Zerto

Zerto Virtual Manager Administration Guide

AWS Environment

Version 7.0

ZVR-ADVA-7.0 Rev01 U2 July2019

Copyright © 2019, Zerto Ltd. All rights reserved.
Information in this document is confidential and subject to change without notice and does not represent a commitment on the part of Zerto Ltd. Zerto Ltd. does not assume responsibility for any printing errors that may appear in this document. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's personal use, without the prior written permission of Zerto Ltd.
All other marks and names mentioned herein may be trademarks of their respective companies.

The scripts are provided by example only and are not supported under any Zerto support program or service.
All examples and scripts are provided "as-is" without warranty of any kind. The author and Zerto further disclaim all implied warranties including, without limitation, any implied warranties of merchantability or of fitness for a particular purpose.
In no event shall Zerto, its authors, or anyone else involved in the creation, production, or delivery of the scripts be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or other pecuniary loss) arising out of the use of or inability to use the sample scripts or documentation, even if the author or Zerto has been advised of the possibility of such damages. The entire risk arising out of the use or performance of the sample scripts and documentation remains with you.

ZVR-ADVA-7.0 Rev01 U2 July2019

Zerto

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION TO THE ZERTO SOLUTION 7

What is the Zerto Solution? 7
The Zerto Solution Architecture 9
How Zerto Recovery Works 11
Zerto Analytics - Overview 12
    Zerto Analytics Secure Architecture 17
Zerto Cloud Control - Overview 21
    Before Getting Started with Zerto Cloud Control 22
    Accessing the Zerto Cloud Control 22
    Using Upgrade Manager 22
    Using Tenants 22
Benefits of Using Zerto 22
Import Methods for AWS 26
How the Zerto Solution for AWS Works 29

CHAPTER 2: ACCESSING THE ZERTO USER INTERFACE 30

Using the Zerto Virtual Manager Web Client 30
Adding a Security Certificate for the Zerto User Interface 30
Working With the Zerto User Interface 32

CHAPTER 3: INTRODUCTION TO PROTECTING VIRTUAL MACHINES 34

CHAPTER 4: OVERVIEW OF RECOVERY FLOWS 37

CHAPTER 5: PROTECTING VIRTUAL MACHINES FROM AWS 39

Replication From a Protected Site AWS to a Recovery Site vCenter Server 40
    Maintained VM Properties when Recovering to a vCenter Server Site from a Protected AWS Site 58
Replication From a Protected Site AWS to a Recovery Site vCloud Director 59
Replication From a Protected Site AWS to a Recovery Site Hyper-V 76
Replication From a Protected Site AWS to a Recovery Site Azure 94
    Requirements for Microsoft Azure Environments 94
    VPGs Recovering to Azure Standard Storage and Premium Managed Disks 94
    Converting Premium Virtual Machines for Protection 94
    Protecting From Amazon Web Services - To a Microsoft Azure Recovery Site 97

CHAPTER 6: MONITORING ZERTO VIRTUAL MANAGER 103

The DASHBOARD Tab 103
Monitoring VPGs – The VPGs Tab 105
    List View - GENERAL 105
    List View - PERFORMANCE 106
    List View - RETENTION STATUS 106
    Additional Fields and Options 107
Monitoring a Single VPG 107
Monitoring Tasks 111
Monitoring Protected Virtual Machines – The VMs Tab 112
Monitoring Peer Sites - The SITES Tab 114

CHAPTER 7: MANAGING VPGS 116

Editing a VPG 116
Modifying the Journal Size Hard Limit 117
Adding Virtual Machines to a VPG 117
Removing Virtual Machines from a VPG 118
Pausing and Resuming the Protection of a VPG 118
Forcing the Synchronization of a VPG 119
Deleting a VPG 119
Deleting a VPG When the Status is Deleting 120
Ensuring Application Consistency – Checkpoints 120
Adding a Checkpoint to Identify a Key Point 121
Running Scripts Before or After Recovering a VPG 121
Creating a Script 122
Example Scripts 124
Exporting and Importing VPG Definitions 124
VPG Statuses and Synchronization Triggers 125
VPG Statuses 126
VPG Synchronization Triggers 130

CHAPTER 8: MANAGING ZERTO VIRTUAL MANAGER 132

Check Connectivity Between Zerto Components 132
Reconfiguring the Zerto Virtual Manager Setup 133
Reconfiguring the Microsoft SQL Server Database Used by the Zerto Virtual Manager 134
Replacing the SSL Certificate 135
Pair to Another Site and Unpairing Sites 135

CHAPTER 9: OVERVIEW OF DISASTER RECOVERY OPERATIONS 137

The Failover Test Operation 137
The Move Operation 138
The Failover Operation 138
The Restore File Operation 139

CHAPTER 10: ADVANCED SITE CONFIGURATION 140

Site Settings 140
Editing Information About a Site 141
Defining Site Policies 143
Configuring Email Settings 144
Defining the Resource Report Sampling Period 145
Seeing What is Licensed 145
Submitting a Support Case 146
Submitting a Feature Request 147
About the Zerto Version 148

CHAPTER 11: TESTING RECOVERY TO AWS 149

The Test Failover Process 149
Starting and Stopping Failover Tests 150
After Starting a Test, What Happens? 152
Viewing Test Results 154
Live Disaster Recovery Testing 154
Basic Verification – User Traffic Is Not Run against the Recovered VMs 155

CHAPTER 12: MIGRATING A VPG TO AWS 157

The Move Process 157
Moving Protected Virtual Machines to a Remote Site 159

CHAPTER 13: MIGRATING A VPG FROM AWS 163

The Move Process 163
Moving Protected Virtual Machines From AWS to a Remote Site 164
Reverse Protection For a Moved VPG 167

CHAPTER 14: MANAGING FAILOVER LIVE TO AWS 168

The Failover Live Process 168
Initiating Failover Live 169
What Happens When the Protected Site is Down 174
Initiating Failover Live During a Test 175

CHAPTER 15: CLONING A VPG TO AWS 176

The Clone Process 176
Cloning Protected Virtual Machines to the Remote Site 176

CHAPTER 16: RECOVERING FILES AND FOLDERS 179

The File and Folder Recovery Process 179
Recovering Files and Folders 180
Mounting the Disk that Contains the Required Files and Folders 180
Downloading the Files and Folders from the Disk 183

CHAPTER 17: FAILING BACK FROM AWS 186

The Failback (Move) Process 186
Failback (Move) After the Original Site is Operational 187
Initiating a Failback (Move) 187
Reverse Protection For a Moved VPG 190

CHAPTER 18: ZERTO REPORTS 192

Recovery Reports 192
Resources Report 193
VPG Performance 197

CHAPTER 19: TROUBLESHOOTING 198

Ensuring the Zerto Virtual Manager is Running 198
Troubleshooting: "Needs Configuration" Problems 198
Troubleshooting VRA Problems 199
Zerto Diagnostics Utility 199
Collecting Zerto Logs 199
Using Remote Log Collection 199
Using the Zerto Diagnostics Tool 200
Understanding the Logs 205

CHAPTER 20: THE ZERTO VIRTUAL MANAGER USER INTERFACE 206

About Dialog 206
Add Checkpoint Dialog 207
Add Site Dialog 208

Pair sites 208
Advanced Journal Settings Dialog 208
Advanced VM Replication Settings Dialog 210
Advanced VM Settings for Cloud Dialog 210
ALERTS 211
Boot Order Dialog 212
Checkpoints Dialog 213
Edit NIC Dialog 214
Edit vNIC Dialog 215
Edit VM Settings Dialog 216
License Dialog 216
New Repository Dialog 217
Offsite Clone Dialog 219
Open Support Case Dialog 220
Remote Support Dialog 221
Restore VPG - NICs Dialog 222
Restore VPG - Volumes Dialog 222
Site Settings Dialog 222
Site Information Dialog 223
Policies Dialog 225
Email Settings Dialog 226
Reports Dialog 226
Stop Failover Test Dialog 226
TASKS 227

CHAPTER 21: GLOSSARY 228

CHAPTER 1: INTRODUCTION TO THE ZERTO SOLUTION

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's Long Term Retention is the additional process of enabling recovery of IT processing tasks after an extended period. This section describes Zerto solution's general concepts to enable replication and recovery in a virtual environment.

The following topics are described in this section:

- "What is the Zerto Solution?", below
- "The Zerto Solution Architecture", on page 9
- "How Zerto Recovery Works", on page 11
- "Zerto Analytics - Overview", on page 12
- "Zerto Cloud Control - Overview", on page 21
- "Benefits of Using Zerto", on page 22
- "Import Methods for AWS", on page 26
- "How the Zerto Solution for AWS Works", on page 29

What is the Zerto Solution?

IT Resilience Platform™

Zerto's IT Resilience Platform™ converges disaster recovery, backup, and workload mobility whether on-premises or to, from and between hybrid and multi-cloud platforms. The platform is built on a foundation of continuous data protection (CDP) with built-in orchestration and automation providing IT leaders with simplicity, enterprise scale and agile data protection to save time, resources and costs. Analytics, with intelligent dashboards and live reports, provides complete visibility across multi-site and multi-cloud environments, giving organizations confidence that business service levels and compliance needs are met.

Disaster Recovery

Zerto is built from the ground up to be the simplest, most powerful disaster recovery solution for virtualized infrastructures. By including all the replication, recovery orchestration and automation in one simple software platform, users can recover one, all or a subset of virtualized applications, from anywhere to anywhere, maximizing the benefits of virtualization and cloud.

Through native integration into all supported platforms, Zerto not only allows replication and recovery between any storage, but it also protects across and between multiple hypervisors and public cloud platforms. This market-leading technology delivers a best of breed Business Continuity/Disaster Recovery (BC/DR) solution irrespective of underlying hypervisor, public cloud or storage.

Zerto supports replication between VMware vSphere, Microsoft Hyper-V, Microsoft Azure, Amazon Web Services (AWS), and IBM Cloud. Zerto is the underpinning technology of hundreds of Cloud Service Providers (CSPs) globally who are leading the way in delivering Disaster Recovery as a Service (DRaaS) capabilities.

Backup

Elastic Journal

The Zerto Elastic Journal is a new concept in data protection bringing together both short-term retention and long-term retention capabilities. Using unified data protection workflows, powered by intelligent index and search, it enables quick recovery of data-regardless of whether it's from a few seconds ago or a few years ago. Data gets copied from the short-term retention storage into the long-term retention repositories. As the short-term retention is already stored on the target side, moving data to long-term retention has no impact on production, so copies can be taken as often as needed, eliminating the concept of "backup windows".

Workload Mobility

Workload mobility is the ability to move workloads to, from and between multiple platforms, on-premises or cloud, with minimal business impact and no traditional infrastructure constraints. With the recent rise of cloud computing, there is now a

plethora of cloud platforms to choose from, with large players such as Microsoft Azure, AWS and IBM Cloud as well as plenty of Cloud Service Providers (CSPs). Each of these platforms have their own attributes which make them suited to certain workloads, whether through compliance, pricing or otherwise. Despite this, the cloud is not always the right option for all workloads, and because of this, the traditional on-premises deployments with VMware vSphere or Microsoft Hyper-V are still prevalent. Ultimately, having all of these options has led to increasing adoption of a multi-cloud and/or hybrid cloud strategy. Organizations are challenged to ensure that workloads, whether in the cloud or on-premises, are not locked into any one platform or vendor, and can be moved around based on where they fit best today, rather than yesterday. This is where workload mobility comes in, removing the traditional lock-in to these platforms and allowing workloads to be moved across platforms seamlessly, on-premises or in the cloud, including the ability to validate mobility prior to the live event without production impact.

Zerto supports any-to-any mobility between VMware vSphere, Microsoft Hyper-V, AWS, IBM Cloud, Microsoft Azure and any of our hundreds of Zerto-powered Cloud Service Providers (CSPs). For a full list of supported platforms and their versions please see our Interoperability Matrix.

The Zerto Solution Architecture

Zerto helps customers accelerate IT transformation by eliminating the risk and complexity of modernization and cloud adoption. By replacing multiple legacy solutions with a single IT Resilience Platform™, Zerto is changing the way disaster recovery, retention and cloud are managed. This is done by providing enterprise-class disaster recovery and business continuity software for virtualized infrastructure and cloud environments.

In **on-premise** environments, Zerto Virtual Replication (ZVR) is installed with virtual machines to be protected and recovered.

In **public cloud** environments, Zerto Cloud Appliance (ZCA) is installed in the public cloud site that is to be used for recovery.

The installation includes the following:

- **Zerto Virtual Manager (ZVM):** The ZVM is a Windows service, running on a dedicated Windows VM, that manages everything required for the replication between the protection and recovery sites, except for the actual replication of data. The ZVM interacts with the hypervisor management user interface, such as vCenter Server or Microsoft SCVMM, to get the inventory of VMs, disks, networks, hosts, etc. It also monitors changes in the hypervisor environment and responds accordingly. For example, a VMware vMotion operation, or Microsoft Live Migration of a protected VM from one host to another is seen by the ZVM, and the Zerto User Interface is updated accordingly.
  - For the maximum number of virtual machines, either being protected or recovered to that site, see Zerto Scale and Benchmarking Guidelines.
- **Virtual Replication Appliance* (VRA)**: A VRA is a virtual machine installed on each hypervisor host where VMs are to be protected from or to. The VRA manages the replication of data from protected virtual machines to the recovery site. The target VRAs are responsible for maintaining any protected VMs disks. VMware limits the number of SCSI Controllers (4 per VM) and targets per Controller (15 per controller), leaving a maximum of 60 SCSI targets per VM. When managing a larger quantity of virtual disks, Zerto utilizes Virtual Replication Appliance Helpers (VRA-H), which act as a disk box with no IP and nearly no resources. VRA-Hs are spun up and down by Zerto automatically when nearing the 60-disk limit of the VRA or last VRA-H.
  - For the maximum number of volumes, either being protected or recovered to that site, see Zerto Scale and Benchmarking Guidelines.

  Note: *In vSphere installations, OVF to enable installing Virtual Replication Appliances.
- **Virtual Backup Appliance (VBA)**: A Windows service that manages File Level Recovery operations within the Zerto solution.
- **Zerto User Interface:** Recovery using the Zerto solution is managed in a browser or, in VMware vSphere Web Client or Client console.

