minutes) = 0.75MB/minute = 13107Bps`

Report output fields

The following describes the fields in the **Details** tab.

PARAMETER	DESCRIPTION
Active Guest Memory (MB)	The active memory of the virtual machine.
Bandwidth (Bps)	The average bandwidth used between two consecutive samples, in bytes per second.

PARAMETER	DESCRIPTION
Consumed Host Memory (MB)	The amount of host memory consumed by the virtual machine.
CPU Limit (MHz)	The maximum MHz available for the CPUs in the virtual machine.
CPU Reserved (MHz)	The MHz reserved for use by the CPUs in the virtual machine.
CPU Used (MHz)	The MHz used by the CPUs in the virtual machine.
CrmId	The CRM identifier specified in Zerto Cloud Manager for an organization that uses a cloud service provider for recovery.
Memory (MB)	The virtual machine defined memory.
Memory Limit (MB)	The upper limit for this virtual machine's memory allocation.
Memory Reserved (MB)	The guaranteed memory allocation for this virtual machine.
Number Of vCPUs	The number of CPUs for the virtual machine.
Number Of Volumes	The number of volumes attached to the virtual machine.
Recovery Journal Provisioned Storage (GB)	The amount of provisioned journal storage for the virtual machine. The provisioned journal size reported can fluctuate considerably when new volumes are added or removed.
Recovery Journal Used Storage (GB)	The amount of journal storage used by the virtual machine.
Recovery Volumes Provisioned Storage (GB)	The amount of provisioned storage for the virtual machine in the target site. This value is the sum of volumes' provisioned size.
Recovery Volumes Used Storage (GB)	The amount of storage used by the virtual machine in the target site.
Service Profile	The service profile used by the VPG.
Source Cluster	The source cluster name hosting the virtual machine.
Source Host	The source host name hosting the virtual machine.
Source Organization VDC	The name of the source vDC organization.
Source Resource Pool	The source resource pool name hosting the virtual machine.
Source Site	The source protected site name, defined in the Zerto User Interface.
Source vCD Organization	The name of the source vCD organization.
Source Volumes Provisioned Storage (GB)	The amount of provisioned storage for the virtual machine in the source site. This value is the sum of volumes' provisioned size.
Source Volumes Used Storage (GB)	The amount of storage used by the virtual machine in the source site. This value is the sum of the volumes' used size.
Source VRA Name	The name of the source VRA used to send data to the recovery site.
Target Cluster	The target cluster name hosting the virtual machine.
Target Datastores	The target storage used by the virtual machine if it is recovered.
Target Host	The target host name hosting the virtual machine when it is recovered.
Target Organization vDC	The name of the target vDC organization.
Target Resource Pool	The target resource pool name where the virtual machine will be recovered.
Target Site	The target site name, defined in the Zerto User Interface.
Target Storage Profile	The target vCD storage profile used.
Target vCD Organization	The name of the target vCD organization.
Target VRA Name	The name of the VRA managing the recovery.
Throughput (Bps)	The average throughput of the VM used between two consecutive samples, in bytes per second.

PARAMETER	DESCRIPTION
Timestamp	The date and time the resource information was collected. The value can be converted to an understandable date using code similar to the following: `var date = new Date(jsonDate);` or code similar to the Perl code example, `jsonDateToString($)`, described in *Zerto Virtual Replication RESTful API Reference Guide*.
VM Hardware Version	The VMware hardware version.
VM Id	The internal virtual machine identifier.
VM Name	The name of the virtual machine.
VPG Name	The name of the VPG.
VPG Type	The VPG type: **VCVpg** – VMware vCenter Server **VCvApp** – Deprecated **VCDvApp** – VMware vCloud Director vApp **PublicCloud** – Amazon WebServices or Microsoft Azure **HyperV** – Microsoft SCVMM
ZORG	The name assigned to an organization using a cloud service provider for recovery. The name is created in the Zerto Cloud Manager. For details, see the *Zerto Cloud Manager Administration Guide*.

Performance Tab

The Performance tab shows bandwidth and throughput information for each virtual machine per sampling period in a table and in a graph. The Performance tab enables the user to view the total bandwidth and throughput per sampling period.
The graph allows the user to view performance trends over time per VM.
For full explanation of the bandwidth and throughput information, refer to the "Details Tab", on page 129.
You can filter information by date and VM name.

The following describes the fields in the **Performance** tab:

PARAMETERS	DESCRIPTION
Time Stamp	For explanation see the Details tab.
Bandwidth (Bps)	The average bandwidth of the VM used between two consecutive samples, in bytes per second.
Throughput (Bps)	The average throughput of the VM used between two consecutive samples, in bytes per second.
Total Bandwidth	The total bandwidth of all VMs during the measured period.
Total Throughout	The total throughput of all VMs during the measured period.

Target Host Tab

The Target Host tab shows information per host in the recovery site. This enables the user to perform capacity planning on the recovery host. You can filter information by time and by host.

The following describes the fields in the **Target Host** tab.

PARAMETERS	DESCRIPTION
Active Guest Memory (MB)	The active memory of the virtual machine.
CPU Used (MHz)	The MHz used by the CPUs in the virtual machine.
Host	The Target Host's IP address or DNS name.
Total Bandwidth	The total bandwidth of all VMs replicating to the host during the measured period.

PARAMETERS	DESCRIPTION
Total Throughput	The total throughput of all VMs replication to the host during the measured period.
vCPUs	The number of CPUs for the virtual machine.
VMs	The number of VMs protected.
Volumes	The number of volumes attached to the virtual machine.

VPG Performance

Performance graphs for all VPGs or for an individual VPG can be seen in the *VPG Performance* report under the *REPORTS* tab. These graphs show more detailed resolution than the corresponding graphs in the *DASHBOARD* tab.

You can specify the VPGs whose performance should be displayed. When you request information about multiple VPGs, each VPG is shown in a different color, with a key at the top of the report that maps each color to the VPG it represents.

Position the cursor on a graph line to see exact information about that point.

Click *APPLY* to apply the selected filtering and produce the report.

Click *RESET* to reset the display to the default values.

Backup Report

Information about offsite backups can be sent as a report every day or weekly on a specified day. To set up the report, select *Site Settings* > *Email Settings*.

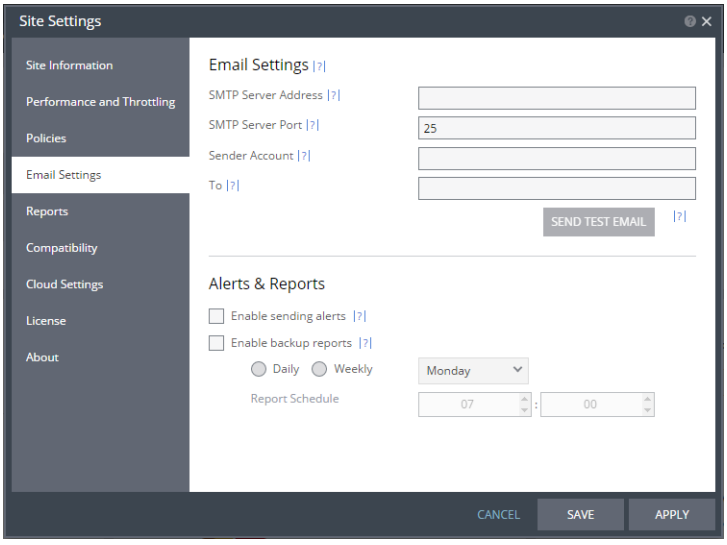

Enter an email address to receive Zerto Virtual Replication backup reports.

SMTP Server Address – The SMTP server address of the hypervisor manager. The Zerto Virtual Manager must be able to reach this address.

SMTP Server Port – The SMTP server port, if it was changed from the default, 25.

Sender Account– A valid email address for the email sender name.

To – A valid email address where that will receive the email containing the backup reports.

SEND TEST EMAIL **button** – Tests that the email notification is set up correctly. A test email is sent to the email address specified in the *To* field.

To configure backup reports:

1. Select *Enable backup reports*.
2. Specify whether you want a backup report sent daily or weekly.

   Daily – Sends a daily backup report.

   Weekly – Sends a weekly backup report. Select the day of the week from the dropdown list.
3. Specify the time of day to send the backup report.

The backup report is sent as HTML with the following information:

- A summary listing every VPG for which an offsite backup job has run. The summary information includes the following:
  - An entry for each backup job that was run.
  - The result of the job: successful, partial successful, or failed.
    A partially successful job means that some, but not all, of the virtual machines were successfully backed up.
  - The time the job started.
  - The time the job completed.
  - The duration of the job.
  - The size of the backup that was stored in the repository.
  - The type of the job: automatic, meaning a scheduled run, or manually initiated.
  - Summary details of the run.
- Specific details about the job, including:
  - The name of the ZORG of the VPG.
  - The protected site.
  - The backup site where offsite backup can be restored.
  - The number of virtual machines backed up from the totally number in the VPG.
  - The number of virtual machines only partially backed up.
  - The start and end times of the run and the run duration.
  - The backup size.
  - Whether the backup was scheduled or initiated manually.
  - The repository name.
  - The next time a backup of the VPG is scheduled.
  - The previous run time.

CHAPTER 18: TROUBLESHOOTING

You can handle problems related to the WAN connecting the protecting or recovery sites, or other problems using a variety of diagnostic and troubleshooting tools.

The following topics are described in this chapter:

For details about Zerto Virtual Manager alarms, alerts, and events, refer to *Zerto Virtual Replication Guide to Alerts and Events*.

Ensuring the Zerto Virtual Manager is Running

If you have problems accessing the Zerto User Interface, check under Windows Services, on the machine where Zerto Virtual Replication is installed, that the *Zerto Virtual Manager* Windows service is started.

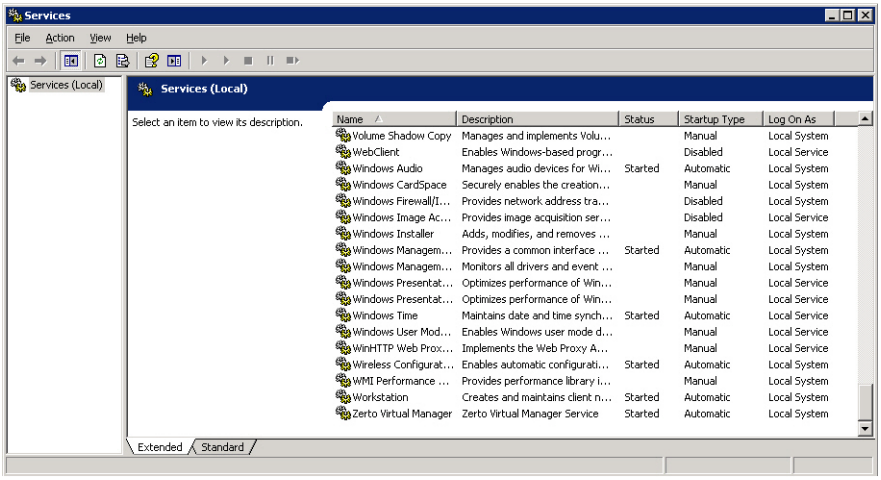

Troubleshooting: "Needs Configuration" Problems

When the VPG status changes to `Needs Configuration`, the settings in the VPG need to be checked and, when necessary, updated.