When Zerto is installed to work with an on-premise hypervisor it also comprises the following component:

- **Data Streaming Service (DSS):** Installed on the VRA machine, and runs in the same process as the VRA. It is responsible for all the retention data path operations.

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

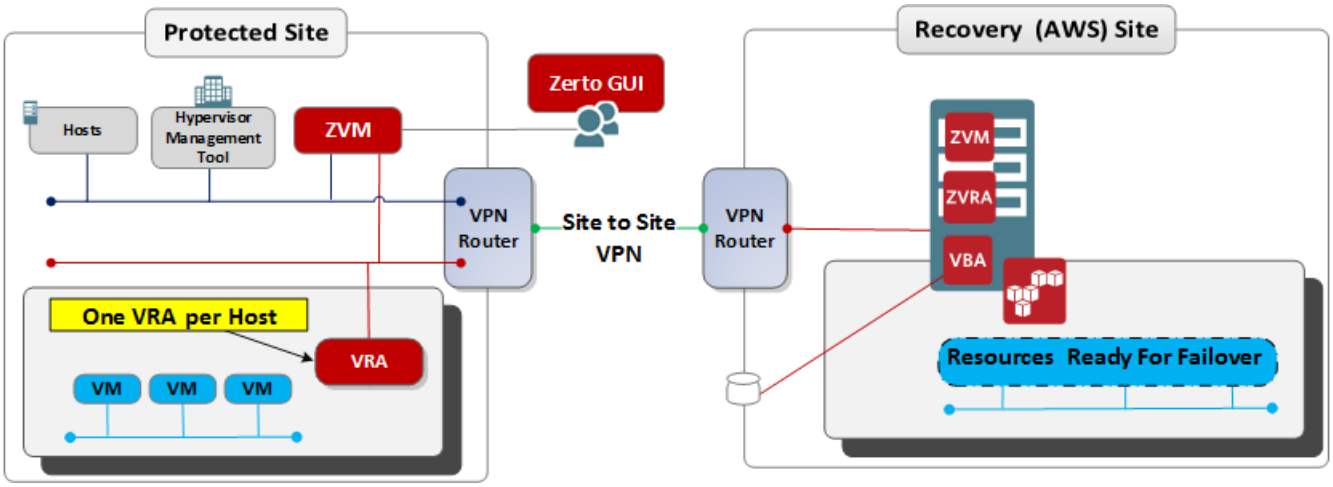

When you plan to recover the enterprise site to a public cloud, ZCA is installed in the cloud environment. ZCA comprises the same components but the VRA runs as a service, so that the ZVM, VRA, and VBA all run as services on a single virtual machine instance in the public cloud.

How Zerto Recovery Works

Installing Zerto installs the Zerto Virtual Manager, which sits in the hypervisor layer and the Zerto Cloud Appliance which sits in AWS. 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 data centers from one location to another.

Note: CD and DVD drives cannot be protected.

Every write is copied by Zerto 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, you can restore from files created during the Retention process. Retention is 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 retention sets at the recovery site.

Zerto Analytics - Overview

Insight-driven data analytics for a new era of data protection

As IT infrastructures become more complex and demands for performance rise, companies require visibility and control over protected IT environments. Visibility of your entire IT infrastructure (both on-premises or cloud) is imperative to monitor, analyze and plan your environment and resource requirements to ensure zero interruptions. To have confidence that business Service level Agreements (SLAs) are met, you need not only visibility and insights to address existing issues, but also to be able to plan for your future data protection needs.

Zerto Analytics delivers these capabilities through a single interface and one user experience for a comprehensive overview of your entire multi-site, multi-cloud environment. Utilizing metrics such as average recovery point objective (RPO), network performance, and storage consumption, Zerto Analytics delivers real-time and historical insights on the health and protection status of your applications and data. Through Intelligent dashboards you can spot trends, identify anomalies, and troubleshoot issues in network, RPO, and other business SLAs. With these insights, you can eliminate inefficiencies and allocate resources effectively to mitigate data loss, reduce downtime and take control of your data.

New Resource Planner

With Resource Planner, you can ensure your data protection needs are met as your IT environment grows and diversifies. Use the Resource Planner to monitor and perform analysis of your protected environment to determine the required compute, storage and networking resources for any or all virtual machines (VMs) within your environment, both on-premises or in the cloud. Build out your data protection strategy with 'what-if' scenario modeling to plan for VM protection across clouds or on-premises, or to transition from on premises to cloud, or to protect additional VMs within your environment.

Zerto Analytics is developed with an API first approach, therefore, everything that is presented in the GUI, is also available with APIs. The APIs are delivered with Swagger open source that help you develop and test REST integration using standardized examples. This allows to easily populate custom portals with Zerto Analytic content.

See also:

"Before Getting Started with Zerto Analytics", on page 12

"Accessing the Zerto Analytics Portal", on page 12

"Accessing Zerto Mobile", on page 13

"Zerto Analytics APIs", on page 13

"Using the Zerto Analytics Portal", on page 13

"End-User Analytics for Service Providers", on page 14

"Zerto Analytics Product Feature Matrix", on page 16

"Zerto Analytics Secure Architecture", on page 17

Before Getting Started with Zerto Analytics

Verify the following:

- At least **1 ZVM** running Zerto **5.0 or higher**.
- The checkbox **Enable Support notification and product improvement feedback** is selected. This is accessed in the ZVM application in **Settings > About**.
- Internet access.
- A **myZerto** account using your corporate email address.

Accessing the Zerto Analytics Portal

Zerto Analytics is accessed from one of the following locations:

- A **URL**: https://analytics.zerto.com
- **myZerto**: www.zerto.com/myzerto. Sign in using your credentials and select the **Analytics** tab.
- From the **ZVM Application Menu** tab: Click

to open Zerto Analytics in a new browser tab.

Accessing Zerto Mobile

The Zerto Mobile app is available for both **iOS** or **Android** operating systems.

Sign in using your myZerto credentials.

Zerto Analytics APIs

Zerto Analytics is developed with an API first approach, therefore, everything that is presented in the GUI, is also available with APIs. APIs are available the same version as their GUI counterparts.

Zerto Analytics APIs are available in OpenAPI Specification.

Zerto Analytics API documentation can be accessed via the link: https://docs.api.zerto.com/

Using the Zerto Analytics Portal

After accessing Zerto Analytics, the **Dashboard** opens by default. This view displays a summary of the entire protected environment, including the average RPO and VPG health, site details and topology, active alerts, running tasks and events.

Use the **What's New** WHAT'S NEW and **Help** widgets in Zerto Analytics to learn more about each of the features available in Zerto Analytics.

In addition to the Dashboard, there are additional views that provide details on the overall health of your environment.

- Use the sites topology to identify sites without remote protection, to identify network issues and to identify cloud issues.
- To handle any issues in the ZVM site, in the **Dashboard** tab > **SITES** area, click the icon **Open ZVM in new tab**. This routes you to the specific ZVM site.

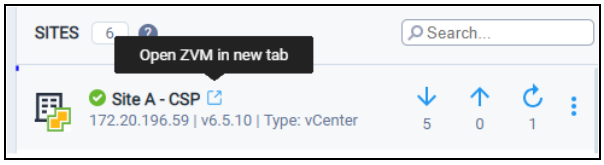

- In the same Sites area, click the Menu button to navigate to the VPGs, Alerts or Tasks for the specified site.

See the following sections:

1. Monitoring Alerts, Events and Tasks
2. Troubleshooting VPGs
3. Reviewing Reports

Monitoring Alerts, Events and Tasks

From the **Monitoring** tab review the active alerts, alerts history, running events and tasks.

To see active alerts, events or tasks from a specific ZVM, use the filter in the top left of the screen Add Filters +.

Events are displayed for the last 24 hours by default. Use the date selector to filter the Events List by date.

To view a history of alerts, select the **Events** tab and **Alert History** in the sub menu. Inactive alerts are displayed for the selected time range.

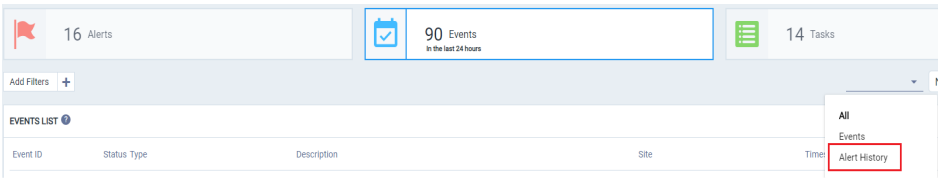

TIP:

Click the **Alert ID** and **Event ID** to open the help and view the full details for that specific alert or event.

Troubleshooting VPGs

From the **VPGs** tab, review the list of VPGs with Errors and Warnings. In the VPGs status area, click either VPGs with Warnings or VPGs with Errors to filter the VPGs list.

To review the status of a VPG, click in the column of the VPG you want to review to open the **VPG Details** page. From here you can view the details of the virtual machines associated with the VPG as well.

If you want to view the details of another VPG, select the VPG from the VPGs drop-down list

TIP:

To further investigate about a selected VPG, click the **VPG History** button, and select to view either the **RPO**, **Journal or Network** Reports page. (See Reviewing Reports for more details about Reports).

Reviewing Reports

The SaaS Analytics reports provide real-time and historical data analysis. ECE and CSP licenses can view up to 90 days of report history. Use the date controller to filter your reports. Statistics are displayed according to the selected time frame.

TIP:

Zoom in to view more granular data by selecting and dragging your mouse over the selected time frame.

From the **Reports** tab, review the RPO, Journal and Network performance history on the VPGs.

Use the **Network** reports for reviewing the network history for any VPG or Site. You can also view the network summary for a selected time frame, the network performance history and IOPs history.

Use the **RPO** reports for viewing a summary to see if RPO SLAs are being met, the RPO history and the RPO breach table for viewing when the specific time breach occurred and the duration of the SLA breach.

Use the **Journal** reports for understanding if resilience is at risk due to journal storage capacity and plan for storage growth.

End-User Analytics for Service Providers

Service providers can filter their customers in Zerto Analytics using the ZORG filter located at the top of the screen.

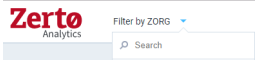

Using the ZORG filter, CSPs can see historical data and status for any individual customer.

Additionally, using the Zerto Analytics **APIs**, CSPs can create custom reports and automate reporting delivery of real-time content to their customers. The APIs can also be used to provide content for customer portals.

Planning your Data Protection Strategy

Zerto Analytics includes the Planning tab; a resource planning capability that performs continuous monitoring and analysis to determine the required storage and networking resources for your data protection strategy. In addition, it provides recommendation for the required number of ZCAs when replicating to Azure. Analysis can be performed for any or all VMs within your VMware environment. This capability uses data collected from VMware sites and does not require any new installation or license.

The planning capability delivers visibility and forecasting into:

- Zerto protected VMs from vSphere to a public cloud.
- Zerto protected VMs from vSphere to an on-premise site.
- The cost benefit analysis between deciding whether to protect VMs on-premises or in public cloud.

The resource Planning capability provides overall sizing recommendations, including number of ZCAs, WAN, Total Journal Size and Total Recovery Disk Size.

It also includes VM configuration details (CPU, RAM and Total Size) collected from vCenter, per VM calculated performance counters (Avg. Throughput, Avg. IOPS) and the Required Journal Size estimation per VM.

To begin planning your data strategy:

1. Click on the Planning tab; the VM Selection window opens.
2. From the VM Selection window, select your Source Site VMs and Target Platform for replication.

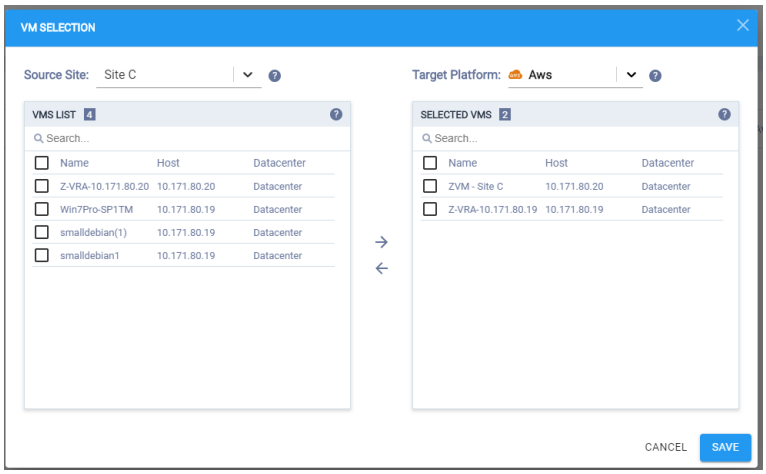

3. Click Save; The Planning view is displayed.

Use the Help widgets to learn more about each of the features available in Planning.

Zerto Analytics Product Feature Matrix

The following table lists the available features and from which ZVM version it's supported:

KEY
✓ Supported
× Not Supported
✓* Supported partially, see comments for further details.

FEATURE	ZVM V5.0 AND ABOVE	ZVM V5.5 AND ABOVE	ZVM V6.0 AND ABOVE	ZVM V6.5 AND ABOVE	ZVM V7.0 AND ABOVE	COMMENTS
Dashboard	✓	✓*	✓	✓	✓	Events are available from v5.5U4 and above.
VPG List	✓	✓	✓	✓	✓	--
VPG Details	✓*	✓	✓	✓	✓	Journal size data is available from v5.5 and above.
Monitoring: Alerts	✓	✓	✓	✓	✓	--
Monitoring: Tasks	✓	✓	✓	✓	✓	--
Monitoring: Events	×	✓*	✓	✓	✓	Events are available from v5.5U4 and above.
Reports: RPO	✓	✓	✓	✓	✓	--
Reports: Journal	✓*	✓	✓	✓	✓	Journal size data is available from v5.5 and above.
Reports: Network	×	×	✓	✓	✓	--
Site List & Topology	✓	✓	✓	✓	✓	--
End-User (ZORG) Filter	×	✓*	✓	✓	✓	ZORG filter is available from v5.5u4 and above.
90 Days History for ECE and Cloud	✓*	✓*	✓	✓	✓	Journal size data is available from v5.5 and above. Network reports are available from v6.0.
Licensing Usage	✓	✓	✓	✓	✓	--
Storage Analytics Tab	×	×	×	✓*	✓	Storage Analytics is available from v6.5u2.
Storage Analytics for VPG Volumes	×	×	×	✓*	✓	Storage Analytics is available from v6.5u2.
Planning	×	×	×	×	✓	Zerto Resource Planner is available from v7.0U1.

Zerto Analytics Secure Architecture

Transmitted Zerto Analytics data is securely stored and can be securely viewed and exported by the user. The data is also used by the analytics engine to provide advanced guidance and analysis. The following section describes the secure architecture that provides the path to transmit, the datastore where the data is housed, and access of that data.

ZVM sends all the data over HTTPS TLS 1.2 only to https://zerto-mobile-data.zerto.com. The data is stored in the Samples Storage in the SaaS backend, located in the public cloud. All communication to and from the cloud are encrypted (HTTPS) using TLS 1.2. Furthermore, the data is isolated in the SaaS backend and accessible only within the backend itself an via VPN from Zerto. End users are authenticated using their myZerto credentials.

Transmission and Storage of Site Status

ZVM transmits current site status over secured HTTPS to the SaaS backend in short (~1 minute) intervals. The SaaS backend collects site status data from all transmitting sites and stores it in the "Samples Storage". The "Samples Storage" is secured inside an internal protected network and is accessible only within the backend itself and via VPN from Zerto. The samples are kept in the "Samples Storage" of the Zerto Analytics cloud for 30 days to support 30 days of report history.

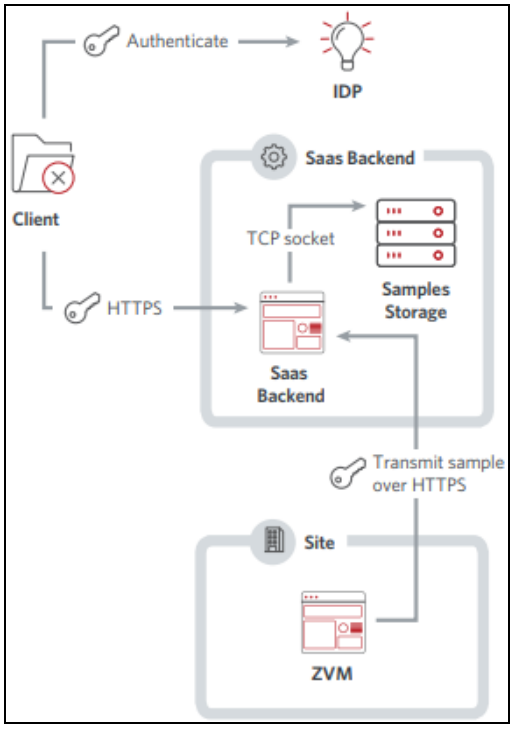

User Authentication and Authorization

Before sending any request to the backend, a client (browser/mobile/RESTful client) must obtain a security token which securely identifies the user. The security token is achieved by authenticating the user against the Identity Provider (IDP). Each client request contains the security token. For each request, the backend validates the security token and authorizes execution of the request according to the user's identity which is part of the security token's payload.

Reporting Historical Data

For each client request, the backend gets the latest data from the relevant samples as stored in the "Samples Storage". The relevant data is sent to the client.

Zerto Analytics Collected Data

The Zerto Analytics dashboard utilizes several metrics to deliver detailed monitoring and reporting of your multi-site Multi-Cloud environment. The following table was created to address privacy and security issues surrounding data collection.

This table outlines each piece of information that is collected and sent from the ZVM for Zerto Analytics.

COMPONENTS	PARAMETER	DESCRIPTION
Alerts		
	affectedZorgs	List of ZORGs affected by the alert.
	affectedZorgIds	List of ZORG IDs affected by the alert.
	alertIdentifier	Internal alert identifier.
	affectedVpgIds	List of VPG IDs affected by the alert.
	alertType	Help identifier associated with the alert.
	createdDateinUTC	The time the alerts was created.
	description	Alert description.
	entityType	The entity for which the alert applies.
	severity	Alert severity (Warning or Error).
	site	The site which is affected by the alert.
Tasks		
	taskName	The task type.
	taskIdentifier	Unique identifier of the task.
	status	The status of the task,
	startedDateInUtc	The time the task started.
	completedDateInUtc	The time the task was completed.
	progress	Progress of the task in percentage of completion.
	sitename	Name of the site.
	protectedDataInMB	Total data protected in the ZVM site.
	siteIdentifier	Internal site identifier.
	transmitterCycleInSec	Number of seconds between two (2) data transmissions.
	type	Site environment.
	initiatedBy	Name of user who initiated task.
	information	Task details.
	relatedVpgsList	List of VPGs related to the task.
	relatedSitesList	List of sites related to the task.

Site		
	collectionTimeInUtc	Time the ZVM data was collected.
	documentVersion	Version of the document.
	hashedLicenseKey	ZVM license hashed.
	isTransmissionEnabled	Whether transmission of data was enabled by the user.
	sitename	Name of the site.
	protectedDataInMB	Total data protected in the ZVM site.
	siteIdentifier	Internal site identifier.
	transmitterCycleInSec	Number of seconds between two (2) data transmissions.
	type	Site environment.
	utcOffsetInMinutes	Offset of the site time from UTC, in minutes.
	version	Zerto Virtual Manager version.
	vpgsAlertsCount	Number of active alerts in ZVM.
	zorgsCount	Number of ZORGs in the ZVM site.
	zvmIp	IP of the ZVM
VPGs		
	actualJournalHistoryInSeconds	The actual journal history, in seconds.
	actualRpoSeconds	Time since the last checkpoint was written to the journal in seconds.
	configuredJournalHistoryInMinutes	Configured journal history, in minutes.
	configuredRpoSeconds	Configured RPO, in seconds.
	earliestRecoveryPointLocalDateInUtc	Time of the earliest checkpoint in the journal.
	priority	Priority specified for the VPG.
	protectedSite	Protected site details (Site ID, Site Name, Site Type and ZVM version).
	recoverySite	Recovery site details (Site ID, Site Name, SiteType and ZVM version).
	status	Status of the VPG.
	subStatus	Substatus of the VPG.
	vpgIdentifier	Identifier of the VPG.
	vpgName	Name of the VPG.
	zorgName	Name of the Zerto organization set up in the Zerto Cloud Manager, which uses this VPG.

VMS		
	IOPS	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.
	journalHardLimit	Maximum journal size in MBs, or in a percentage of the virtual machine volume size.
	journalWarningThreshold	Journal size in MBs, or in a percentage of the virtual machine volume size, that generates a warning when the journal is nearing its hard limit.
	outgoingBandWidthInMbps	Bandwidth throttling defined for the virtual machines.
	provisionedStorageMb	The storage provisioned for the virtual machine in the recovery site.
	throughputInMB	MBs of all the applications running on the virtual machine being protected.
	usedJournalStorageMb	Storage used by the virtual machine at the recovery site for Journals.
	usedStorageMb	Storage used by the virtual machine at the recovery site.
	vpgName	The name of the VPG, of which the VM is associated.
	vmIdentifier	ID of the VM.
	vmName	Name of the VM.
VRAs		
	vraName	Name of the VRA.
	vraVersion	Version of the VRA.
Events		
	eventIdentifier	The identifier of the event.
	eventCategory	Category of the event. (Possible values: VPG, VRA, Site).
	eventType	The type of event.
	siteIdentifier	The internal site identifier related to the event.
	zorgIdentifier	The identifier of the Zerto organization (ZORG), defined in the Zerto Cloud Manager.
	zorgName	The name of the Zerto organization (ZORG) set up in the Zerto Cloud Manager.
	description	The event description.
	occurredOnInUTC	The date the event occurred.
	eventCode	The event code.
	eventCompletedSuccessfully	Whether the event completed successfully or not.
ZORGs		
	zorgIdentifier	The identifier of the Zerto organization (ZORG), defined in the Zerto Cloud Manager.
	zorgName	The name of the Zerto organization (ZORG) set up in the Zerto Cloud Manager.

Datastores		
	datastoreIdentifier	The datastore identifier.
	datastoreName	The datastore name.
	devices	The datastore device.
	type	The datastore type.
	capacityInBytes	The datastore capacity.
	usedInBytes	The datastore usage. (Possible values: journal usage; scratch usage; protected usage; recovery usage; appliance usage; other non-Zerto usage)
Volumes		
	datastore	The owning datastore.
	isThisProvisioned	Provision type.
	owningVm	The virtual machine on which the volume is attached.
	path	The volume path.
	protectedVm	The protected virtual machine to which the volume is related.
	size	The volumes used storage.
	volumeType	The volume type.
	vpg	The VPG to which the volume is related.
Planning Metrics are collected for all VMs from the source site.		
	WriteKBps - Hourly Avg.	The hourly average data write rate.
	WriteIOPS - Hourly Avg.	The hourly average IO count per second.
	VM ID	The virtual machine identifier.
	Datacenter Name	The datacenter name.
	Datacenter ID	The datacenter identifier.
	Host Name	The host name.
	Host ID	The host identifier.
	Cluster ID	The cluster identifier.
	Cluster Name	The cluster name.
	Folder	The VM folder.
	Configured CPU	The configured CPU.
	Configured RAM	The configured RAM.
	Provisioned Size (GB)	The provisioned storage .
	Used Size (GB)	The used storage.
	Disk Count	Number of virtual disks.

Zerto Cloud Control - Overview

The Zerto Cloud Control includes the following features:

- Upgrade Manager
- Tenants

See also:

- "Before Getting Started with Zerto Cloud Control", on page 22
- "Accessing the Zerto Cloud Control", on page 22
- "Using Upgrade Manager", on page 22
- "Using Tenants", on page 22

Before Getting Started with Zerto Cloud Control

Verify the following:

- **Enable Zerto SaaS features** check box is selected. This is accessed in the ZVM application in **Settings > About.**
- Internet access.
- A **myZerto** account using your corporate email address.

Accessing the Zerto Cloud Control

Zerto Cloud Control is accessed from one of the following locations:

- From a URL: https://cloudcontrol.zerto.com
- From myZerto: www.zerto.com/myzerto. Sign in using your credentials and select the **Cloud Control** tab.

Using Upgrade Manager

The Upgrade Manager feature allows Cloud Service Providers to remotely upgrade their end customer sites. This allows CSPs to keep their customers Zerto software always up to date. It also allows Cloud Service Providers an overall view of all their end customers' ZVR version.

The remote upgrade functionality has the following pre-requisites:

- The customers ZVR instance is paired to a CSP-deployed ZCC and does not have a perpetual license applied.
- The end customer Zerto instance should be able to transmit over port 443.

Use the **Help** widgets in Upgrade Manager to learn more about each of the features available in the Upgrade Manager.

Using Tenants

The Tenants feature allows Cloud Service Providers to register new and existing enterprise customers to the respective ZORG for that customer. In the ZORG registration, the CSP selects the Product Type. The product type can include features such as One-To-Many and Long Term Retention.

Use the **Help** widgets in Tenants to learn more about each of the features available in the Upgrade Manager.

Benefits of Using Zerto

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. The Zerto solution has been designed to resolve these issues by being fully virtualization aware.

See the following topics:

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

Hardware Agnostic

Because the Zerto solution 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 the Zerto solution, 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 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, 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, retention, testing, and development. With Zerto 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 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 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 supports the hypervisor versions, virtual machines can be protected across versions.

Efficient Asynchronous Replication

Writes are captured by the Zerto 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 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 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.

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

The Zerto solution 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.

The Zerto solution 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 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

The Zerto solution 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 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 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 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 the Zerto solution 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

  Note: The ImportImage API is not used by Zerto.
- **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.
  - F535Zerto can set default encryption on the S3 bucket so that all objects are encrypted when they are stored in the bucket. To enable S3 encryption please contact support.
  - The default zImporter instance type is c5.4xlarge and the AWS EC2 default maximum instance quota is 10. 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.
    - Each zImporter VM is responsible for the import process of a single volume. Therefore, it is recommended to contact AWS and increase the maximum instance quota of the c5.4xlarge instance type to the maximum number of volumes you are planning to failover to AWS at once.
  - 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 26
  - "Zerto Import for All Volumes", on page 27
  - "AWS Import", on page 28

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.
- 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 Zerto Import for All Volumes 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 ENA (Elastic Network Adapter) Drivers

When recovering to **C5/M5** instances using Zerto Import for All Volumes import method, only **Windows 2012R2** and **Windows 2016** are supported. The following drivers must be installed on the protected virtual machine before starting recovery operations **when recovering to C5/M5 instances**:

- Windows ENA (Elastic Network Adapter) Drivers
- NVMe driver

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.

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 ENA Drivers:**

NOTE:
If you are recovering to **C5/M5 instances**, skip this step and go to step 2.

   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.
2. When using **C5/M5 instance types only**, the following drivers need to installed:
   a) Download ENA (Enhanced Network Adapter) from this link.
   b) Extract and execute "InstallEna.ps1"
   c) Download the NVMe driver from this link.

d) Extract and execute "Dpinst.exe"

NOTES:

C5/M5 instance types are supported with the Zerto Import for All Volumes import method only.

IMPORTANT:

When using this import method for Windows machines, **ZertoTools** for AWS needs to be run on the **protected** Windows virtual machine in VMware **before VPG creation.** For more information, see ZertoTools for AWS.

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

How the Zerto Solution for AWS Works

Zerto provides a business continuity (BC) and disaster recovery (DR) solution by enabling failing over to AWS and failing back from AWS. The solution enables failback and migration from AWS to a hypervisor based datacenter or to the cloud, allowing protection of native AWS workloads and the ability to migrate that workload out of AWS.

Using Zerto scale-out technology, data is protected in AWS using Zerto Satellite (**zSATS**) for reading data in order to replicate the data to the recovery site.

To failback from AWS, a snapshot is created in the Amazon S3 bucket for each EBS disk. A new EBS disk is then created from the snapshot and attached to the zSAT EC2 instance. The zSAT EC2 instances are automatically created and deleted by the ZCA.

The following Zerto components enable the protection of workloads in AWS:

- **Zerto AWS Snapshot Adapter (zASA):** This is a snapshots lifecycle manager that returns the entire disk for syncing. The zASA is an EC2 instance that is created once the system detects there is a VPG for protected workloads in AWS. The zASA instance remains up and running as long as a VPG exists and the ZCA is installed.
- **Zerto Satellite (zSATS):** Scale-out solution with EC2 instances for reading data on protected EBS disks. The zSAT is an EC2 instance for reading data from the EBS disk that is created from the snapshot of the protected EBS disk. Once that disk is read, the snapshot, EBS disk and zSAT instance from the previous sync are deleted.