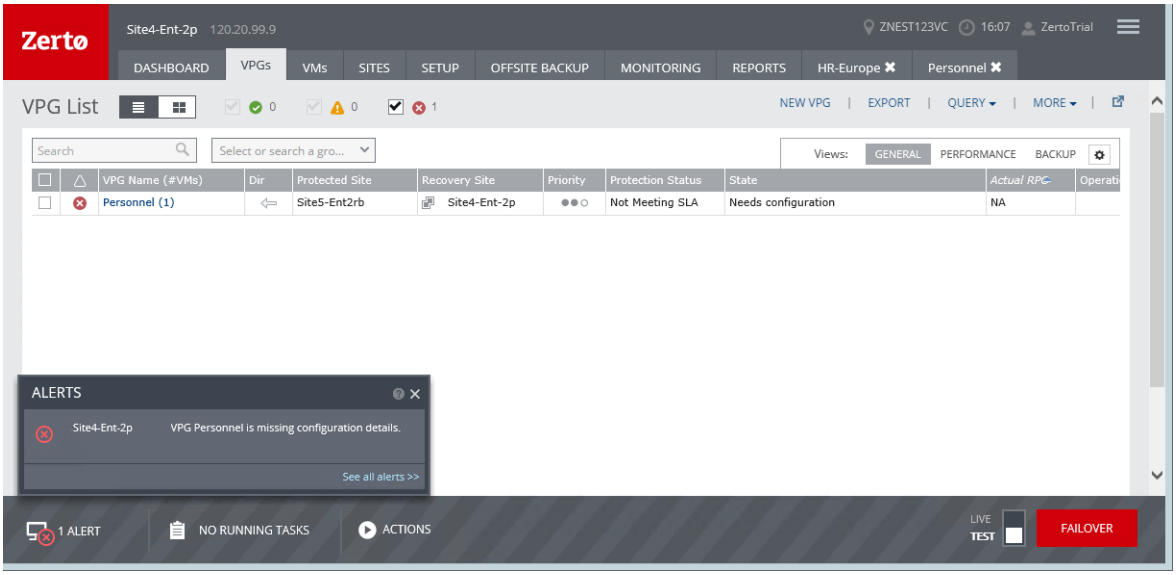

The following scenarios result in the VPG status changing to *Needs Configuration*:

- A protected disk resize operation fails, for example when there is not enough disk space.
- When a volume is added to a protected virtual machine and the VPG settings are not updated because of a site disconnection or a AWS error. In some situations, after the sites reconnect, the state corrects itself automatically.

Troubleshooting VRA Problems

VPG Syncing Takes a Long Time – Network Problems

Check the network. If the firewall configuration is modified, the VRA TCP connections have to be reset. After a VRA disconnect and reconnect the system can wait for up to fifteen minutes before syncing the sites after the reconnection.

Zerto Virtual Replication Diagnostics Utility

Zerto Virtual Replication includes a diagnostics utility to help resolve actual and potential problems. Using the diagnostics tool, you can do the following:

- Collect logs to help Zerto support resolve problems. The Zerto Virtual Manager must be running on each site for which you want logs. See "To collect logs for Zerto support to use when troubleshooting:", below.
- Collect local Zerto Virtual Manager logs. Use this option if the Zerto Virtual Manager is not running. See "To collect local Zerto Virtual Manager logs when the Zerto Virtual Manager is not running:", on page 139.
- Check the connectivity between Zerto Virtual Replication components. See "Reconfiguring the Zerto Virtual Manager Setup", on page 62.
- Export VPG settings to an external file and import these settings.
- Reconfigure access to the Microsoft SQL Server that is used by the Zerto Virtual Manager. This database was specified during the installation of Zerto Virtual Replication. See "Reconfiguring the Microsoft SQL Server Database Used by the Zerto Virtual Manager", on page 70.

Note: A separate installation kit is available for download from the Zerto Support Portal downloads page that installs the Zerto Virtual Replication Diagnostics utility as a standalone utility on any Windows machine that has Microsoft .NET Framework 4 installed[1].

Collecting Zerto Virtual Replication Logs

Virtual replication logs can be collected to help Zerto support resolve problems related to Zerto Virtual Replication. Virtual replication logs can be collected in the following ways:

"Using Remote Log Collection", below

"Using the Zerto Diagnostics application", on page 137

Using Remote Log Collection

Remote Log Collection allows customers to authorize Zerto support engineers to collect logs from their environment. By using remote log collection customers can avoid having to use the Diagnostic Tool on their ZVM server in order to collect logs for analysis, a potentially complex and time-consuming procedure.

To enable Remote Log Collection:

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

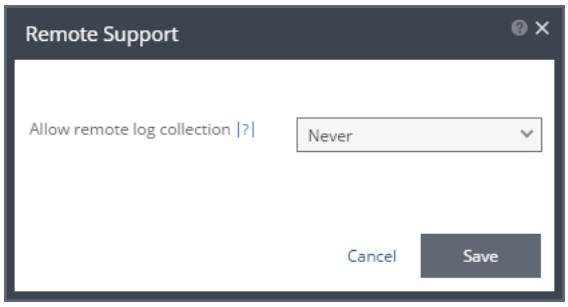

2. Click the drop down menu to display the remote log collection options.

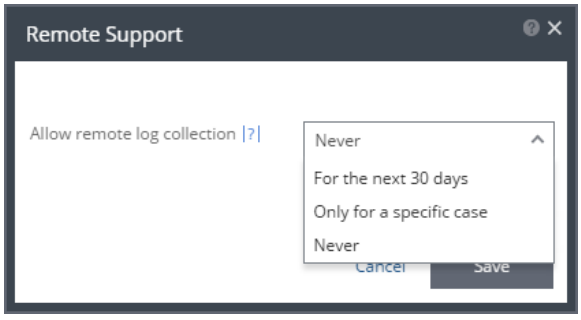

3. Select the remote log collection option you wish to allow:
   Never – Remote log collection is not allowed (default). If remote log collection is currently is allowed, it will be terminated if you select this option.
   For the next 30 days -Remote log collection is allowed. This permission will automatically terminate in 30 days unless terminated by selecting the *Never* option.

1. The installation executable is included as part of the standalone utility installation kit and it requires an additional 1.8GB of free disk space.

Only for a specific case - You will be prompted to enter the *Case number* opened via the *Salesforce Self-service Portal*. Remote log collection will be allowed for as long as the case is active or until remote log collection is terminated by selecting the *Never* option.

4. Click *Save*.

Using the Zerto Diagnostics application

You can collect logs using the diagnostics tool to help Zerto support resolve problems when the Zerto Virtual Manager is running or when the Zerto Virtual Manager is not running.

- When the Zerto Virtual Manager is running, see "To collect logs for Zerto support to use when troubleshooting:", below. This option enables you to specify the logs that you want to collect, generated by Zerto Virtual Replication, for example VRA logs, as well as logs generated by VMware, for example, vCenter Server logs or host logs. The Zerto Virtual Replication generated logs can be filtered by any alerts issued and by the VPGs that require analysis to identify problems.
- When the Zerto Virtual Manager is not running, see "To collect local Zerto Virtual Manager logs when the Zerto Virtual Manager is not running:", on page 139.

To collect logs for Zerto support to use when troubleshooting:

1. Open the *Zerto Diagnostics* application. For example, via *Start > Programs > Zerto Virtual Replication > Zerto Diagnostics*.

   The *Zerto Virtual Replication Diagnostics* menu dialog is displayed.

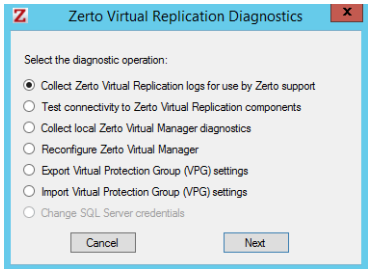

2. Select the *Collect the Zerto Virtual Replication logs for use by Zerto support* option.
3. Click *Next*.

   The *Initialize* dialog is displayed.

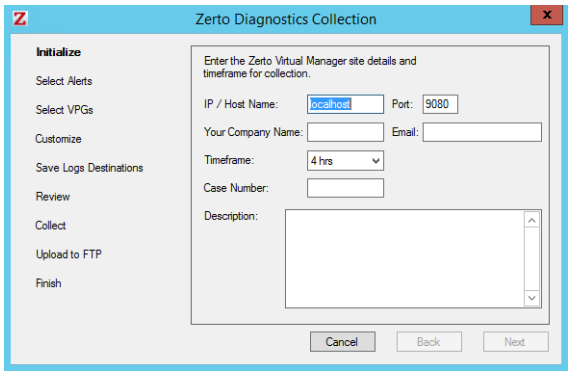

4. Specify the following and click *Next*.

   IP / Host Name – The IP of the Zerto Virtual Manager where the log collection runs from. Logs are collected from this site and from the paired site.

   Port – The port used for inbound communication with the Zerto Virtual Manager.

   Your Company Name – A name to identify the log collection for the customer. This information is used by Zerto support. An account name must be entered. After this information is added, it is displayed in subsequent uses of the diagnostics utility.

   Email – An email address for use by Zerto support when analyzing the logs. An email address must be entered. After this information is added, it is displayed in subsequent uses of the diagnostics utility.

Timeframe – The amount of time you want to collect logs for. The more time, the bigger the collection package.

Case Number – The case number assigned by Zerto support, if one already exists. Optional.

Description – An optional free text description of the reason for collecting the logs.

After clicking *Next* the utility connects to the Zerto Virtual Replication and if any alerts have been issued, they are displayed in the `Select Alerts` dialog.

If there are no alerts, this dialog is skipped.

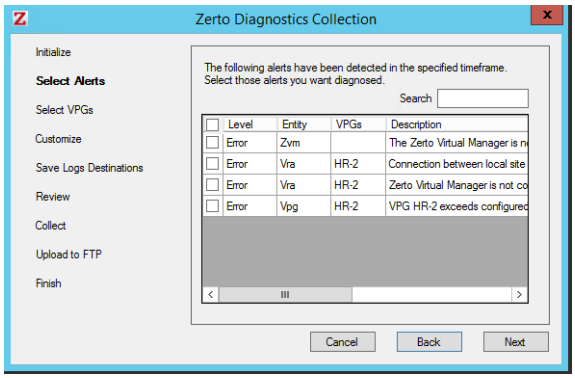

5. Select any alerts that need analyzing from the list and click *Next*.

   The *Select VPGs* dialog is displayed.

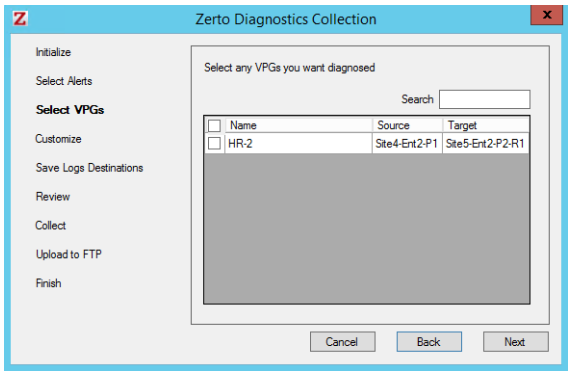

6. Select the VPGs that you want analyzed and click *Next*.

   The *Customize* dialogs are displayed. With AWS these values are not relevant.