Tagging of Zerto AWS Resources

F533Below is a table showing the tag values assigned to Zerto AWS resources:

ZERTO AWS RESOURCE	ZERTO_TAG	OWNER
ZImporter Instance	ZERTO_EC2_RESOURCE	ZCA Site ID
ZImporter Volume	ZERTO_EBS_RESOURCE	ZCA Site ID
ZImporter NIC	ZERTO_EC2_RESOURCE	ZCA Site ID
ZImporter Security Group	ZERTO_VPC_RESOURCE	ZCA Site ID
ZASA Instance	ZERTO_EBS_RESOURCE	ZCA Site ID
ZASA Volume	ZERTO_EBS_RESOURCE	ZCA Site ID
ZASA NIC	ZERTO_EC2_RESOURCE	ZCA Site ID
ZASA Security Group	ZERTO_VPC_RESOURCE	ZCA Site ID
ZSAT Instance	ZERTO_EC2_RESOURCE	ZCA Site ID
ZSAT Volume	ZERTO_EBS_RESOURCE	ZCA Site ID
ZSAT NIC	ZERTO_EC2_RESOURCE	ZCA Site ID
Recovered VM Instance	ZERTO_EC2_RESOURCE	ZCA Site ID
Recovered VM Volume	ZERTO_EBS_RESOURCE	ZCA Site ID
Recovered VM NIC	ZERTO_EC2_RESOURCE	ZCA Site ID
S3 Bucket	ZERTO_S3_RESOURCE	ZCA Site ID

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:
- For supported browsers, see Zerto Interoperability Matrix.
- The **lowest** supported screen resolution is **1366x768**.

Note: You must **exclude** the following folders from **antivirus scanning**:

Zerto Virtual Replication
%ProgramData%\Zerto\Data\zvm_db.mdf
C:\Program Files\Zerto\Zerto Virtual Replication\Zerto.Zvm.Service.exe
C:\Program Files\Zerto\Zerto Virtual Replication\Zerto.Vba.VbaService.exe
C:\Program Files\Zerto\Zerto Virtual Replication\Zerto Online Services Connector\Zerto.Online.Services.Connector.exe
C:\Program Files\Zerto\Zerto Virtual Replication\Embedded DB Manager Service\Zerto.LocalDbInstanceManagerService.exe

Failure to do so may lead to the Zerto Virtual Replication folder being incorrectly identified as a threat and in some circumstances corrupt the Zerto Virtual Replication folder.

The following topics are described in this section:

- "Using the Zerto Virtual Manager Web Client", below
- "Adding a Security Certificate for the Zerto User Interface", on page 30
- "Working With the Zerto User Interface" on page 32

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**.

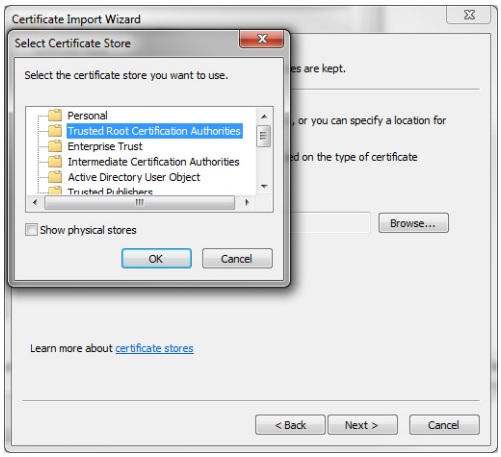

4. Continue to the end of the wizard. Click **Yes** when the Security Warning is displayed.

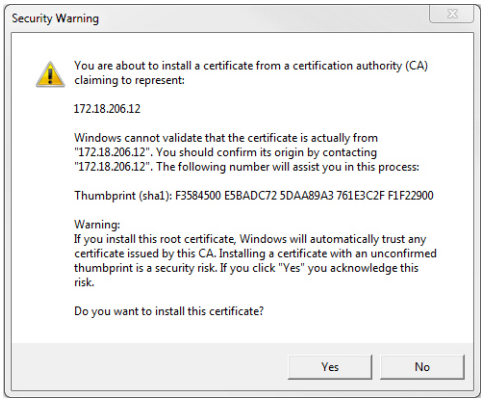

5. Click **OK** that the installation was successful.
6. Click **OK** when prompted and then **Yes** in the **Security Alert** dialog to continue.

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:

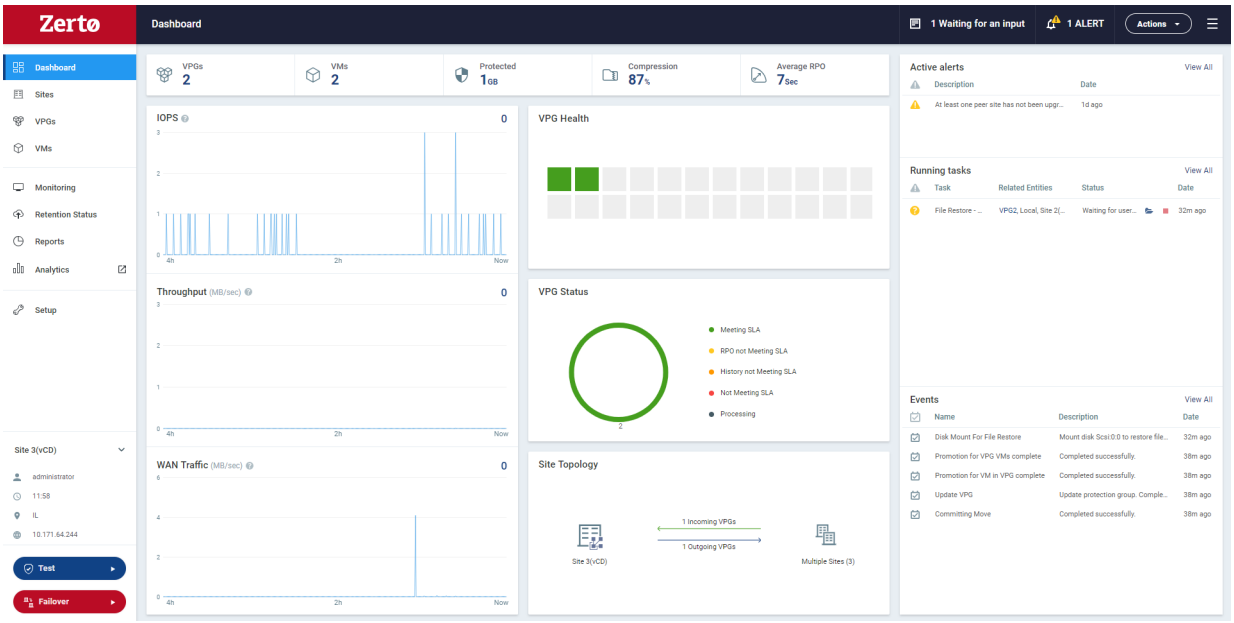

Use the tabs to access the specific information you want:

- **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:** *(On-Premise only)* Details about VRAs, storage and repositories.
- **RETENTION STATUS:** Details of the retention repository jobs either by VPG or virtual machine. This tab lists all the defined retention sets and their statuses.
- **MONITORING:** Details about the alerts, events and tasks for the site.
- **REPORTS:** General reports.

Zerto User Interface - Subtabs

From the SETUP, RETENTION STATUS and MONITORING tabs you can see and manage details for a specific VPG and VRA from different perspective.

For example, under SETUP, from the subtabs, you can manage VRAs, Storage/Datastores and Repositories.

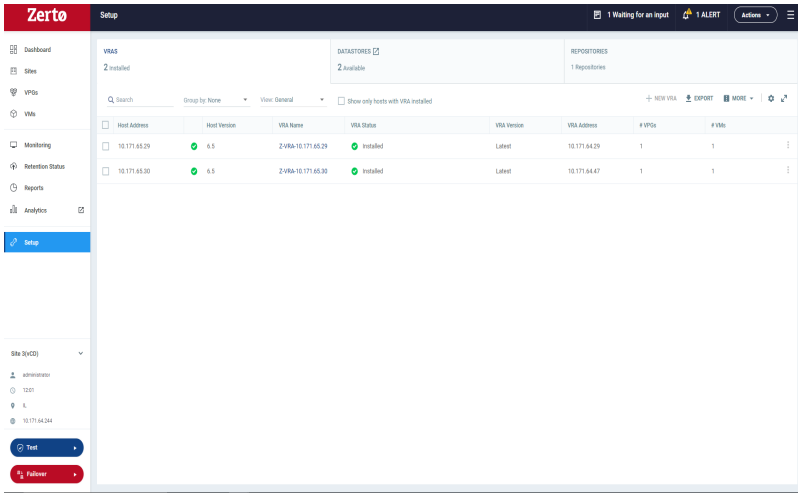

Zerto User Interface - Views

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: 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, a retention set 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.

For the maximum number of virtual machines, either being protected or recovered to a site, see Zerto Scale and Benchmarking Guidelines

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 section:

- "Configuring Virtual Protection Groups", below
- "The Role of the Journal During Protection", on page 35
- "What Happens After the VPG is Defined", on page 35

Configuring Virtual Protection Groups

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

For the maximum number of virtual machines, either being protected or recovered to a site, see Zerto Scale and Benchmarking Guidelines

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

VPGs can be created on the on-premise site or in AWS. The virtual machines on the on-premise 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.
- **Retention Status:** The properties that govern the VPG retention including the repository where the retention sets are saved.
- **Summary:** The details of the VPG configuration defined in the previous components.

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 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.

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.

CHAPTER 4: OVERVIEW OF RECOVERY FLOWS

The Zerto solution enables protecting virtual machines, for both disaster recovery or for extended, longer term recovery from a retention repository, 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.

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.

In addition to disaster recovery and recovery from retention, the Zerto solution enables recovery of individual files or folders from a certain point of time.

The following are described in this section:

- "What is Zerto's Disaster Recovery Operation?", below
- "What is Zerto's File or Folder Level Restore?", on page 38

What is Zerto's Disaster Recovery Operation?

Disaster recovery using the Zerto solution 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.

What is Zerto's Disaster Recovery Operation to AWS?

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 a retention set.

Virtual machines are protected in VPGs, which are defined in the protected site. Once a VPG is created, Zerto 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 35 for more details about the journal.

The following link references the appropriate procedure to protect virtual machines:

Recovery from: **Amazon Web Services (AWS)**	Protecting Virtual Machines From AWS

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

- "Overview of Disaster Recovery Operations", on page 137
- "Managing Failover Live to AWS", on page 168
- "Migrating a VPG to AWS", on page 157
- "Cloning a VPG to AWS", on page 176

What is Zerto's Test Failover Operation in AWS Environments?

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 137
- "Testing Recovery to AWS", on page 149

What is Zerto's File or Folder Level Restore?

You can recover specific files and folders from the recovery site for virtual machines that are being protected by Zerto 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 179.

CHAPTER 5: PROTECTING VIRTUAL MACHINES FROM AWS

Virtual Machines created in AWS can be protected to the following platforms:

PROTECTED SITE	TO RECOVERY SITE	SEE PROCEDURE...
AWS	To a **vCenter Server** site	"Replication From a Protected Site AWS to a Recovery Site vCenter Server", on page 40
	To **vCloud Director**	"Replication From a Protected Site AWS to a Recovery Site vCloud Director", on page 59
	To a **Hyper-V** site	"Replication From a Protected Site AWS to a Recovery Site Hyper-V", on page 76
	To a **Microsoft Azure** site	"Replication From a Protected Site AWS to a Recovery Site Azure", on page 94

Requirements for Amazon Web Services (AWS) Environments

Before protecting your virtual machines, review Zerto Virtual Replication - Prerequisites & Requirements for Amazon Web Services (AWS)

Retained VM Properties

For settings that are **retained** when recovering from AWS, see the following:

- When recovering to a **vCenter Server** site: Maintained VM Properties when Recovering to a vCenter Server Site from a Protected AWS Site

Replication From a Protected Site AWS to a Recovery Site vCenter Server

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

When creating a VPG in an AWS site, to recover to a vCenter Server site, the recovery operations bring up the recovered machines on VMware vCenter Server hosts.

When recovering to a vCenter Server site, the following is required:

- The operating systems of the protected machines must be supported by vCenter Server. Refer to the VMware documentation for a list of supported operating systems.

Zerto uses SCSI for vCenter Server virtual machine disks.

When protecting virtual machines from AWS to vCenter Server, the operating systems of the protected machines must be supported by vCenter Server. Refer to the VMware documentation for a list of supported operating systems.

To review **which settings are retained**, see Maintained VM Properties when Recovering to a vCenter Server Site from a Protected AWS Site, on page 58.

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

To create a virtual protection group (VPG) to recover in a vCenter Server:

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

   The GENERAL step of the Create VPG wizard is displayed.

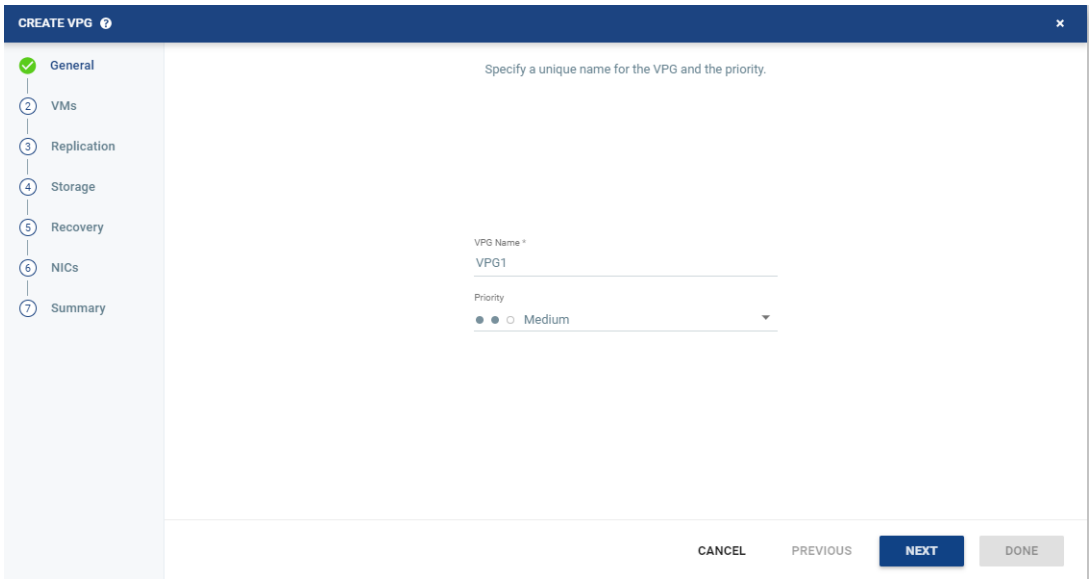

2. Specify the name of the VPG and the priority of the VPG.
   - **VPG Name:** The VPG name must be unique. The name cannot be more than 80 characters.
   - **Priority:** Determine the priority for transferring data from the protected site to the recovery site when there is limited bandwidth and more than one VPG is defined on the protected site.
     - **High Priority**: When there are updates to virtual machines protected in VPGs with different priorities, updates from the VPG with the highest priority are passed over the WAN first.
     - **Medium Priority:** Medium priority VPGs will only be able to use whatever bandwidth is left after the high priority VPGs have used it.
     - **Low Priority:** Low priority VPGs will use whatever bandwidth is left after the medium priority VPGs have use it.