7. Click *Next* until the *Save Log Destinations* dialog is displayed.

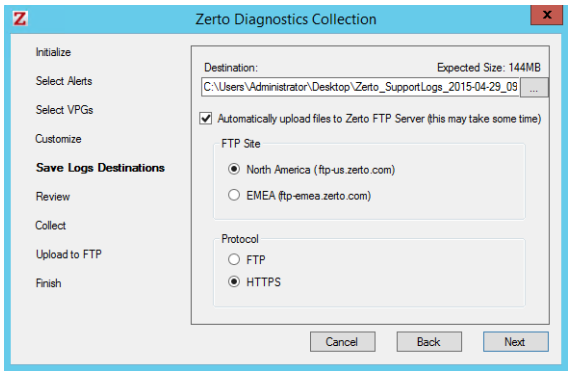

8. Specify destination for the files that you want collected.

   Destination – The name and location where the log collection will be saved.

   Automatically upload files to Zerto FTP Server – When this option is checked, the log collection is automatically uploaded to a specified FTP site.

   If you choose to upload the log collection to a site that you specify, make sure that the site is up.

9. Specify the FTP site to send the collection and the protocol to use, either FTP or HTTP.
10. Click *Next*.

    The *Review* dialog is displayed.

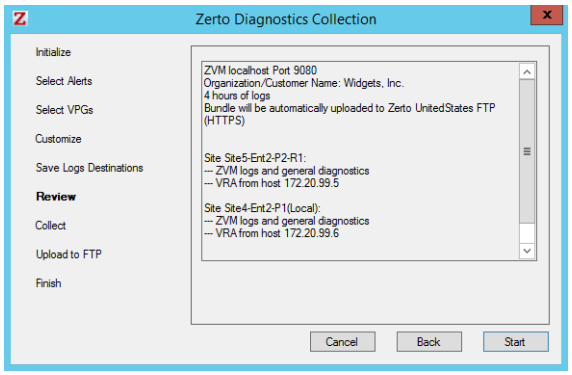

    Check that you have specified everything you want to collect and if you want to make changes, click *Back* to change the selection.
11. Click *Start*.

    The data is collected and stored in the destination file which, by default, is timestamped. If specified, the collection is also sent to an FTP site.

    Note: The log collection is performed on the server. Canceling the collection in the GUI does not stop the collection from continuing on the server and a new log collection cannot be run until the running collection finishes.

    When the log collection has completed the result is displayed. For example:

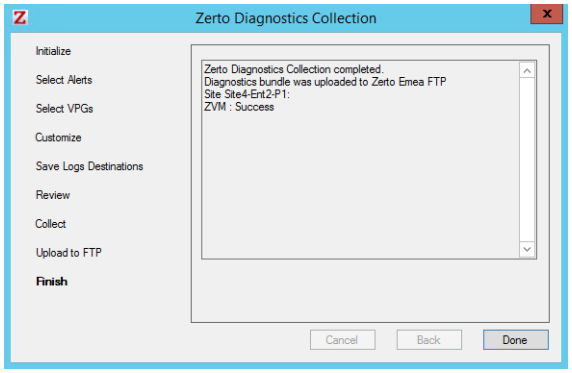

12. Click *Done* to return to the *Zerto Virtual Replication Diagnostics* menu dialog.
13. Send the log to Zerto support, unless the `Automatically upload files to Zerto FTP Server` option was specified, in which case it is automatically sent to Zerto.

To collect local Zerto Virtual Manager logs when the Zerto Virtual Manager is not running:

1. Open the *Zerto Diagnostics* application. For example, via *Start* > *Programs* > *Zerto Virtual Replication* > *Zerto Diagnostics*.

   The *Zerto Virtual Replication Diagnostics* menu dialog is displayed.
2. Select the *Local Zerto Virtual Manager diagnostics* option and click *Next*.

You are prompted to use the first option to collect more comprehensive diagnostics. If you continue, the *Initialize* dialog is displayed.

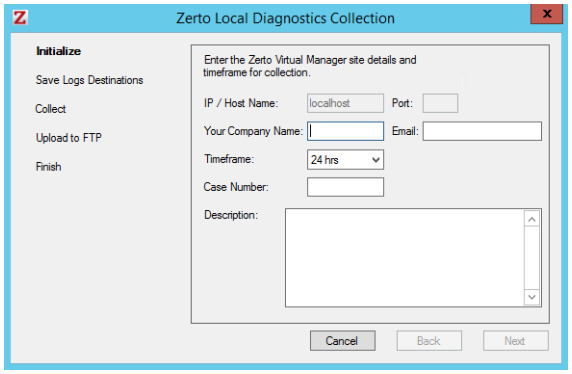

3. Specify the details that you want collected.

   IP / Host Name – The IP of the Zerto Virtual Manager where the log collection runs from. Logs are collected from this site and from the paired site.

   Port – The port used for inbound communication with the Zerto Virtual Manager.

   Your Company Name – A name to identify the log collection for the customer account. This information is used by Zerto support. An account name must be entered.

   Email – An email address for use by Zerto support when analyzing the logs. An email address must be entered.

   Timeframe – The amount of time you want to collect logs for. The more time, the bigger the collection package.

   Case Number – An optional field for the case number assigned to the issue by Zerto.

   Description – An optional free text description of the reason for collecting the logs.

4. Click *Next*.

   The *Save Log Destinations* dialog is displayed.

5. Specify the details that you want collected.

   Destination – The name and location where the log collection will be saved.

   Automatically upload files to Zerto FTP Server – When this option is checked, the log collection is automatically uploaded to a specified FTP site.

   If you choose to upload the log collection to a site that you specify, make sure that the site is up before clicking *Finish*.

   The data is collected and stored in the destination file which, by default, is timestamped. If specified, the collection is also sent to an FTP site.

6. Click *Next*.

   The collection progress is displayed. When the log collection has completed the result is displayed.

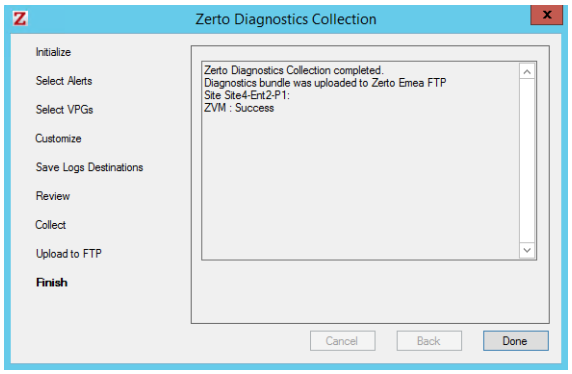

7. Click *Done* to return to the `Zerto Virtual Replication Diagnostics` menu dialog.
8. Send the log to Zerto support, unless the `Automatically upload files to Zerto FTP Server` option was specified, in which case it is automatically sent to Zerto.

Understanding the Logs

If problems arise with Zerto Virtual Manager, you can view the Zerto Virtual Manager logs to see what is happening.

The current log is called `logfile.csv` and resides in the <*Zerto_Install_Dir*>`\Zerto Virtual Replication\logs` folder, where *Zerto_Install_Dir* is the folder specified during the installation.

When the log reaches 10MB its name is changed to `log.`*nnnn*`.csv`, where *nnnn* is a number incremented by one each time `logfile.csv` reaches 10MB. Up to 150 log files are kept.

The log file has the following format:
`FFFF, yyyy-mm-dd hh:mm:ss, ####, LVL, Component, API, Message`

where:

FFFF – A HEX code. For internal use.

yyyy-mm-dd hh:mm:ss – Timestamp for the message.

– Number for internal use.

LVL – Severity level of the message. The more messages written to the log the bigger the impact on performance. The number of messages written to the log decreases from `Debug` to `Error`. The level can be one of the following:

Debug – All messages are written to the log. This level should only be specified during testing.

Info – Information messages.

Warn – Warning messages such as a reconnect ion occurred.

Error – Error messages that need handling to find the problem.

Component – The specific part in the Zerto Virtual Manager that issued the message.

API – The specific API that issued the message.

Message – The message written in the log.

The following is a sample from a log:

```
07f4c878,2010-12-01 19:54:41.4237,Debug,5,
Zerto.Zvm.RemoteZvmConnector.ResyncingRemoteZvmConnector,
TestConnectivity,TestConnectivity returning true,
07f4c878,2010-12-01 19:54:41.7362,Info,11,
Zerto.Zvm.ZvmServices.Protection.PromotionMonitor,
PromotionMonitoringThreadFunc,Promoting protection groups: ,
07f4c878,2010-12-01 19:54:42.7987,Info,9,
Zerto.Infra.ZvmReaderWriterLock,LogLock,Synchronizer: Enter Writer,
07f4c878,2010-12-01 19:54:42.7987,Info,9,
Zerto.Zvm.ZvmServices.ReconnectingConnectorProxy,
GetConnector,"Connecting IP=106.16.223.86, PORT=4005, attempt (1/3)",
07f4c878,2010-12-01 19:54:42.7987,Debug,9,
Zerto.Zvm.VraConnector.VraNetworkConnector,
Connect,try to connect 106.16.223.86:4005 ...,
07f4c878,2010-12-01 19:54:43.0643,Debug,17,
Zerto.Zvm.ZvmServices.CrossSiteService,Ping,Ping,
07f4c878,2010-12-01 19:54:43.0643,Debug,17,
Zerto.Zvm.ZvmServices.PingService,Ping,Ping called,
07f4c878,2010-12-01 19:54:43.8612,Error,9,
Zerto.Zvm.VraConnector.VraNetworkConnector,
ClearAndThrow,connection is closed: No connection could be made because the target
machine actively refused it 106.16.223.86:4005,
07f4c878,2010-12-01 19:54:43.8612,Warn,9,
Zerto.Zvm.ZvmServices.ReconnectingConnectorProxy,GetConnector,failed,
```

CHAPTER 19: THE ZERTO VIRTUAL MANAGER USER INTERFACE

Configuration and management of disaster recovery for a site is performed in the Zerto User Interface.

The following dialogs and tabs are described in this chapter:

Add Checkpoint Dialog

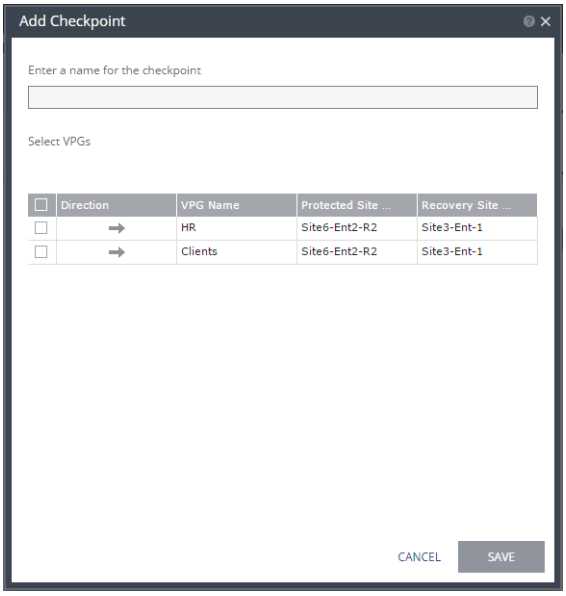

Checkpoints are recorded automatically every few seconds in the journal. These checkpoints ensure crash-consistency and are written to the virtual machine journals by the Zerto Virtual Manager. Each checkpoint has a timestamp set by the Zerto Virtual Manager. In addition to the automatically generated checkpoints, you can manually add checkpoints to identify events that might influence the recovery, such as a planned switch over to a secondary generator.

The list of VPGs is displayed. You can select more VPGs to add the same checkpoint.

Enter a name for the checkpoint – The name to assign to the checkpoint.

Dir – The direction of the protection.

VPG Name – The name of the VPG.

Protected Site Name – The name of the site where virtual machines are protected.

Recovery Site Name – The name of the site where protected virtual machines are recovered.

You can filter columns in the list via the filter icon next to each column title. You can also sort the list by each column. Clicking the cog on the right side of the table enables you to change the columns that are displayed and to create a permanent view of the columns you want displayed.

Add Site Dialog

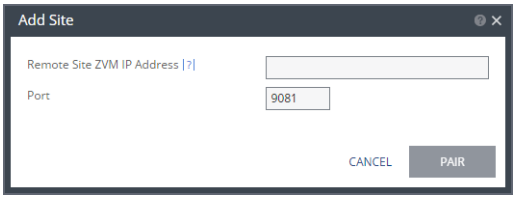

Pair sites

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.

Advanced VM Settings for Cloud Dialog

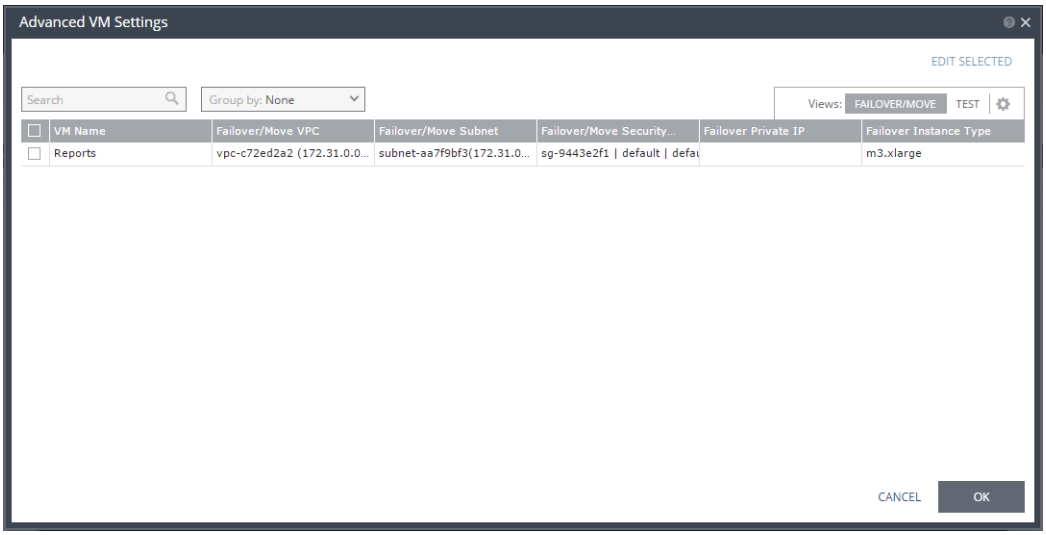

Displays the recovery settings for the virtual machines in the VPG. By default, the recovery settings for failover and move are displayed. You can display the recovery settings for failover tests by selecting *TEST*. You can choose to edit information in one field by clicking the field and updating the information. You can choose to edit information for several virtual machines at the same time by selecting the virtual machines and clicking *EDIT SELECTED*.

ALERTS

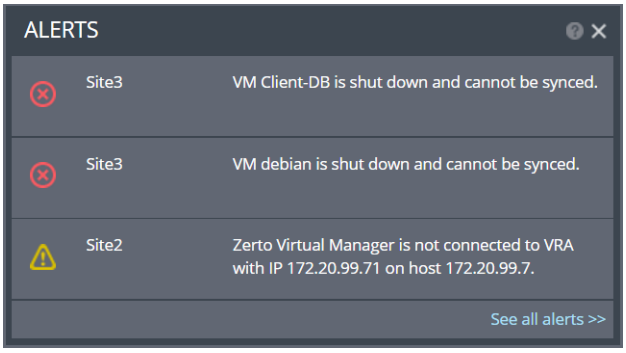

Monitor the recent alerts by clicking the ALERTS area in the status bar at the bottom of the Zerto User Interface. The following information is displayed for the most recent alerts:

- The alert status.
- The site where the alert is issued.
- A description of the alert.

Click *See All Alerts* to access *MONITORING > ALERTS*.

Boot Order Dialog

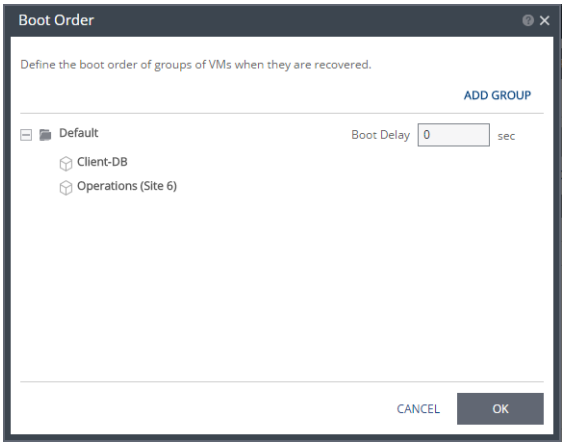

To specify the boot order of virtual machines in a VPG. When machines are started up on recovery, for example, after a move operation, the virtual machines in the VPG are not started up in a particular order. If you want specific virtual machines to start before other machines, you can specify a boot 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.

There is no boot order for virtual machines within a group, only between groups.

ADD GROUP button – Adds a group. After adding a group you can edit the group name by clicking the Edit icon at the right of the group name and remove the group via the delete icon at the right of the group. You cannot remove the `Default` group nor a group that contains a virtual machine.

Boot Delay – Specifies 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 Start-up 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 up.

Checkpoints Dialog

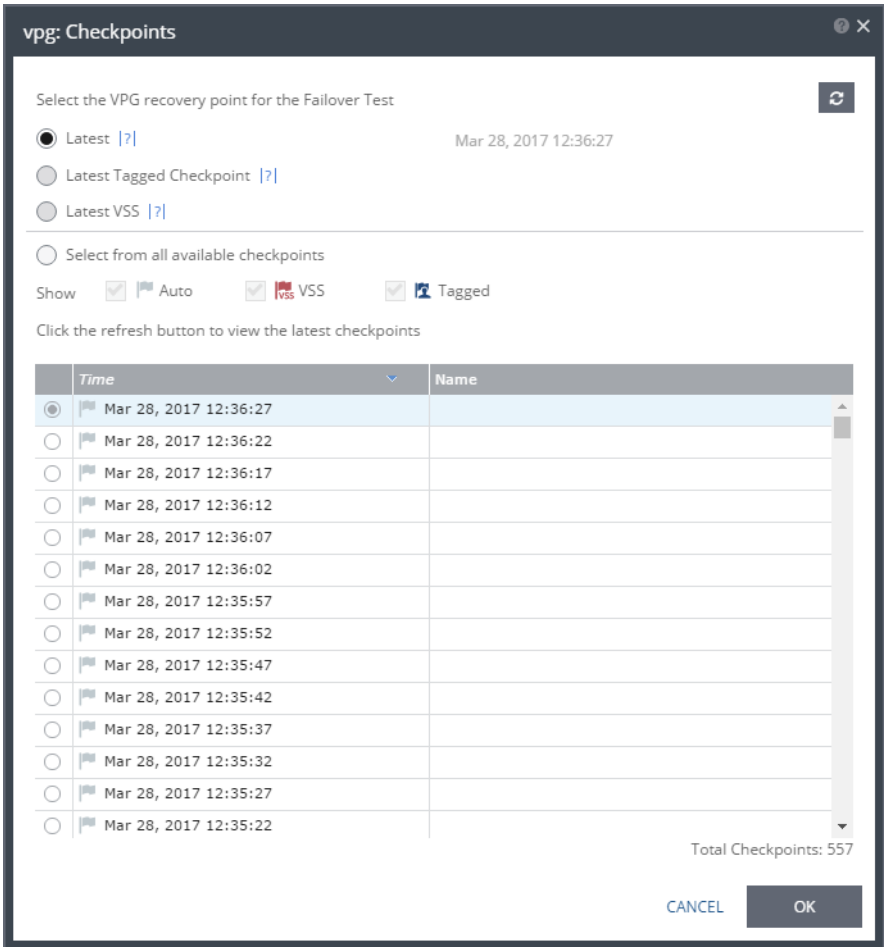

When selecting the point to recover to:

- The refresh button is initially grayed out and is enabled for clicking after 5 seconds. It is also grayed out for 5 seconds after being clicked, before being re-enabled.
- A ***Click the refresh button to view the latest checkpoints*** reminder is displayed 10 seconds after the refresh button is clicked to remind the user that there is a new *Latest Checkpoint*.
- If the user has scrolled to, and selected, a checkpoint anywhere in the checkpoints list, clicking the refresh button will automatically return the user to the selected checkpoint in the list.

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

Latest Tagged Checkpoint – The recovery operation is to the latest checkpoint added in one of the following situations:

- By a user.
- When a failover test was previously performed on the VPG which includes the virtual machine.
- When the virtual machine was added to an existing VPG after the added virtual machine was synchronized.

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

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

Edit VM Network Dialog

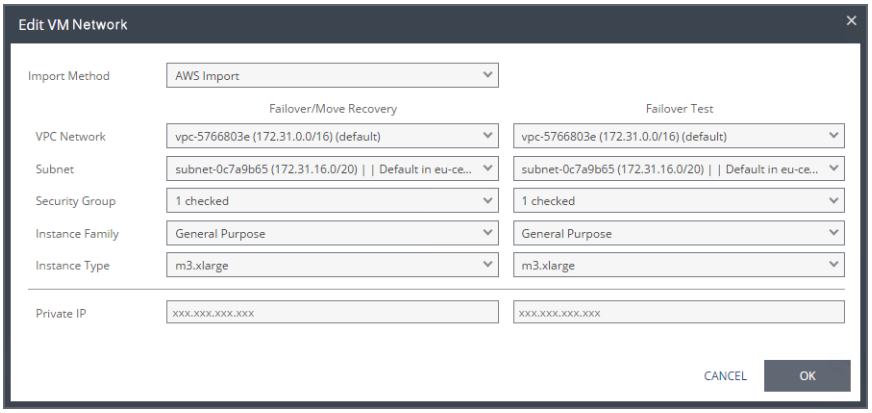

Edit the network settings for one or more virtual machines in a VPG that will be recovered to AWS. There are recovery settings for failovers and moves, and for failover tests.

Import Method

- **AWS Import** - The method that has been used in all past implementations.
- **Zerto Import for all volumes** – This is the fastest import method and uses a ZImport VM for all volumes. A ZImport virtual machine is created, per volume. For each recovered volume, the ZImport virtual machine is terminated when all the data has been imported and its disk has been attached to the recovered virtual machine.
- **Zerto Import for data volumes** – Uses the ZImport method for data volumes only. This is the default setting and has a faster RTO than AWS Import.

VPC Network – A virtual network dedicated to your AWS account.

Subnet – A range of IP addresses in your VPC.

Security Group – The AWS security to be associated with the virtual machines in this VPG. You can associate one or more security groups with the virtual machines.

Instance Family – The instance family from which to select the type. AWS instance families are optimized for different types of applications. Choose the instance family appropriate for the application in the virtual machine protected in the VPG.

Instance Type – The instance type, within the instance family, to assign to recovered instances. Different types within an instance family vary primarily in vCPU, ECU, RAM, and local storage size. The price per instance is directly related to the instance size.