   Updates to the protected virtual machines are always sent across the WAN before synchronization data, such as during a bitmap or delta sync.

During synchronization, data from the VPG with the highest priority is passed over the WAN before data from medium and low priority VPGs.

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

3. Click **NEXT**.

   The VMs step is displayed.

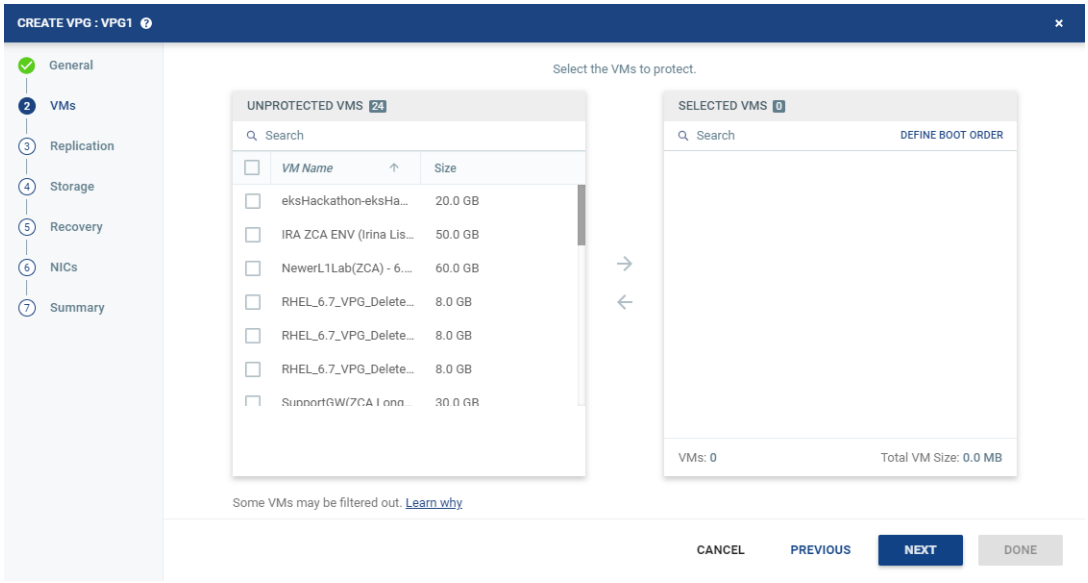

- Unprotected VMs with all its disks that were in the ZCA's **account and region** during installation, are included in the list and can be selected for protection.
- A VPG can include virtual machines that are already protected by another ZCA in the same account and region.
- The size of the disk is defined by data disks and the OS disk.
- AWS Instance Store disks (Temp disks) are not included in the list of unprotected VMs.

4. From the **Unprotected VMs** list, select the VMs to include in this VPG and click the arrow to move them to the **Selected VMs**.
   - When using the **Search** field, you can use the wildcards; ***** or **?**
5. To define the boot order of the virtual machines in the VPG, click **Define Boot Order**, otherwise go to the next step.

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

   Note: Up to five (5) virtual machines may boot on a host simultaneously. Following the boot, a 300 second (default) delay occurs until the next boot batch.

   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.

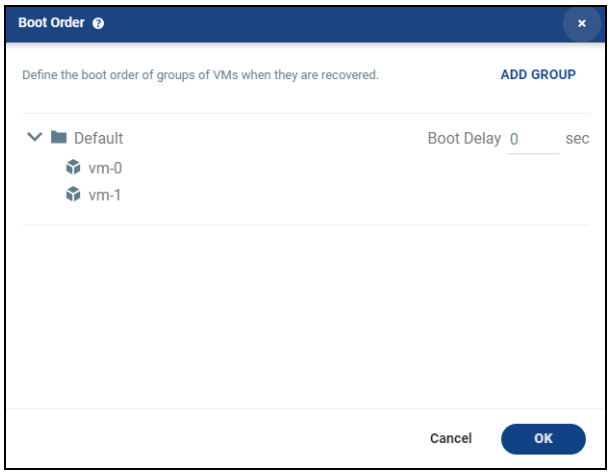

6. Click **ADD GROUP** to add a new group. Then, do the following:
   a) To change the name of a group, click the Pencil icon next to the group.
   b) To delete a group, click the delete icon on the right side. You cannot delete the Default group nor a group that contains a virtual machine.
   c) Drag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

7. Click **NEXT**.
   The **REPLICATION** step is displayed.

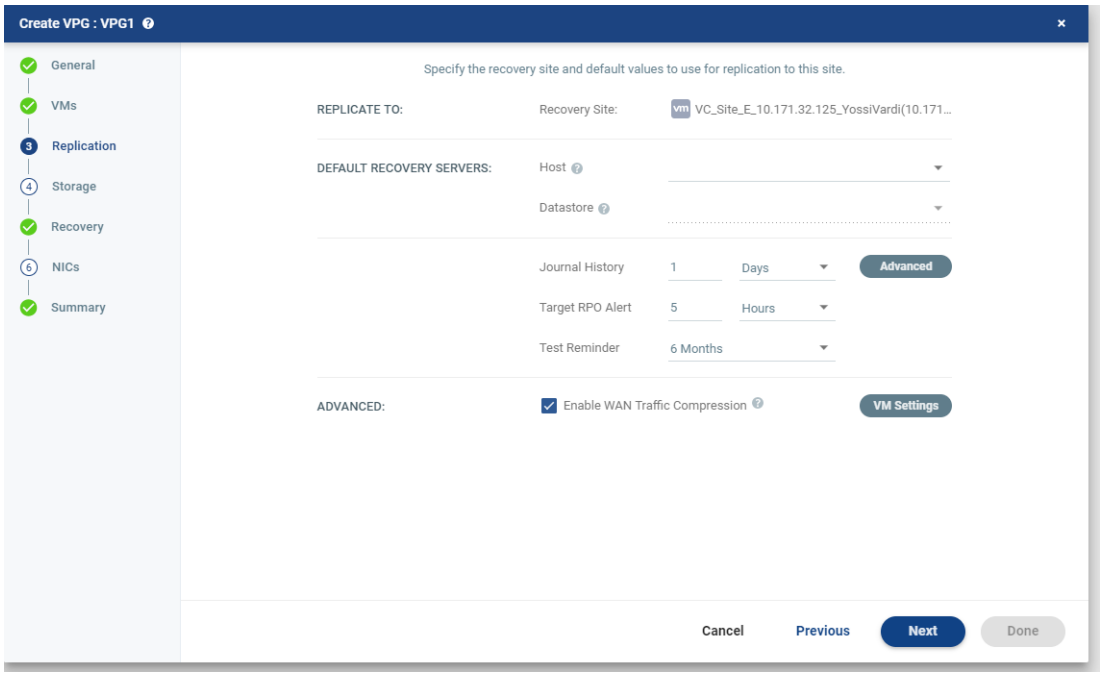

8. Specify the **Recovery Site**. This is the site to which you want to recover the virtual machines.
   - If the protected site is paired with **more than one** recovery site, the Recovery Site needs to be selected. Proceed to step 9.
   - If the protected site is paired with **only one** recovery site, and that recovery site is defined in **Zerto Cloud Manager**, the Recovery Site field is automatically filled in, as well as the defaults set for the SLA and Advanced settings. Proceed to step 10.
   - If the protected site is paired with **only one** recovery site, the Recovery Site field is automatically filled in, as well as the defaults set for the SLA and Advanced settings. Proceed to step 11.
9. Select the **Recovery Site** from the drop-down list.

   After specifying the Recovery Site, other fields are displayed including the Host and Storage to use for replication.

   If the selected virtual machines are **already in VPGs** that recover to that site, select another recovery site.
10. (**If the site is defined in** Zerto Cloud Manager) **ZORG:** Specify the name used by the **cloud service provider** to identify you as a Zerto Organization, ZORG.

    If the site is defined in Zerto Cloud Manager, vSphere Standard edition **cannot** be used.

    For details about Zerto Cloud Manager, refer to *Zerto Cloud Manager Administration Guide*.
11. After selecting the Recovery Site, define the following parameters:

PARAMETER	FIELD	DESCRIPTION
Default Recovery Services	**Host**	The default **cluster**, **resource pool** or **host** in the recovery site that handles the replicated data. ■ If the site is defined in **Zerto Cloud Manager**, only a resource pool can be specified and the resource pool must also have been specified as a resource in Zerto Cloud Manager. ■ Zerto checks that the resource pool capacity is enough for any virtual machine specified in the VPG. ■ All resource pool checks are made at the VPG level and do not take into account multiple VPGs using the same resource pool. ■ If the resource pool CPU resources are specified as unlimited, the actual limit is inherited from the parent. If the inherited value is too small, failover, move, and failover test operations can fail, even without a warning alert being issued by Zerto Virtual Manager. ■ If a resource pool is specified and DRS is disabled for the site later on, all the resource pools are removed by VMware and recovery will be to any one of the hosts in the recovery site with a VRA installed on it.
	Datastore	The default datastore to use for recovered virtual machine files and for their data volumes. ■ Every datastore for the selected recovery host is included in the drop-down list. ■ If a cluster or resource pool is selected for the host, only datastores that are accessible by every host in the cluster or resource pool are displayed.

PARAMETER	FIELD	DESCRIPTION
SLA	**Service Profile**	(Only if the site is defined in **Zerto Cloud Manager**) Specify the name of the service profile. The service profile determines the VPG SLA settings for the group, and is applied to every virtual machine in the group. To change the VPG SLA settings, select the **Custom** Service Profile. To change VPG SLA for individual virtual machines, proceed to step 15.
	Journal History	Specify these settings for the group, which apply to every virtual machine in the group. The longer the information is saved in the journal, the more space is required for each journal in the VPG. Journal history can be defined either in hours or in days. **Select** one of the following: ■ Number of **hours** from 1 to 24, ■ Number of **days** from 2 to 30. For **additional fields**, click **ADVANCED**, and proceed to step 12.
	Target RPO Alert	The maximum desired time between each automatic checkpoint write to the journal before an alert is issued. Default value is X hours.
	Test Reminder	The time recommended between testing the integrity of the VPG. A warning is issued if a test is not done within this time frame.
Advanced	**VM SETTINGS**	To change the replication settings **per virtual machine**, click **VM SETTINGS** button, and proceed to step 15.

12. For additional journal-related fields, click **ADVANCED**.
    The Advanced Journal Settings window is displayed.

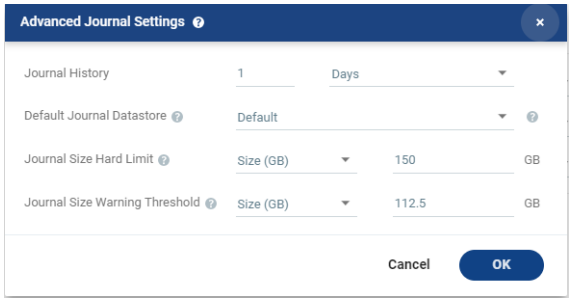

13. Select the journal settings.

SETTING & DESCRIPTION	SELECT...
Journal History	
The time that all write commands are saved in the journal. The longer the information is saved in the journal, the more space is required for each journal in the VPG.	■ Number of **hours** from **1 to 23** ■ Number of **days** from **1 to 30**

SETTING & DESCRIPTION	SELECT...
Default Journal Storage (Hyper-V), or **Default Journal Datastore** (vSphere)	
The storage/datastore used for the journal data for each virtual machine in the VPG. **Note:** This field is **not** relevant when replicating to a **vCD** recovery site.	■ Select the storage/datastore accessible to the host. When you select a specific journal storage/datastore, the journals for each virtual machine in the VPG are stored in this storage/datastore, regardless of where the recovery storage/datastore is for each virtual machine. All protected virtual machines are recovered to the hosts that can access the specified journal storage/datastore.
Journal Size Hard Limit	

The maximum size that the journal can grow, either as a percentage or a fixed amount.
 	
The journal is always **thin-provisioned**.
 	
Note: The Journal Size Hard Limit applies independently **both** to the Journal History **and** also to the Scratch Journal Volume.
 	
For Example: If the Journal Size Hard Limit is configured to a maximum size of 160 GB limit, then during Failover Test, both the Journal History and the Scratch Journal Volume together can take up to 320 GB. Each one with a maximum size of 160 GB limit.
	■ **Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore.
	
	If Unlimited is selected, Size and Percentage options are **not** displayed.
	
	■ **Size (GB):** The maximum journal size in GB.
	  ■ The **minimum** journal size, set by Zerto, is **8GB** for Hyper-V and vSphere environments, and **10GB** for Microsoft Azure environments.
	
	■ **Percentage:** The percentage of the virtual machine volume size to which the journal can grow.
	  ■ This value can be configured to more than 100% of the protected VM's volume size.

Journal Size Warning Threshold	
The size of the journal that triggers a warning that the journal is nearing its hard limit.	■ **Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore. If Unlimited is selected, Size and Percentage options are **not** displayed. ■ **Size* (GB):** The size in GB that will generate a warning. ■ **Percentage*:** The percentage of the virtual machine volume size that will generate a warning. ***The values of **Size** and **Percentage** must be **less** than the configured **Journal Size Hard Limit** so that the warning will be generated when needed. In addition to the warning threshold, Zerto will issue a message when the free space available for the journal is almost full.