Private IP – The private IP of an instance from the selected subnet. If you do not set the private IP, during recovery, AWS sets the private IP from the defined subnet range.

New Repository Dialog

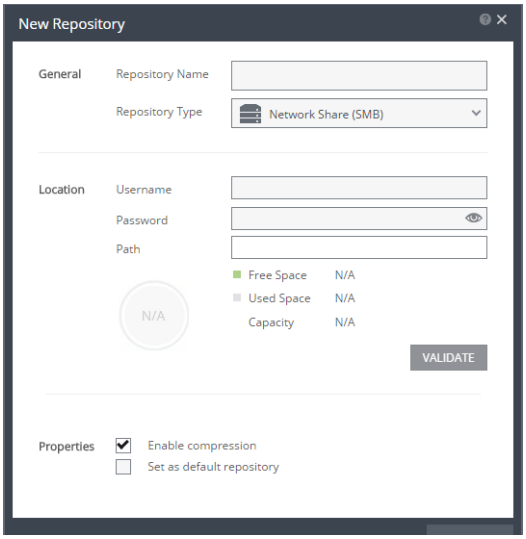

To create a new repository for backups.

Repository Name – The name of the repository.

Repository Type – The type of repository. The options are *Local* or *Network Share (SMB)*. If *Local* is specified, backups are stored on the local machine where the Zerto Virtual Manager is installed. If *Network Share (SMB)* is specified, the network share drive must be an SMB drive and if specified the username and password to access the drive must be provided.

Username – Username to access the Network Share drive. The name can be entered using either of the following formats:

- `username`
- `domain\username`

This field is not displayed when the type is `Local`.

Password – Password to access the Network Share drive. This field is not displayed when the type is `Local`.

Path – The path where the repository will reside. The path must be accessible from the Zerto Virtual Manager, so if the repository is on a different domain to the Zerto Virtual Manager, the domain must be included in the path.

Free Space – The amount of free space currently available for the repository.

Used Space – The amount of space currently used in the repository.

Capacity – The overall capacity of the repository.

VALIDATE – Click to validate the path. The path must be valid in order to save the information.

Enable compression – Check this option to compress backups stored in the repository. Compression is done using zip compression, set to level six. If you want better compression, which requires more CPU, or less compression to reduce the CPU overhead, contact Zerto support.

Note: Compression usually reduces the effectiveness of deduplication on stored data. If the backup repository resides on a deduplication-enabled storage appliance, it is recommended that the data be stored uncompressed.

Set as default repository – Check this option to make the repository the default repository.

Offsite Clone Dialog

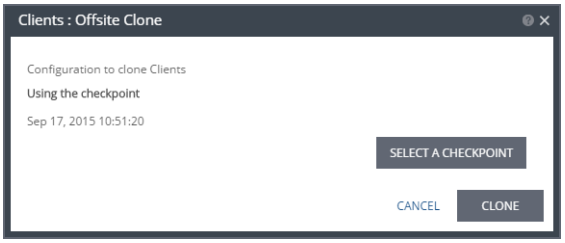

Create a clone of each virtual machine in a VPG on the recovery site in the production network. The clone is a copy of the protected virtual machines on the recovery site, while the virtual machines on the protected site remain protected and live.

SELECT A CHECKPOINT button – Opens the Checkpoints Dialog dialog to select the checkpoint to use to make the clone.

Open Support Ticket Dialog

Support tickets can be opened directly in the Zerto User Interface.

Creating a support ticket in the Zerto User Interface simplifies the submission process since much of the information that is required when entering a ticket using the Zerto Support Portal, such as the version and build numbers, is automatically added to the ticket when it is submitted via the Zerto User Interface.

In addition, when the ticket is submitted, a snapshot of the current environment is also attached to the ticket. The snapshot information includes the lists of alerts, events, tasks, VPGs, and virtual machines that are protected.

This information is used to help Zerto resolve the ticket quickly and, whenever possible, without the need to request more information from you.

Note: 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. If the clocks are not synchronized with UTC, submitting a support ticket can fail.

To open a support ticket:

1. In the Zerto User Interface, click **SETTING** ( ) in the top right of the header and select **Submit Support Ticket**.

The Open Support Ticket dialog for the site is displayed.

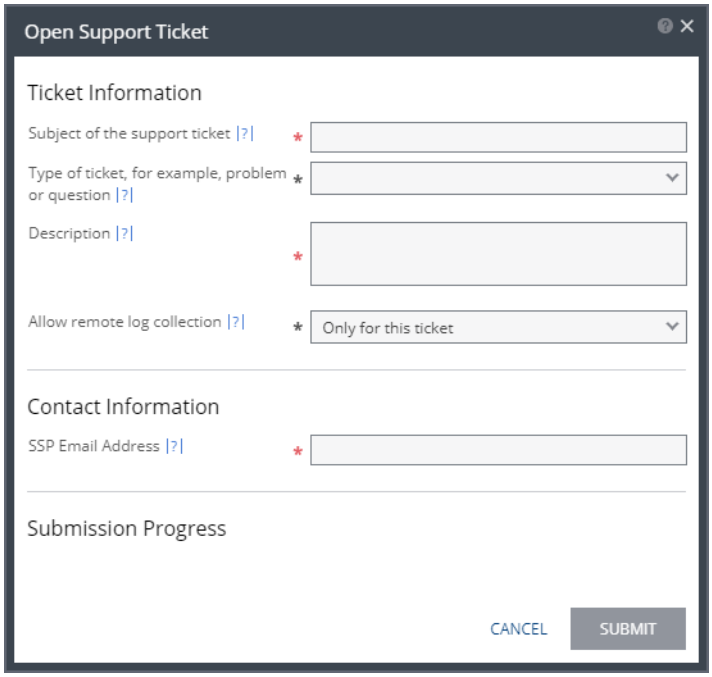

2. Specify the ticket details:
   - **Subject:** The subject of the support ticket.
   - **Type:** The type of ticket being opened. Available options are:
     - Problem
     - Feature Request
     - Question
   - **Description:** A description of the problem or question in addition to the information supplied in the subject.
   - **Allow remote log collection:** How many logs is Zerto allowed to collect. Available options are:
     - Only for this ticket
     - For the next 30 days
     - Never
   - **SSP Email Address:** A valid email address registered with Zerto, with permission to open tickets.
3. Click **SUBMIT**.

The ticket is processed and its progress is displayed. If the email address is not valid, the ticket is rejected. Once the ticket submission starts, it cannot be canceled.

Remote Support Dialog

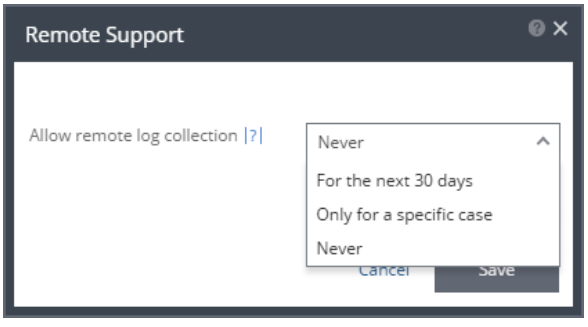

Remote Log Collection allows customers to authorize Zerto support engineers to collect logs from their environment. By using remote log collection customers can avoid having to use the Diagnostic Tool on their ZVM server in order to collect logs for analysis, a potentially complex and time-consuming procedure.

Never – Remote log collection is not allowed (default). If remote log collection is currently is allowed, it will be terminated if you select this option.

For the next 30 days -Remote log collection is allowed. This permission will automatically terminate in 30 days unless terminated by selecting the *Never* option.

Only for a specific case - You will be prompted to enter the *Case number* opened via the *Salesforce Self-service Portal*. Remote log collection will be allowed for as long as the case is active or until remote log collection is terminated by selecting the *Never* option.

Site Settings Dialog

Contains site-wide settings:

Site Information Dialog

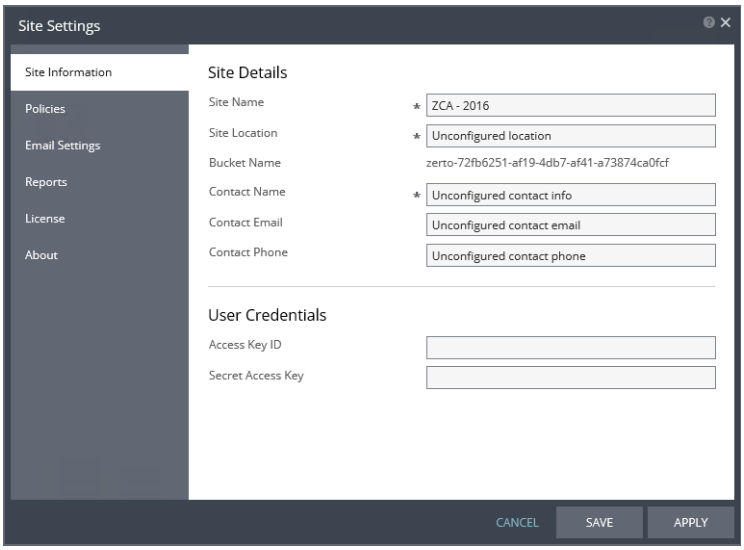

During installation, information about the site is entered to identify the site in the user interface and to identify the contact person at the site. After installation you can update this information.

Site Name – The name used to identify the site.

Site Location – Information such as the address of the site or a significant name to identify it.

Bucket Name – The name of the bucket that was created when Zerto Virtual Replication was installed. This cannot be changed.

Contact Name – The name of the person to contact if a need arises.

Contact Email – An email address to use if a need arises.

Contact Phone – A phone number to use if a need arises.

Access Key ID – A unique identifier that is associated with a secret access key.

Secret Access Key – A key that is used with the access key ID.

Policies Dialog

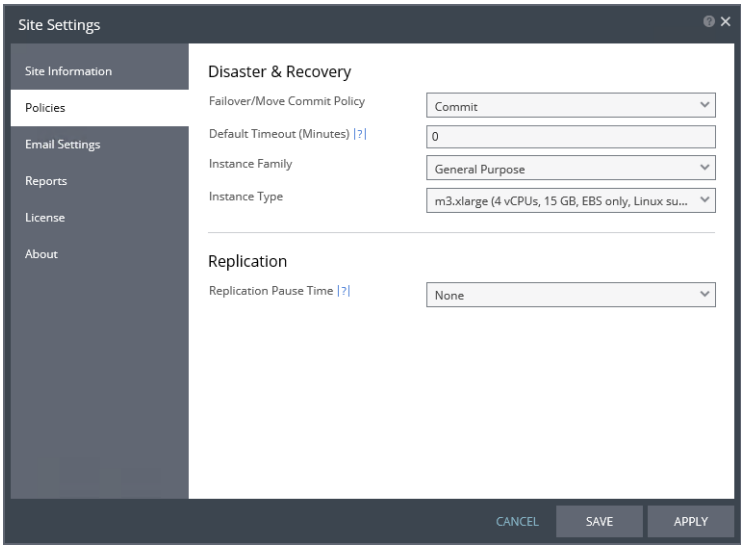

Failover/Move Commit Policy – The commit policy to use during a failover or move operation. The value set here is the default for all failover or move operations from this point on but can be changed when defining a failover or move operation. The following options are available:

- **None** – The failover or move operation must be manually committed or rolled back by the user.
- **Commit** – After the time specified in the `Default Timeout` field, the failover or move operation is committed. During the specified time you can check the recovered VPG virtual machines, and you can manually commit or roll back.
- **Rollback** – After the time specified in the `Default Timeout` field the failover or move operation is rolled back, unless you manually commit it or roll it back before the time out value is reached. During the specified time you can check the recovered VPG virtual machines.

Default Timeout – The time-out in minutes after which a `Commit` or `Rollback` is performed. A value of zero indicates that the system automatically performs the commit policy, without waiting for any user interaction.

Instance Family – The instance family from which to select the type. AWS instance families are optimized for different types of applications.

Instance Type – Within the instance family, the types of instances that can be chosen for recovered instances. Different types within an instance family vary primarily in vCPU, ECU, RAM, and local storage size. The price per instance is directly related to the instance size.

Replication Pause Time – The time to pause when synchronizing a VPG if continuing the synchronization will cause all the checkpoints in the journal to be removed.

Email Settings Dialog

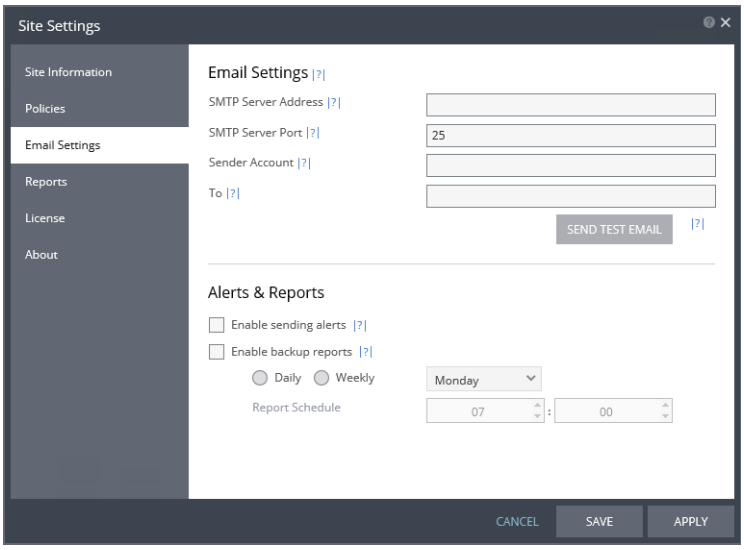

Define an email address to receive Zerto Virtual Replication alerts and backup reports.

SMTP Server Address – The SMTP server address. The Zerto Virtual Manager must be able to reach this address.

SMTP Server Port – The SMTP server port, if it was changed from the default, 25.

Sender Account– A valid email address for the email sender name.

To – A valid email address where you want to send the email.

SEND TEST EMAIL **button** – Tests that the email notification is set up correctly. A test email is sent to the email address specified in the `To` field.

Enable sending alerts – Check to be notified by email about any Zerto Virtual Replication alerts issued. An email is sent when the alert is issued, and after it has been successfully handled and the alert is no longer valid.

Enable backup reports – Defines when backup reports will be emailed.

Reports Dialog

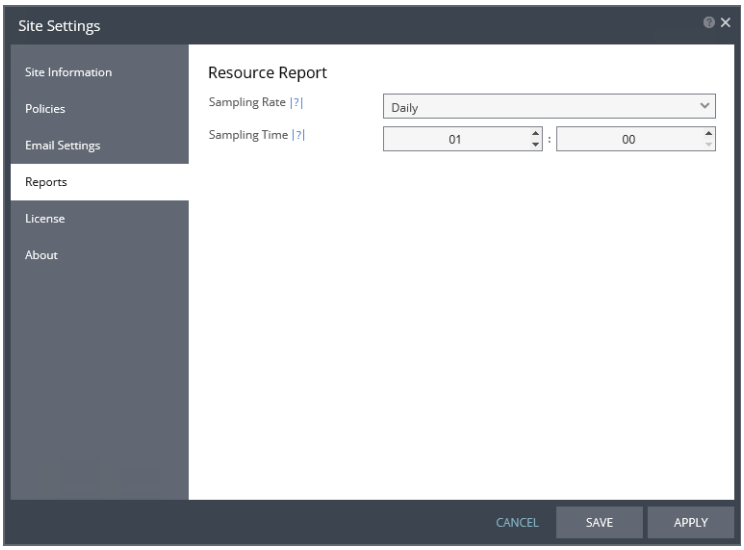

Configures the Resource Report.

Sampling Rate – When to take resource samples to identify resource usage, either daily at a specific hour and minute or hourly at a specific minute within each hour. Note that collecting a sample hourly provides a higher resolution picture of replication traffic than if samples are only collected once a day.

Sampling Time – The time that the sample is taken.

License Dialog

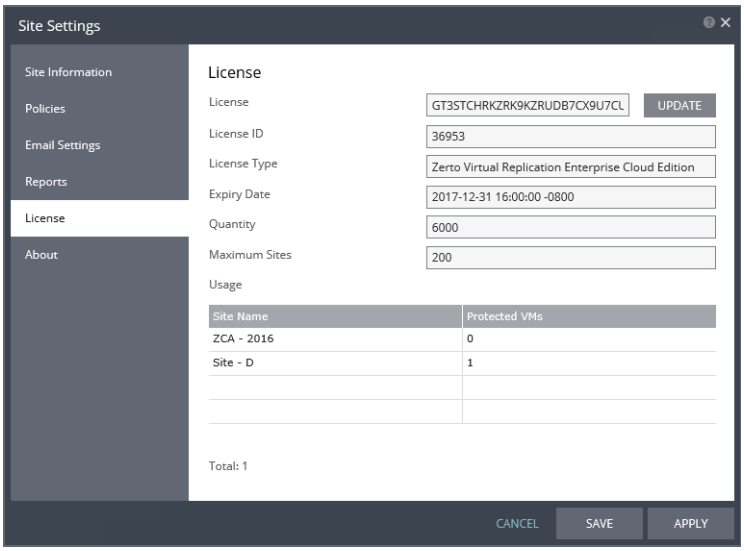

The Zerto license includes the following details:

License – The license key itself.

License ID – An identifier for the license.

License Type – What is licensed: whether the license restricts the number of virtual machines that can be protected or the number of sockets used.

Expiry Date – The license expiry date.

Quantity – The maximum amount licensed, either virtual machines or sockets, based on the license type. If blank, the quantity is unlimited.

Maximum Sites – The maximum number of sites allowed.

Usage – The sites using the license and number of protected virtual machines in each site. In VMware, the number of virtual machines is independent of whether they are in vApps.

About Dialog

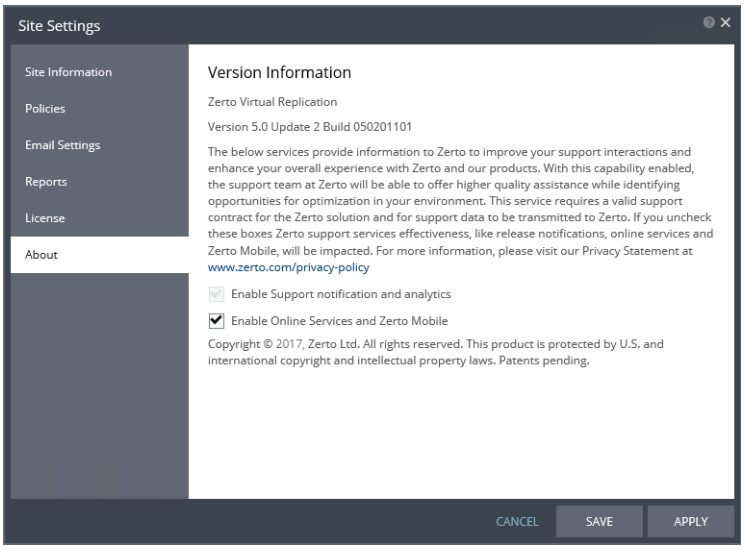

The *About* dialog includes the Zerto Virtual Replication version number and whether to allow analytics to be sent to the ZertoCall Home server to help improve Zerto Virtual Replication. You can also enable or disable the Zerto Virtual Manager to send data to the SaaS platform for monitoring purposes.

Send analytics to Zerto – When selected, analytics are sent to Zerto that are used to improve Zerto Virtual Replication and to automatically update Zerto Virtual Replication when a new version of a hypervisor is released that is supported by Zerto.

Send Data to Cloud – Allows licensed Zerto Virtual Manager users to enable or disable data being sent from the Zerto Virtual Manager to the SaaS platform thereby enabling site monitoring using the Zerto Mobile App.

Note: The Enable Online Services and Zerto Mobile option is enabled by default.

Stop Failover Test Dialog

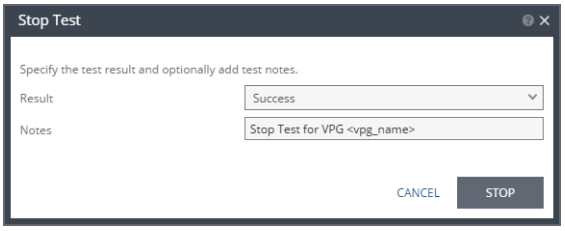

Enables stopping the testing of the selected VPG.

Result – Whether the test passed or failed.

Notes – A description of the test. For example, defines where external files that describe the tests are saved. Notes are limited to 255 characters.

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

TASKS

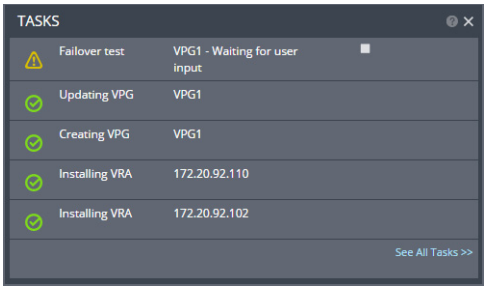

Monitor the recent tasks by clicking the TASKS area in the status bar at the bottom of the Zerto User Interface. The following information is displayed for the most recent tasks:

- The task status.
- The name of the task.
- A description of the task.

Also, actions, such as stopping a failover test, can be performed from this dialog.

Click *See All Tasks* to access *MONITORING > TASKS*.

Restore Volumes Dialog

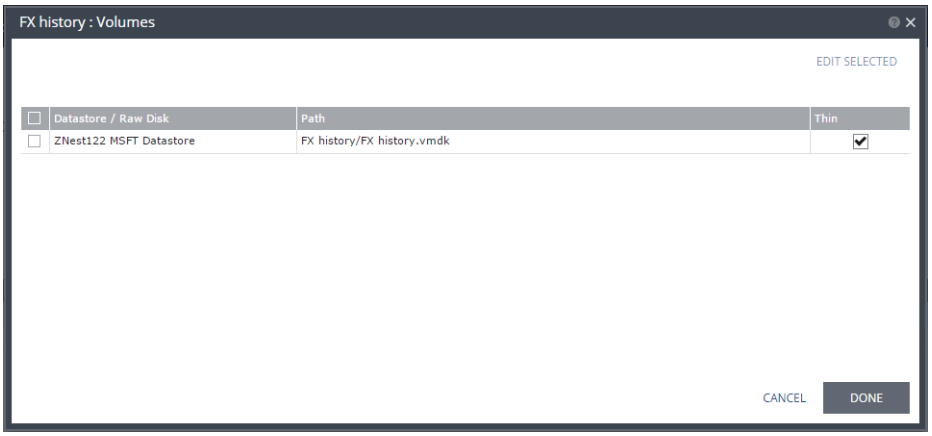

When restoring an offsite backup to the recovery site, this dialog shows the datastores for a selected virtual machine in the VPG. You can choose to edit information in one field by clicking the field and updating the information. You can choose to edit several datastore settings at the same time by selecting the datastores and clicking *EDIT SELECTED*.