14. Click **OK**.

15. To change the replication settings per virtual machine, click **VM SETTINGS**.

The Advanced VM Replication Settings window is displayed.

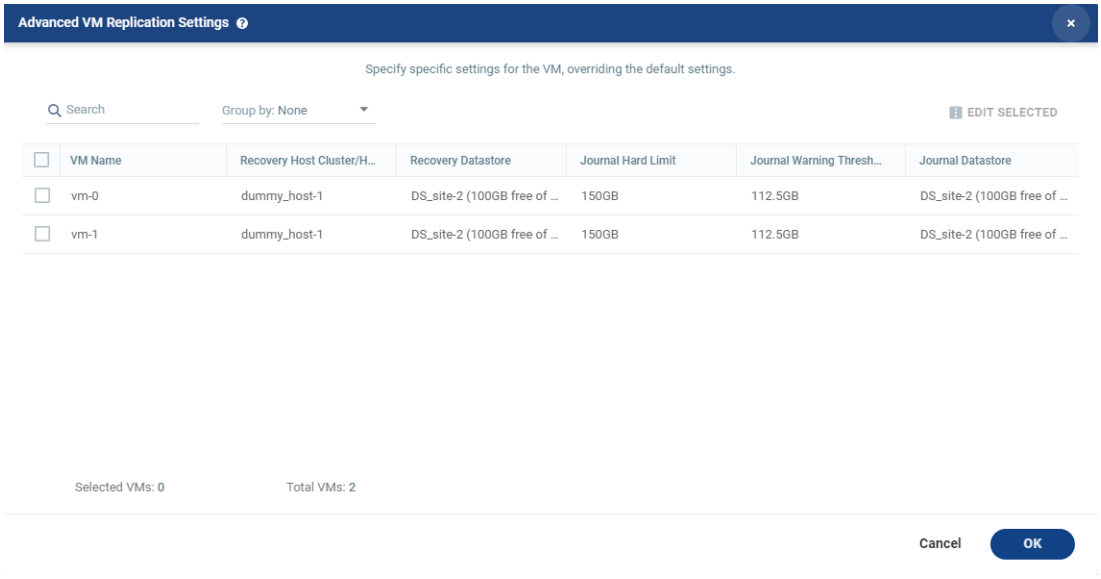

In this window, you can edit the values of one or more of the virtual machines in the VPG.

16. To edit information in **one field**, click the field and update the information.
17. To edit information for **several virtual machines** at the same time, select the virtual machines and click **EDIT SELECTED**.
    The **Edit VM** window is displayed.

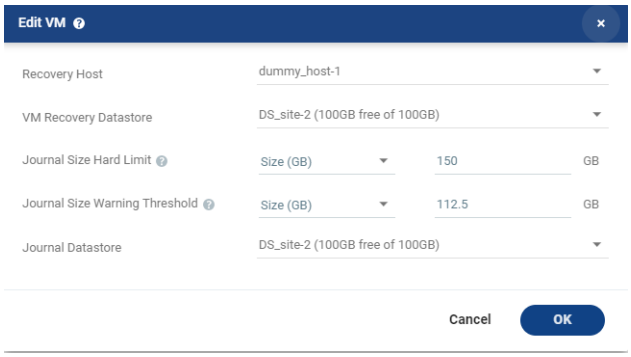

Define as follows:

SETTING & DESCRIPTION	SELECT...
Recovery Host (not relevant when replicating to vCD)	
(Hyper-V) The cluster or host that will host the recovered virtual machine.	

SETTING & DESCRIPTION	SELECT...
(vSphere) The cluster, resource pool, or host that will host the recovered virtual machine. If the site is defined in Zerto Cloud Manager, only a resource pool can be specified and the resource pool must also have been defined in Zerto Cloud Manager. For details about Zerto Cloud Manager, see *Zerto Cloud Manager Administration Guide*. When a resource pool is specified, Zerto checks that the resource pool capacity is enough for all the virtual machines specified in the VPG	When a resource pool is specified, Zerto checks that the resource pool capacity is enough for all the virtual machines specified in the VPG. If a resource pool is specified and DRS is disabled for the site later on, all the resource pools are removed by VMware and recovery is to any one of the hosts in the recovery site with a VRA installed on it. All resource pool checks are made at the level of the VPG and do not take into account multiple VPGs using the same resource pool. If the resource pool CPU resources are defined as unlimited, the actual limit is inherited from the parent but if this inherited value is too small, failover, move, and failover test operations can fail, even without a warning alert being issued by Zerto Virtual Manager.
VM Recovery Datastore (vSphere) (not relevant when replicating to vCD)	
The datastore where the VMware metadata files for the virtual machine are stored, such as the VMX file.	If a cluster or resource pool is selected for the host, only datastores that are accessible by every ESX/ESXi host in the cluster or resource pool are displayed. This is also the datastore where RDM backing files for recovery volumes are located.
Recovery Storage (Hyper-V) (not relevant when replicating to vCD)	
The location where the metadata files for the virtual machine are stored, such as the VHDX file. I	If a cluster is selected for the host, only storage that are accessible by every host in the cluster are displayed.
Journal Size Hard Limit	
The maximum size that the journal can grow, either as a percentage or a fixed amount. ■ The journal is always **thin-provisioned**. ■ The Journal Size Hard Limit applies independently **both** to the Journal History **and** also to the Scratch Journal Volume. *For Example:* If the Journal Size Hard Limit is configured to a maximum size of 160 GB limit, then during Failover Test, both the Journal History and the Scratch Journal Volume together can take up to 320 GB. Each one with a maximum size of 160 GB limit.	**Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore. If Unlimited is selected, Size and Percentage options are **not** displayed. **Size (GB):** The maximum journal size in GB. ■ The **minimum** journal size, set by Zerto, is **8GB** for Hyper-V and vSphere environments, and **10GB** for Microsoft Azure environments. **Percentage:** The percentage of the virtual machine volume size to which the journal can grow. ■ This value can be configured to more than 100% of the protected VM's volume size.

SETTING & DESCRIPTION	SELECT...
Journal Size Warning Threshold	
The size of the journal that triggers a warning that the journal is nearing its hard limit.	**Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore. If Unlimited is selected, Size and Percentage options are **not** displayed.
	Size* (GB): The size in GB that will generate a warning.
	Percentage*: The percentage of the virtual machine volume size that will generate a warning.
	***The values of **Size** and **Percentage** must be **less** than the configured **Journal Size Hard Limit** so that the warning will be generated when needed. In addition to the warning threshold, Zerto will issue a message when the free space available for the journal is almost full.
Journal Storage (Hyper-V), or **Journal Datastore** (vSphere) (not relevant when replicating to vCD)	
The storage/datastore used for the journal data for each virtual machine in the VPG.	**(vSphere)** To change the default, specify a host and then select one of the datastores accessible by this host to be used as the journal datastore. When you select specific journal datastore, the journals for each virtual machine in the VPG are stored in this datastore, regardless of where the recovery datastores are for each virtual machine. In this case, all the protected virtual machines must be recovered to hosts that can access the specified journal datastore. **(Hyper-V)** To change the default, specify a host and then select the storage location accessible by this host to be used as the journal storage. When you select specific journal storage, the journals for each virtual machine in the VPG are stored in this storage, regardless of where the recovery storage is for each virtual machine. In this case, all the protected virtual machines must be recovered to hosts that can access the specified journal storage.

18. Click **OK**.
19. In the Advanced VM Replication Settings window, click **OK**.
20. Click **NEXT**.

    The **Storage** step is displayed.

    By default the storage used for the virtual machine definition is also used for the virtual machine data.

    For each virtual machine in the VPG, Zerto displays its storage-related information.

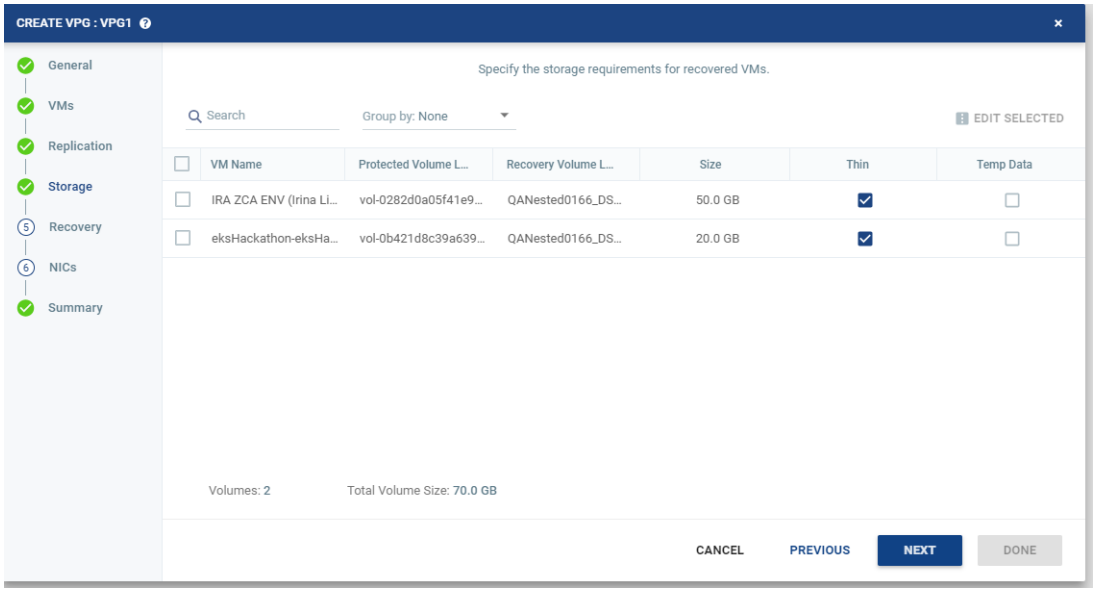

You can define **Thin** provisioning and **Temp Data** in this window, **or** you can alternatively define them when you separately select and **edit each VMs volume**.

IMPORTANT:

Changing the VPG recovery volume from thin-provisioned to thick-provisioned or vice versa, results in **volume initial synchronization**.

See the following considerations regarding Thin provisioning:

- Unless the user **explicitly** requests Thin provisioning, provisioning type is the same type as provisioning in the **source VM**.
- If the **source** disk is Thin provisioned, the default for the **recovery** volume is also Thin provisioned.
- If the user uses preseed disks, Zerto maintains the provisioning types of the disks, so they can have other provisioning types.

PRESEED	PROVISIONING IN THE RECOVERY VM
Not selected	■ User **can** select Thin provisioning
Selected	■ User **cannot** select Thin provisioning ■ Provisioning is the **same** as defined in **source VMs**

21. To define whether the recovery volumes are thin-provisioned or not, select the **Thin** checkbox.
22. If the virtual machine to be replicated includes a temp data disk as part of its configuration, select the **Temp Data** checkbox to mark the recovery disk for this disk as a temp data disk. In this case, data is not replicated to the temp data disk after initial synchronization.

23. To edit storage information for one of the virtual machines' volume location, first select the virtual machine, then click **EDIT SELECTED**. The Edit Volumes window is displayed.

    - For **Hyper-V** recovery environments, the following window appears. For details, click here.

- For **vSphere** recovery environments, the following window appears. For details, click here.

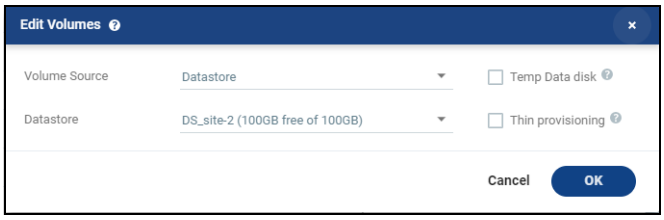

- For **vCD** recovery environments, the following window appears. For details, click here.

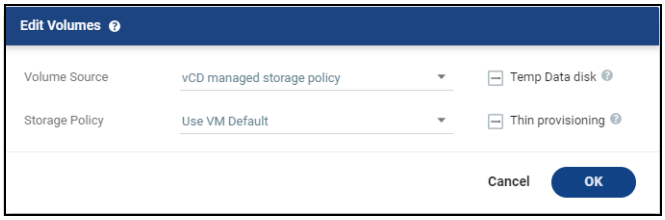

SETTING & DESCRIPTION	SELECT...
Volume Source	
■ (Hyper-V) Select a **Volume Source** for recovery from one of the drop-down options: ■ **Storage** ■ **Preseeded volume**	■ **Volume Source > Storage:** A new volume is used for replicated data. ■ From the **Storage** drop-down list, specify the storage to use to create disks for the replicated data. The storage specified for the replication must have at least the same amount of space as the protected volume and then an additional amount for the journal. The amount of additional space needed for the journal can be fixed by specifying a maximum size for the journal, or can be calculated as the average change rate for the virtual machines in the VPG, multiplied by the length of time specified for the journal history. For more details, see *Zerto Scale and Benchmarking Guidelines*, in the section **Estimating WAN Bandwidth for VMware and Hyper-V**.
	■ **Volume Source > Preseeded volume:** Whether to copy the protected data to a virtual disk in the recovery site. Zerto **recommends** using this option particularly for large disks so that the initial synchronization will be faster since a **Delta Sync** can be used to synchronize any changes written to the recovery site after the creation of the preseeded disk. When **not** using a preseeded disk, the initial synchronization phase must copy the whole disk over the WAN. When using a preseeded virtual disk, you select the storage and exact location, folder, and name of the preseeded disk. Zerto takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. Only disks with the same size as the protected disk can be selected when browsing for a preseeded disk. The storage where the preseeded disk is placed is also used as the recovery storage for the replicated data.