CHAPTER 20: GLOSSARY

Access Key (AWS)	An alphanumeric text string that uniquely identifies the AWS account owner. No two accounts can have the same AWS Access Key.
Amazon Web Services (AWS)	A collection of remote computing services, also called web services, that make up a cloud computing platform by Amazon.com. The most central and well-known of these services are Amazon EC2 and Amazon S3. The service is advertised as providing a large computing capacity (potentially many servers) much faster and cheaper than building a physical server farm.
Asynch Replication	See Replication, Asynchronous.
Backup	See Extended Recovery.
Bare Metal	A computer system or network in which a virtual machine is installed directly on hardware rather than within the host operating system (OS).
Bitmap Sync[1]	A change tracking mechanism of the protected machines during a disconnected state when Zerto Virtual Replication starts to maintain a smart bitmap in memory to track and record changed storage areas. Since the bitmap is kept in memory, Zerto Virtual Replication does not require any LUN or volume per VPG at the source side. The bitmap is small and scales dynamically, containing references to the areas of the source disk that have changed but not the actual I/O. The bitmap is stored locally on the VRA within the available resources. For example, when a VRA goes down and is then rebooted. When required, Zerto Virtual Replication starts to maintain a smart bitmap in memory, to track and record storage areas that change. When the issue that caused the bitmap sync is resolved, the bitmap is used to check updates to the source disks and send any updates to the recovery site. A bitmap sync occurs during the following conditions: ■ Synchronization after WAN failure or when the load over the WAN is too great for the WAN to handle, in which case the VPGs with the lower priorities will be the first to enter a `Bitmap Sync`. ■ When there is storage congestion at the recovery site, for example when the VRA at the recovery site cannot handle all the writes received from the protected site in a timely fashion. ■ When the VRA at the recovery site goes down and is then rebooted. During the synchronization, new checkpoints are not added to the journal but recovery operations are still possible. If a disaster occurs requiring a failover during a bitmap synchronization, you can recover to the last checkpoint written to the journal. **Note:** For the synchronization to work, the protected virtual machines must be powered on. The VRA requires an active IO stack to access the virtual machine data to be synchronized across the sites. If the virtual machine is not powered on, there is no IO stack to use to access the source data to replicate to the target recovery disks.
Bucket (AWS)	Amazon buckets are like a container for your files. You can name your buckets the way you like but it should be unique across the Amazon system.
Business Continuity & Disaster Recovery (BC/DR)	An organization's ability to recover from a disaster and/or unexpected event and resume or continue operations. A disaster recovery, DR, plan is a subset of a Business Continuity plan. Organizations should have a business continuity, BC, plan in place that outlines the logistics and business operations. The key metrics to be measured in a disaster recovery environment are the Recovery Point Objective (RPO) and Recovery Time Objective (RTO).
Business Continuity Management (BCM)	Holistic management process that identifies potential threats to an organization and the impacts to business operations that those threats, if realized, might cause, and which provides a framework for building organizational resilience with the capability for an effective response that safeguards the interests of its key stakeholders, reputation, brand and value-creating activities. (ISO 22313, formerly BS 25999-1).

Business Continuity Plan	Contains the instructions, procedures and guidelines that are developed and maintained in readiness for use during and after any potentially disruptive event in order to enable the organization to continue to deliver its critical activities at an acceptable, predefined level.
Business Impact Analysis (BIA)	The process of analyzing business functions and processes and the effects that a business disruption might have upon them.
Checkpoint	Zerto Virtual Replication ensures crash consistency by writing checkpoints to the journal every few seconds. These checkpoints ensure write order fidelity and crash-consistency to each checkpoint. During recovery you pick one of these crash-consistent checkpoints and recover to this point. Additionally, checkpoints can be manually added by the administrator, with a description of the checkpoint. For example, when an event is going to take place that might result in the need to perform a recovery, you can pinpoint when this event occurs as a checkpoint in each journal.
Cloud Service Provider (CSP)	A service provider that offers customers storage or software services available via a private (private cloud) or public network (cloud). Usually, it means the storage and software is available for access via the Internet. Typically Infrastructure as a Service (IaaS), Software as a Service (SaaS), or Platform as a Service (PaaS) – are offered to their customers. Zerto enables them to offer Disaster Recovery As A Service (DRaaS) and In-Cloud DR (ICDR), too.
Crisis Management Plan	Provides the overall coordination of the organization's response to a crisis (which is a critical event that needs to be handled appropriately to prevent a damaging impact to the organization's profitability, reputation or ability to operate).
Data Deduplication	A specialized data compression technique for eliminating duplicate copies of repeating data.
Delta Sync[1]	The *Delta Sync* uses a checksum comparison to minimize the use of network resources. A Delta Sync is used when the protected virtual machine disks and the recovery disks should already be synchronized, except for a possible few changes to the protected disks, for example, when the target recovery disk is defined as a preseeded (not available in the cloud) disk or after a VRA upgrade, or for reverse protection after a move or failover. During the synchronization, new checkpoints are not added to the journal but recovery operations are still possible. If a disaster occurs requiring a failover during a delta synchronization, you can recover to the last checkpoint written to the journal. **Note:** For the synchronization to work, the protected virtual machines must be powered on. The VRA requires an active IO stack to access the virtual machine data to be synchronized across the sites. If the virtual machine is not powered on, there is no IO stack to use to access the source data to replicate to the target recovery disks.
Disaster	The occurrence of one or more events which, either separately or cumulatively, activate disaster recovery.
Disaster Recovery	The ability to restart operations after an interruption to the business according to a plan that ensures an orderly and timely restoration.
Disaster Recovery Plan	The disaster recovery, DR, plan is a component of the Business Continuity plan that details the process and procedures to recover the organization's resources to continue business operations. The Technology DR plan focuses on the IT disaster recovery. Also see Business Continuity Plan.
Disaster Recovery As A Service (DRaaS)	A disaster recovery solution that incorporates a service provider to replace or augment the organization's data protection implementation. In a DRaaS scenario, the customer may manage and have complete control over the production data. The Cloud Service Provider (CSP) may provide a partial or completely managed service. In either case, the CSP must ensure the availability of the data and adapt as the customers infrastructure changes. An advantage of this model is the CSP has dedicated resources skilled in DR operations.
DRS (vSphere)	Enables balancing computing workloads with available resources in a VMware vCenter cluster.
Emergency Management	Covers the immediate response to a situation or set of circumstances that present a clear and present threat to the safety of personnel or other assets of the organization.
Estimated Recovery Time (ERT)	This is the estimated timings based on full resource provision available during a live invocation. This time typically sits between the Net Recovery Time and the Recovery Time Achieved (RTA) time.

ESX/ESXi (vSphere)	*Bare-metal* hypervisor from VMware, meaning it installs directly on top of the physical server and partitions it into multiple virtual machines that can run simultaneously, sharing the physical resources of the underlying server. ESXi is the most recent version.
Extended Recovery	Extended DR includes the ability to configure both disaster recovery and offsite backups for the protected virtual machines in the VPG, according to a user-defined data retention policy.
High Availability (VMHA)	VMware high availability decreases downtime and improves reliability with business continuity by enabling another ESX/ESXi host to start up virtual machines that were running on another ESX/ESXi host that went down. High availability is automatically disabled by Zerto Virtual Replication while updating recovered virtual machines in the recovery site from the VRA journal. After the promotion of the data from the journal to the virtual machine completes, high availability is automatically re-enabled. The HA configuration can include admission control to prevent virtual machines being started if they violate availability constraints. If this is the case, then a failover, test failover or migration of the virtual machines in a VPG to the cluster with this configuration will fail, if the availability constraints are violated when the virtual machines are recovered.
Hyper-V	A hybrid hypervisor, which is installed in the operating system. However, during installation it redesigns the operating system architecture and becomes just like a next layer on the physical hardware.
Hypervisor	The host for multiple VMs in a virtualized environment. vSphere, ESX/ESXi, is the VMware brand hypervisor. The hypervisor is the virtualization architecture layer that allows multiple operating systems, termed guests, to run concurrently on a host computer.
Hypervisor Manager	The tool used to manage the host. For example VMware vCenter Server and Microsoft SCVMM.
I/O (Input/Output)	Describes any operation, program, or device that transfers data to or from a computer. Typical I/O devices are printers, hard disks, keyboards, and mouses. In fact, some devices are basically input-only devices (keyboards and mouses); others are primarily output-only devices (printers); and others provide both input and output of data (hard disks, diskettes, writable CD-ROMs). In computer architecture, the combination of the CPU and main memory (memory that the CPU can read and write to directly, with individual instructions) is considered the brain of a computer, and from that point of view any transfer of information from or to that combination, for example to or from a disk drive, is considered I/O.
In-Cloud DR (ICDR)	A disaster recovery solution that incorporates a service provider to replace or augment the organization's data protection implementation. When customers leverage an ICDR service, the CSP hosts the production and DR sites. The virtual machines (VMs) are typically replicated from one CSP datacenter to another CSP datacenter as a managed service or as managed co-located datacenters. The customers have the ability to interact with their applications as if they were locally hosted.
Initial Sync[1]	Synchronization performed after creating the VPG to ensure that the protected disks and recovery disks are the same. Recovery operations cannot occur until after the initial synchronization has completed. Adding a virtual machine to a VPG is equivalent to creating a new VPG and an initial synchronization is performed. In this case, any checkpoints in the journal become unusable and only new checkpoints added after the initial synchronization completes can be used in a recovery. The data in the journal however remains and is promoted to the recovered virtual machine as part of a recovery procedure. **Note:** For the synchronization to work, the protected virtual machines must be powered on. The VRA requires an active IO stack to access the virtual machine data to be synchronized across the sites. If the virtual machine is not powered on, there is no IO stack to use to access the source data to replicate to the target recovery disks.
iSCSI	An Internet Protocol (IP)-based storage networking standard for linking data storage facilities. By carrying SCSI commands over IP networks, iSCSI is used to facilitate data transfers over intranets and to manage storage over long distances.

Journal	Every write to a protected virtual machine is intercepted by Zerto Virtual Replication and a copy of the write is sent, asynchronously, to the recovery site, while the write continues to be processed on the protected site. On the recovery site the write is written to a journal managed by the Virtual Replication Appliance. Each protected virtual machine has its own journal. Each journal can expand to a size specified in the VPG definition and automatically shrinks when the expanded size is not needed.
LUN	Disk drives are the foundation of data storage, but operating systems cannot use physical disk storage directly. The platters, heads, tracks and sectors of a physical disk drive must be translated into a logical space, which an OS sees as a linear address space comprised of fixed-size blocks. This translation creates a logical entity that allows operating systems to read/write files. Storage networks must also partition their physical disks into logical entities so that host servers can access storage area network (SAN) storage. Each logical portion is called a logical unit number (LUN). A LUN is a logical entity that converts raw physical disk space into logical storage space, which a host server's OS can access and use. Any computer user recognizes the logical drive letter that has been carved out of their disk drive. For example, a computer may boot from the C: drive and access file data from a different D: drive. LUNs do the same basic job.
Level of Business Continuity	The reduced level of service that has been agreed if there is an interruption to business operations.
Managed Service Provider (MSP)	See Cloud Service Provider (CSP).
Maximum Tolerable Data Loss	The maximum tolerable data loss an organization can endure without compromising its business objectives.
Maximum Tolerable Outage (MTO)	The maximum time after which an outage will compromise the ability of the organization to achieve its business objectives.
Maximum Tolerable Period of Disruption	The duration after which an organization's viability will be irrevocably threatened if product and service delivery cannot be resumed.
NAS	A network-attached storage (NAS) device is a server that is dedicated to nothing more than file sharing. NAS does not provide any of the activities that a server in a server-centric system typically provides, such as e-mail, authentication or file management. NAS allows more hard disk storage space to be added to a network that already utilizes servers without shutting them down for maintenance and upgrades. With a NAS device, storage is not an integral part of the server. Instead, in this storage-centric design, the server still handles all of the processing of data but a NAS device delivers the data to the user. A NAS device does not need to be located within the server but can exist anywhere in a LAN and can be made up of multiple networked NAS devices.
Net Recovery Time	The net time achieved in recovering one or more VPGs after a disaster.
Offsite Backup	See Extended Recovery.
Operational Level Agreement (OLA)	The agreement between the service management and the Service Provision Partners. It defines the responsibilities for support and delivery of the services provided.
Pair	Zerto Virtual Replication can be installed at one or more sites and each of these sites can connect to any of the other sites enabling enterprises to protect virtual machines across multiple vCenters or within the same vCenter. Two sites connected to each other are considered *paired*. Also see Replication to Self.
Preseed	A virtual disk (a. vmdk flat file and descriptor or a .vhdx file) in the recovery site that has been prepared with a copy of the protected data. Using this option is recommended particularly for large disks so that the initial synchronization is much faster. When not using a preseeded disk the initial synchronization phase has to copy the whole disk over the WAN. Zerto Virtual Replication takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. Note that preseeding is not available in the cloud.

Quiesce	Pausing or altering the state of running processes on a computer, particularly those that might modify information stored on disk during a backup, in order to guarantee a consistent and usable backup. Critical applications, such as databases have quiescent mechanisms that Zerto Virtual Replication can use to get application consistent checkpoints.
RBAC	Role-based Access control, available in the Zerto Cloud Manager via the *Permissions* tab.
RDM (vSphere)	RDM is a mapping file in a separate VMFS volume that acts as a proxy for a raw physical storage device. The RDM allows a virtual machine to directly access and use the storage device. The RDM contains metadata for managing and redirecting disk access to the physical device. The file gives you some of the advantages of direct access to a physical device while keeping some advantages of a virtual disk in VMFS. As a result, it merges VMFS manageability with raw device access. Zerto Virtual Replication supports both physical and virtual mode RDMs.
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.
Recovery Time Achieved (RTA)	The actual times achieved during a DR test.
Recovery Time Objective (RTO)	Related to downtime. The metric refers to the amount of time it takes to recover from a data loss event and how long it takes to return to service. The metric is an indication of the amount of time the system's data is unavailable or inaccessible, thus preventing normal service.
Replication, Asynchronous	Technique for replicating data between databases or file systems where the system being replicated does not wait for the data to have been recorded on the duplicate system before proceeding. Asynchronous Replication has the advantage of speed, at the increased risk of data loss during due to communication or duplicate system failure.
Replication to Self	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 have to enable replication to the same vCenter Server and pairing is not required.
Resource	The elements (such as staff, site, data, IT systems) that are required to deliver an activity or service.
Resource Recovery Plan	Contains the instructions, procedures and guidelines to recover one or more resources and return conditions to a level of operation that is acceptable to the organization. Recovery Plans include detailed recovery procedures for IT equipment and infrastructure.
Rolling Back	Rolling back to an initial status, for example, after canceling a cloning operation on the VPG.
RPO	See Recovery Point Objective (RPO).
RTO	See Recovery Time Objective (RTO).
SAN	A storage area network (SAN) is any high-performance network whose primary purpose is to enable storage devices to communicate with computer systems and with each other. A storage device is a machine that contains nothing but a disk or disks for storing data. A SAN's architecture works in a way that makes all storage devices available to all servers on a LAN or WAN. As more storage devices are added to a SAN, they too will be accessible from any server in the larger network. In this case, the server merely acts as a pathway between the end user and the stored data. Because stored data does not reside directly on any of a network's servers, server power is utilized for business applications, and network capacity is released to the end user.
SCSI	Acronym for Small Computer System Interface. SCSI is a parallel interface standard used by many servers for attaching peripheral devices to computers. SCSI interfaces provide for faster data transmission rates (up to 80 megabytes per second) than standard serial and parallel ports. In addition, you can attach many devices to a single SCSI port, so that SCSI is really an I/O bus rather than simply an interface.
SCVMM	A Microsoft management solution for the virtualized datacenter, enabling you to configure and manage your virtualization host, networking, and storage resources in order to create and deploy virtual machines and services to private clouds that you have created.

Secret Access Key (AWS)	A password. The Secret Access Key with the Access Key forms a secure information set that confirms the user's identity.
Security Group	A virtual firewall that controls the traffic for one or more instances.
Service Continuity Plan	The continuity plan that acts as an umbrella document for a service, referencing other plans as required and providing service-specific emergency management and recovery plans.
Service Level Agreement (SLA)	The agreement between the customer and service provider which defines the service that is to be delivered to the customer.
Service Profile	A predefined set of default properties to use when VPGs are defined or edited. Zerto provides a default service profile and the option for the organization to specify their own requirements. The cloud service provider can define service profiles to manage specific service level agreements (SLAs) with its customers.
Service Test Plan	Detailed plan defining the activities required to test the recovery of an individual IT service to meet business requirements documented in the RTO and RPO.
Shadow VRA	During normal operation, a VRA might require more disks than a single virtual machine can support. If this situation arises, the VRA creates new shadow VRA virtual machines, used by the VRA to maintain additional disks. These virtual machines must not be removed. A VRA can manage a maximum of 1500 volumes, whether these are volumes being protected or recovered.
Snapshots	A snapshot is a block device which presents an exact copy of a logical volume, frozen at some point in time. Typically this would be used when some batch processing, a backup for instance, needs to be performed on the logical volume, but you don't want to halt a live system that is changing the data. Zerto does NOT use a snapshot mechanism, but is constantly replicating data writes.
Storage Account (Azure)	Storage accounts are like a container for your files. You can name your storage account the way you like but it should be unique across the Azure system.
Subnet	A logical, visible subdivision of an IP network.[1] The practice of dividing a network into two or more networks is called subnetting.
Subscription (Azure)	The description uses information derived from the following site: https://blogs.msdn.microsoft.com/arunrakwal/2012/04/09/create-windows-azure-subscription/ An Azure subscription grants access to Azure services and Platform Management Portal. A subscription has two aspects: ■ The Windows Azure account, through which resource usage is reported and services are billed. ■ The subscription itself, which governs access to and use of the Azure services that are subscribed to.
System Center Virtual Machine Manager	See SCVMM.
Virtual Machine (VM)	A virtual machine (VM) is an environment, usually a program or operating system, which does not physically exist but is created within another environment. In this context, a VM is called a *guest* while the environment it runs within is called a *host*.
Virtual Network (VNet) (Azure)	A virtual network dedicated to an Azure subscription.
Virtual Private Cloud (VPC) (AWS)	An on demand configurable pool of shared computing resources allocated within a public cloud environment, providing a certain level of isolation between the different organizations (denoted as users hereafter) using the resources. The isolation between one VPC user and all other users of the same cloud (other VPC users as well as other public cloud users) is achieved normally through allocation of a Private IP Subnet and a virtual communication construct (such as a VLAN or a set of encrypted communication channels) per user.
Virtual Protection Group	See VPG.
Virtual Replication Appliance	See VRA.

VMDK, Virtual Machine Disk	Virtual Machines created with VMware products typically use virtual disks. The virtual disks, stored as files on the host computer or remote storage device, appear to the guest operating systems as standard disk drives.
Volume Delta Sync[1]	Synchronization when only delta changes for a volume needs synchronizing, for example, when a virtual machine is added to a VPG using a preseeded disk. During the synchronization, new checkpoints are not added to the journal. Also, recovery operations are not possible during a `Volume Delta Sync`. For the synchronization to work, the protected virtual machines must be powered on. The VRA requires an active IO stack to access the virtual machine data to be synchronized across the sites. If the virtual machine is not powered on, there is no IO stack to use to access the source data to replicate to the target recovery disks. Preseeding is not available in the cloud.
Volume Full Sync[1]	Synchronization when a full synchronization is required on a single volume. During the synchronization, new checkpoints are not added to the journal. Also, recovery operations are not possible during a `Volume Full Sync`. **Note:** For the synchronization to work, the protected virtual machines must be powered on. The VRA requires an active IO stack to access the virtual machine data to be synchronized across the sites. If the virtual machine is not powered on, there is no IO stack to use to access the source data to replicate to the target recovery disks.
Volume Initial Sync[1]	Synchronization when a full synchronization is required on a single volume, for example, when changing the target datastore or adding a virtual machine to the VPG without using a preseeded (not available in the cloud) disk. During the synchronization, new checkpoints are not added to the journal. Also, recovery operations are not possible during a `Volume Initial Sync`. For the synchronization to work, the protected virtual machines must be powered on. The VRA requires an active IO stack to access the virtual machine data to be synchronized across the sites. If the virtual machine is not powered on, there is no IO stack to use to access the source data to replicate to the target recovery disks.
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 for the Web Server, require that all three virtual machines are replicated to maintain data integrity.
VRA	A virtual machine installed on each hypervisor hosting virtual machines to be protected or recovered, that manages the replication of protected virtual machine writes across sites. A VRA must be installed on every hypervisor that hosts virtual machines that require protecting in the protected site and on every hypervisor that will host the replicated virtual machines in the recovery site.
vSphere	VMware's server virtualization platform for building a cloud infrastructure.
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).
Zerto User Interface	Recovery using Zerto Virtual Replication is managed via a user interface: in a browser via the Zerto Virtual Manager Web Client, or in either the vSphere Web Client or vSphere Client console in the Zerto tab.
Zerto Self-service Portal (ZSSP)	An out-of-the-box DR portal solution with a fully functioning browser-based service portal to enable cloud service providers to quickly introduce disaster recovery as part of their portal offering.
Zerto Virtual Backup Appliance (VBA)	A Zerto Virtual Replication service that manages the offsite backup.
Zerto Virtual Manager (ZVM)	A Windows service, which manages everything required for the replication between the protection and recovery sites, except for the actual replication of data. The ZVM interacts with the vCenter Server to get the inventory of VMs, disks, networks, hosts, etc. The ZVM also monitors changes in the VMware environment and responds accordingly. For example, a vMotion operation of a protected VM from one host to another is intercepted by the ZVM so the Zerto User Interface is updated accordingly.
ZORG, Zerto Organization	Cloud customers are defined to Zerto Cloud Manager as Zerto organizations, ZORGs. A ZORG is defined with the cloud resources it can use, the permissions that it has to perform operations, such as testing a failover or defining a VPG.

1. Synchronization after a recovery starts after the promotion of data from the journal to the virtual machine disks ends. Thus, synchronization of virtual machines can start at different times, dependent on when the promotion for the virtual machine ends. All synchronizations are done in parallel, whether a delta sync or full sync, etc.

F

G

H

I

J

K

L

M

N

O

T

V

W

Z