SETTING & DESCRIPTION	SELECT...
■ (vSphere) Select a **Volume Source** for recovery from one of the drop-down options: ■ **Datastore** ■ **RDM** ■ **Preseeded volume**	**Volume Source > Datastore:** A new volume is used for replicated data. 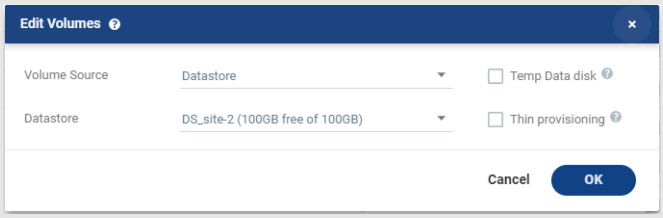 ■ Specify the **Datastore** to use to create disks for the replicated data. ■ If the **source disk is thin provisioned**, the default for the recovery volume is also thin provisioned. ■ The datastore specified for replication must have at least the same amount of space as the **protected volume** and an additional amount for the **journal**. ■ The amount of additional space needed for the journal can be fixed by specifying a maximum size for the journal, or can be calculated as the average change rate for the virtual machines in the VPG, multiplied by the length of time specified for the journal history. ■ Zerto supports the SCSI protocol. Only disks that support this protocol can be specified. Then, define the following: ■ **Datastore:** The Datastore where the preseeded disk is located. Only disks with the same size as the protected disk can be selected when browsing for a preseeded disk. For more details, see *Zerto Scale and Benchmarking Guidelines*, in the section **Estimating WAN Bandwidth for VMware and Hyper-V**.
	Volume Source > RDM: The VMware RDM (Raw Device Mapping) which will be used for the replication. 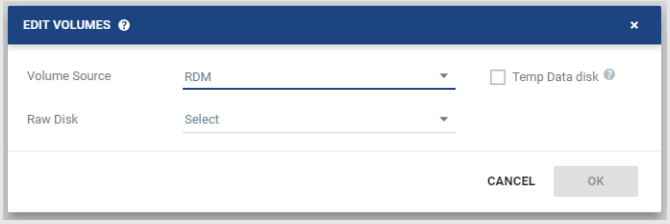 By default, **RDM is recovered as thin-provisioned VMDK** in the datastore specified in the **VM Recovery Datastore/Storage** field in the Edit VM dialog, and not to RDM. Only a raw disk with the **same size as the protected disk** can be selected from the list of available raw disks. Other raw disks with different sizes are not available for selection. The RDM is always stored in the recovery datastore, used for the virtual machine. The following **limitations** apply to protecting RDM disks: ■ RDM disks with an even number of blocks can replicate to RDM disks of the same size with an even number of blocks and to VMDKs. ■ RDM disks with an odd number of blocks can only replicate to RDM disks of the same size with an odd number of blocks and not to VMDKs. ■ You cannot define an RDM disk to be protected to a cloud service provider via a Zerto Cloud Connector nor if the virtual machine uses a BusLogic SCSI controller, nor when protecting or recovering virtual machines in an environment running vCenter Server 5.x with ESX/ESXi version 4.1 hosts.

SETTING & DESCRIPTION	SELECT...
■ (vSphere) **Volume Source** continued	**Volume Source > Preseeded volume**: Select this when you want to copy the protected data to a virtual disk in the recovery site. 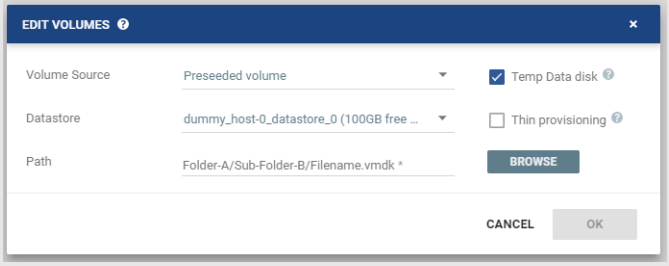 Consider the following, then proceed to define the Datastore and the Path: ■ Zerto **recommends** using this option particularly for **large disks** so that the initial synchronization is **faster** since a Delta Sync can be used to synchronize any changes written to the recovery site after the creation of the preseeded disk. ■ If a preseeded disk is **not** selected, the initial synchronization phase must copy the **whole disk** over the WAN. ■ If you use a preseeded virtual disk, you select the datastore and exact location, folder, and name of the preseeded disk, which cannot be an IDE disk. Zerto takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. ■ The datastore where the preseeded disk is placed is also used as the recovery datastore for the replicated data. ■ If the preseeded disk is **greater than 1TB on NFS storage**, the VPG creation might fail. This is a known VMware problem when the NFS client does not wait for sufficient time for the NFS storage array to initialize the virtual disk after the RPC parameter of the NFS client times out. The timeout default value is 10 seconds. See VMware documentation, http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1027919, which describes the configuration option to tune the RPC timeout parameter by using the command: `esxcfg-advcfg -s <Timeout> /NFS/SetAttrRPCTimeout` ■ If the protected disks are **non-default geometry**, configure the VPG using preseeded volumes. ■ If the protected disk is an **RDM disk**, it can be used to preseed to a recovery VMDK disk. Zerto makes sure that the VMDK disk size is a correct match for the RDM disk. ■ If the VPG is being defined for a Zerto Organization, ZORG, the location of the preseeded disk must be defined in the Zerto Cloud Manager. See *Zerto Cloud Manager Administration Guide*. Then, define the following: ■ **Datastore:** The Datastore where the preseeded disk is located. Only disks with the same size as the protected disk can be selected when browsing for a preseeded disk. ■ **Path:** The full path to the preseeded disk.

SETTING & DESCRIPTION	SELECT...
■ (vCD) Select a **Volume Source** for recovery from one of the drop-down options: ■ **vCD managed storage policy** ■ **Preseeded volume**	**Volume Source > vCD managed storage policy:** Zerto will select a datastore, from the list of available datastores, in the selected Storage Policy in which to place the Volume, unless the datastore is excluded in the Configure Provider vDCs Dialog. ■ If there are several valid datastores, the datastore with the most available space is selected. ■ Zerto recalculates the datastore available space for each volume sequentially, taking into consideration previously allocated volumes. **Volume Source > Preseeded volume:** Select this when you want to copy the protected data to a virtual disk in the recovery site. ■ Zerto **recommends** using this option particularly for **large disks** so that the initial synchronization is **faster** since a Delta Sync can be used to synchronize any changes written to the recovery site after the creation of the preseeded disk. ■ If a preseeded disk is **not** selected, the initial synchronization phase must copy the **whole disk** over the WAN. ■ The preseeded volume is a virtual disk (the VMDK flat file and descriptor) in the recovery site that has been prepared with a copy of the protected data. ■ Browse to the preseed folder configured for the customer and the disk name, of the preseeded disk. In order to use a preseeded VMDK, do the following: a) Create a folder in vCD to use for the preseeded disks in the datastore you want to use for the customer. b) Specify this datastore as a provider datastore for preseeded disks in the Configure Provider vDCs window, from the Advanced Settings window, as described in *Zerto Cloud Manager Administration Guide*. c) In the Zerto Cloud Manager specify the Preseed Folder Name for the ZORG, in the Manage ZORG tab. ■ Zerto searches for the preseeded folder in the available datastores in the Org vDCs specified in the vCD Cloud Resources for the ZORG in the Zerto Cloud Manager and takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. If the virtual machine has more than one preseeded disk, these disks must reside on the same datastore. ■ If the preseeded disk is **greater than 1TB on NFS storage**, the VPG creation might fail. This is a known VMware problem when the NFS client does not wait for sufficient time for the NFS storage array to initialize the virtual disk after the RPC parameter of the NFS client times out. The timeout default value is 10 seconds. See VMware documentation, http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1027919, which describes the configuration option to tune the RPC timeout parameter by using the command: **esxcfg-advcfg -s <Timeout> /NFS/SetAttrRPCTimeout** ■ If the VPG is being defined for a Zerto Organization, ZORG, the location of the preseeded disk must be defined in the Zerto Cloud Manager. See *Zerto Cloud Manager Administration Guide*. ■ Zerto supports the SCSI protocol. Only disks that support this protocol can be specified. Virtual machine RDMs in a vCenter Server are replicated as VMDKs in a vCD environment.

SETTING & DESCRIPTION	SELECT...
(vCD) continued	**Storage Policy**: Specify the Storage Policy for recovery from one of the options: ■ Storage Policy per volume is supported only in vCD supported versions, and when the selected Orgvdc is not configured for fast provisioning. ■ Zerto will select a datastore from the selected Storage Policy in which to place these files, unless the datastore is excluded in the Configure Provider vDCs Dialog. ■ The Storage Policies which appear in the drop-down list: ■ Include the **Use VM Default** option (default), which will apply the VM's storage policy to this volume. This is also the Storage Policy default value. ■ Were defined in VMware vCloud Director and are configured in the Orgvdc. ■ Have at least one Datastore that was not excluded as a Recovery Volume in the Configure Provider vDCs Dialog.
Temp Data disk	
If the virtual machine to be replicated includes a temp data disk as part of its configuration.	Specify a mirror disk for replication that is marked as a temp data disk. In this case, data is not replicated to the temp data disk after initial synchronization.
Thin provisioning (vSphere)	
If the recovery volumes are thin-provisioned or not.	If the source disk is thin provisioned, the default for the recovery volume is also thin provisioned. vCD only: Unless the Org vDC only supports thin-provisioned volumes
Dynamic provisioning (Hyper-V)	
If the recovery volumes are dynamic-provisioned or not.	If the source disk is dynamic provisioned, the default for the recovery volume is also dynamic provisioned. vCD only: Unless the Org vDC only supports dynamic-provisioned volumes

24. Click **OK**.
25. Click **NEXT**.
26. The **RECOVERY** step is displayed. Recovery details include the networks to use for failover, move, and for testing failover, and whether scripts should run as part of the recovery operation.

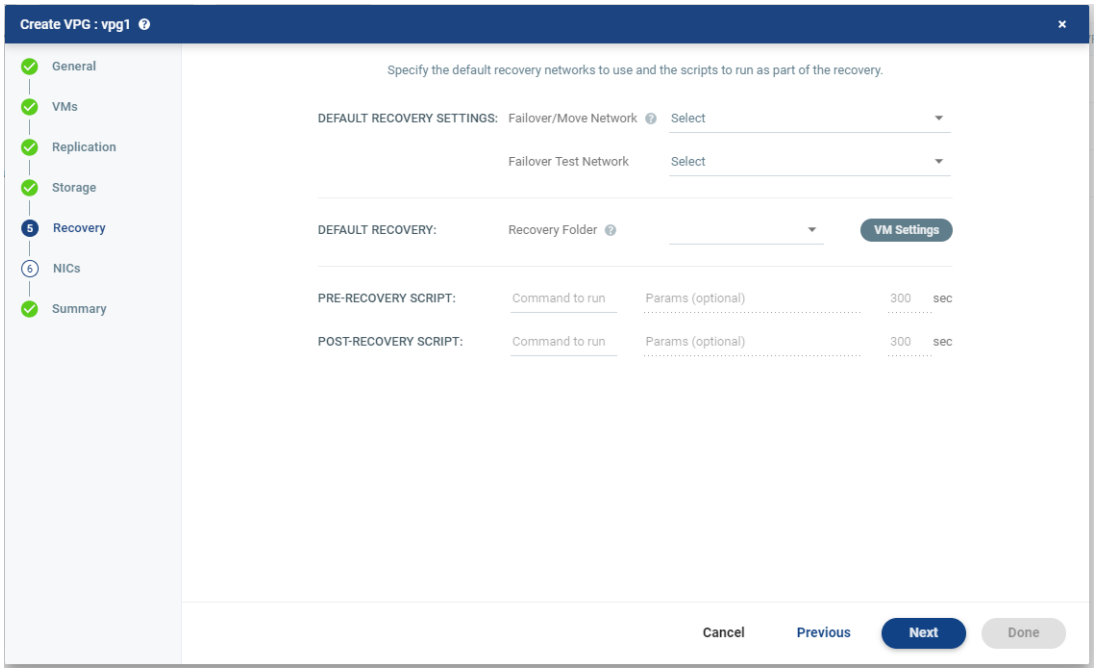

27. Select the default recovery settings. These are applied to every virtual machine in the VPG.
    - **Failover/Move Network:** The network to use during a failover or move operation in which the recovered virtual machines will run.
    - **Failover Test Network:** The network to use when testing the failover of virtual machines in the recovery site. Zerto recommends using a fenced-out network so as not to impact the production network at this site.
    - **Recovery Folder:** The folder to which the virtual machines are recovered.
28. To specify a recovery folder for each virtual machine in the VPG, click **VM SETTINGS.**

    The Advanced VM Recovery Settings window is displayed.

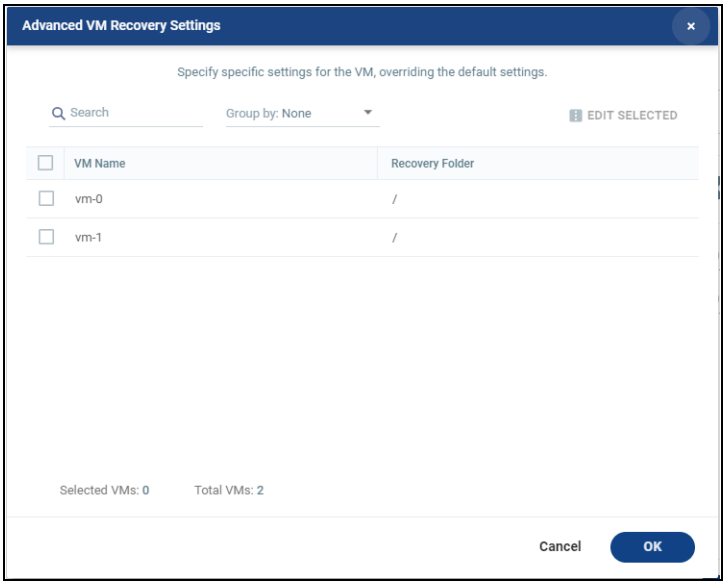

    In this window, you can edit the values of one or more of the virtual machines in the VPG.
    - To edit information in **one field**, click the field and update the information.
29. To edit information for **several virtual machines** at the same time, select the virtual machines and click **EDIT SELECTED**.

The Edit VM window is displayed.

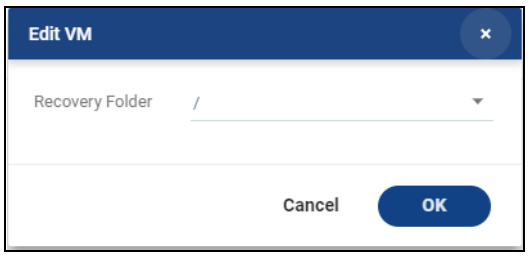

- **Recovery Folder:** The folder to which the virtual machine is recovered.

30. Click **SAVE**.
31. In the Advanced VM Recovery Settings window, click **OK**.
32. Enter the name of the script to run in the **Command to run** text box. You can then enter details about the script.

    Pre-recovery Script: The information about a script that should run at the beginning of the recovery process.

    Post-recovery Script: The information about a script that should run at the end of the recovery process.

    For both types of scripts, enter the following information:

TEXT BOX	DESCRIPTION
Command to run	The full path of the script. The script must be located on the same machine as the Zerto Virtual Manager for the recovery site.
Params	The parameters to pass to the script. Separate parameters with a space.
Timeout	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 the 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 time-out value is specified in Performance and Throttling tab in the Site Settings window.

33. Click **NEXT**.

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

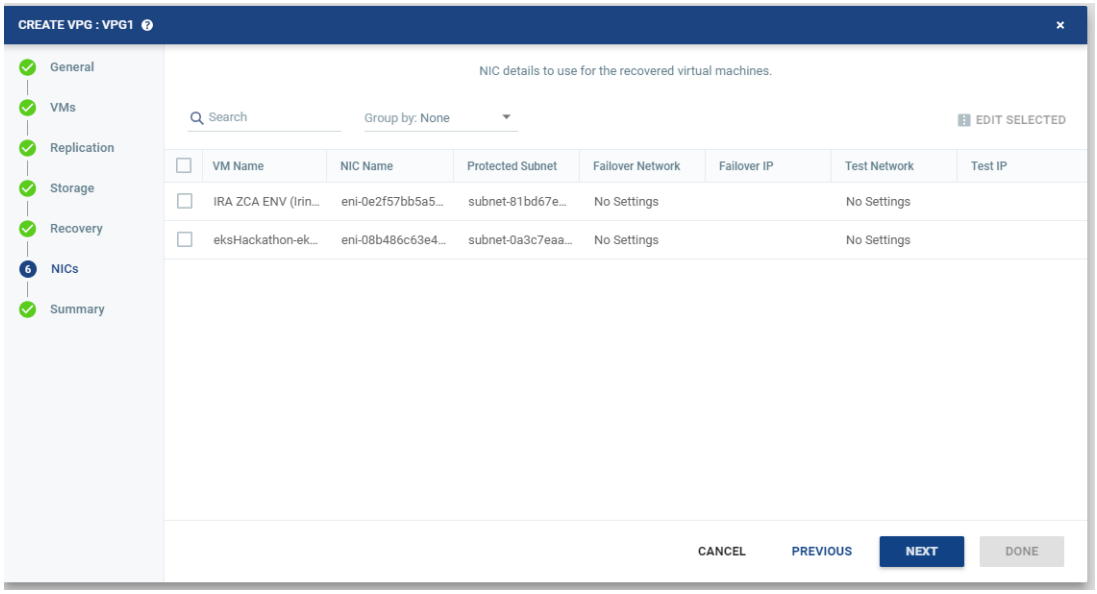

34. To edit information in one field, click the required field and update the information.
35. To edit information for several virtual machines at the same time, select the virtual machines and click **EDIT SELECTED**.
36. Otherwise, go to step 38.

The Edit vNIC window is displayed.

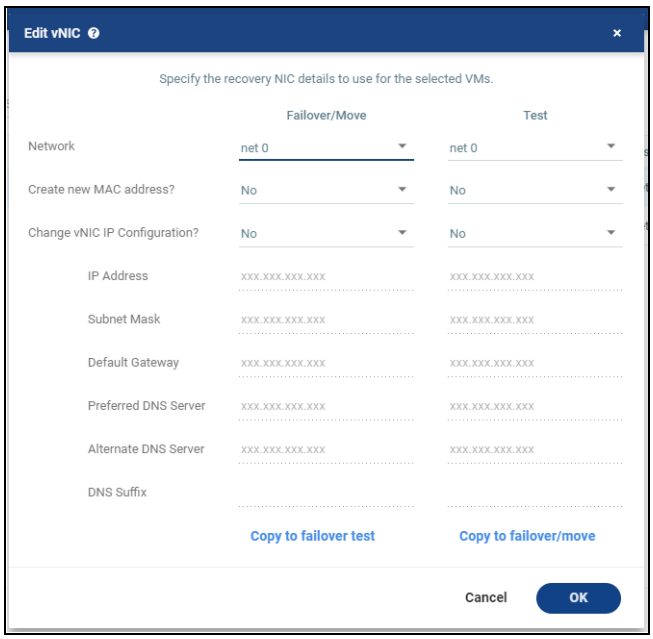

Specify the network details to use for the recovered virtual machines after a failover or move operation, in the Failover/Move column, and for the recovered virtual machines when testing replication, in the Test column.

37. In each column, **specify** the following:
    - **Network:** The network to use for this virtual machine.
    - **Create new MAC address?:** Whether the Media Access Control address (MAC address) used on the protected site should be replicated on the recovery site. The default is to use the same MAC address on both sites. Note that if you check this option, to create a new MAC address, and the current IP address is not specified, the protected virtual machine static IP address might not be used for the recovered virtual machine.
    - **Change vNIC IP Configuration?**: Whether or not to keep the default virtual NIC (vNIC) IP configuration. The vNIC IP is only changed after recovery for virtual machines with VMware Tools running.
      See the Zerto Interoperability Matrix for the list of operating systems for which Zerto supports Re-IPing.
      - To change the vNIC IP, in the Failover/Move or Test column, select **Yes**. If you select to use a static IP connection, set the **IP address**, **subnet mask**, and **default gateway**.
      - Optionally, change the preferred and **alternate DNS server IPs** and the DNS suffix.
      - If you leave the DNS server and suffix entries empty, or select to use DHCP, the IP configuration and DNS server configurations are assigned automatically, to match the protected virtual machine. You can change the DNS suffix.
      - If the virtual machine has multiple NICs but is configured to only have a single default gateway, fill in a **0** for each octet in the Default gateway field for the NICs with no default gateway.

      During a failover, move, or test failover, if the recovered virtual machine is assigned a different IP than the original IP, then after the virtual machine has started, it is automatically rebooted so that it starts up with the correct IP. If the same network is used for both production and test failovers, Zerto recommends changing the IP address for the virtual machines started for the test, so that there is no IP clash between the test machines and the production machines.
    - **Copy to failover test:** Select this to copy the settings in the Failover/Move column to the **Test** column.
    - **Copy to failover/move:** Select this to copy the settings in the Test column to the **Failover/Move** column.

    Click **OK**.
38. Click **NEXT**.

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

For details of what happens after saving the VPG, see "What Happens After the VPG is Defined", on page 35.

Maintained VM Properties when Recovering to a vCenter Server Site from a Protected AWS Site

The following conversions are done to a protected virtual machine when it is recovered in vSphere:

- The SCSI controller type is operating system dependent.
- All disks are thin provisioned.
- Recovered virtual machines use the VMware Virtual E1000 network adapter.
- Operating systems will be either Windows 2012 (64-bit) or Linux Other (64-bit)
- Memory and CPU properties will be extracted from the instance type in AWS.

The following tables display settings that are retained when recovering to a vCenter Server site from AWS:

Properties Extracted From Instance Size in AWS:

SETTINGS	SETTING RETAINED?
CPU Cores	Yes
Memory	Yes

Recovered VM Properties

SETTINGS	SETTING RETAINED?
HW Version	Highest version supported by the hosting ESXi
SCSI Controller Type	Depends on the operating system. Linux: LSI Logic Parallel Windows: LSI Logic SAS
Disks Location	■ Root device /dev/xvda --> SCSI 0:0 ■ /dev/xvdg/ --> SCSI 0:6
Disk Type	The default is thin provision, but configurable in the VPG screen.
Boot Firmware	BIOS.
vNIC type	E1000.
VM IP	Per NIC, according to the configuration in the VPG.

Replication From a Protected Site AWS to a Recovery Site vCloud Director

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

When creating a VPG from AWS to a vCloud Director, all recovery operations bring up the recovered machines on VMware vCenter Server hosts.

The hardware version of the virtual machine must be the same or less than the hardware version supported by the vDC in vCloud Director (vCD), otherwise recovery of the virtual machine in vCD is not permitted.

Zerto uses SCSI for vCenter Server virtual machine disks.

When protecting virtual machines from AWS to vCloud Director, the operating systems of the protected machines must be supported by vCenter Server. Refer to the VMware documentation for a list of supported operating systems.

The following conversions are done to a protected virtual machine in when it is recovered in vCD:

- The SCSI controller type is operating system dependent.
- All disks are thin provisioned.
- Recovered virtual machines use the VMware Virtual E1000 network adapter.
- Operating systems will be either Windows 2012 (64-bit) or Linux Other (64-bit)
- Memory and CPU properties will be extracted from the instance type in AWS.
-

To create a virtual protection group (VPG) to recover in vCD:

1. In the Zerto User Interface, select **ACTIONS > CREATE VPG**.
   The GENERAL step of the Create VPG wizard is displayed.

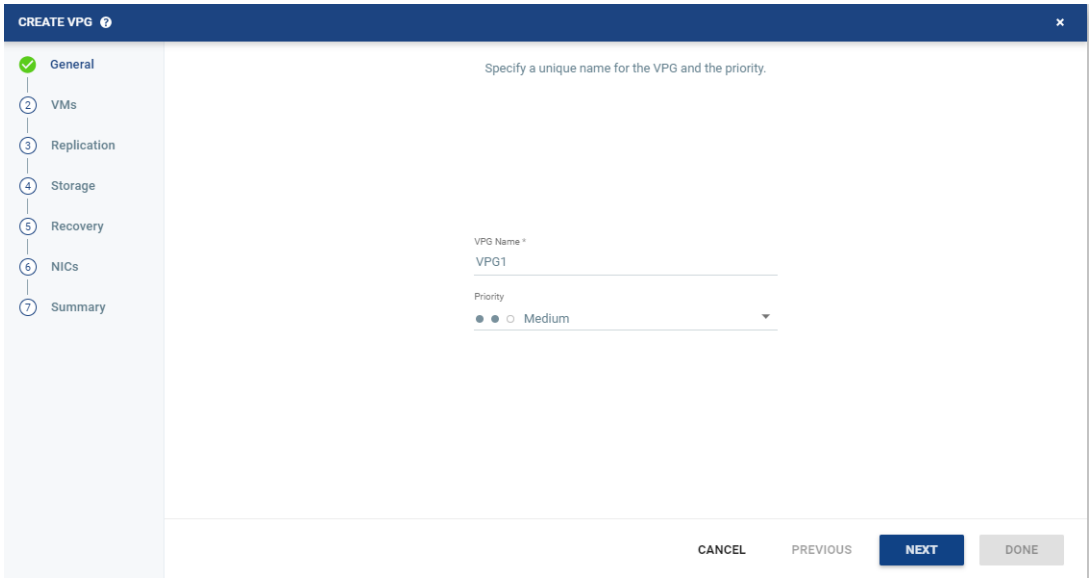

2. Specify the name of the VPG and the priority of the VPG.
   - **VPG Name:** The VPG name must be unique. The name cannot be more than 80 characters.
   - **Priority:** Determine the priority for transferring data from the protected site to the recovery site when there is limited bandwidth and more than one VPG is defined on the protected site.
     - **High Priority**: When there are updates to virtual machines protected in VPGs with different priorities, updates from the VPG with the highest priority are passed over the WAN first.
     - **Medium Priority:** Medium priority VPGs will only be able to use whatever bandwidth is left after the high priority VPGs have used it.
     - **Low Priority:** Low priority VPGs will use whatever bandwidth is left after the medium priority VPGs have use it.

     Updates to the protected virtual machines are always sent across the WAN before synchronization data, such as during a bitmap or delta sync.

During synchronization, data from the VPG with the highest priority is passed over the WAN before data from medium and low priority VPGs.

3. Click **NEXT**.

   The VMs step is displayed.

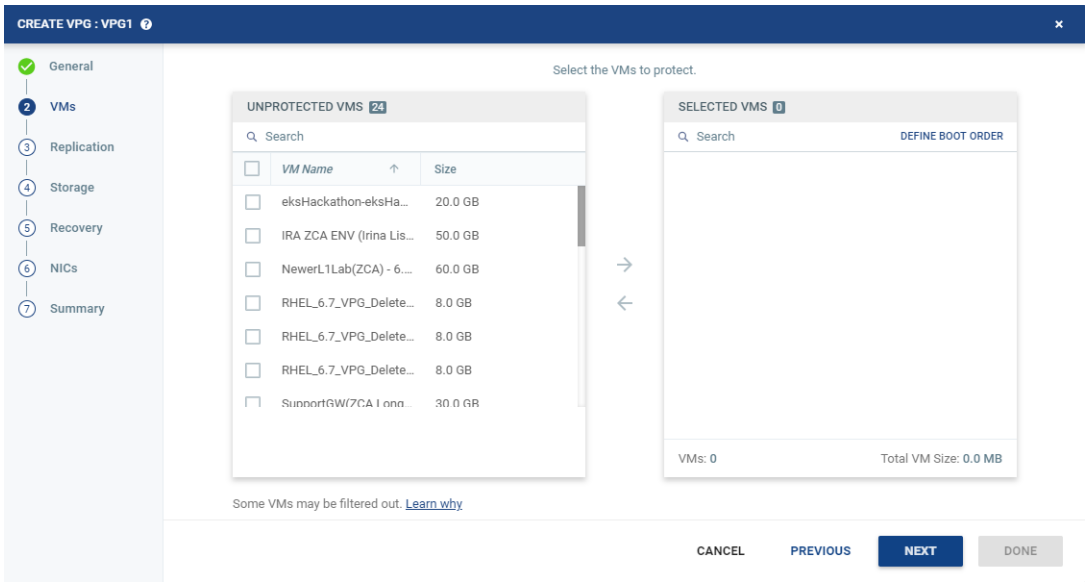

- The list of unprotected VMs includes VMs that were in the ZCA's **account and region** during installation.
- A VPG can include virtual machines that are already protected by another ZCA in the same account and region.
- The size of the disk is defined by data disks and the OS disk.
- AWS Instance Store disks (Temp disks) are not included in the list of unprotected VMs.

4. From the **Unprotected VMs** list, select the VMs to include in this VPG and click the arrow to move them to the **Selected VMs**.
   - When using the **Search** field, you can use the wildcards; ***** or **?**
5. To define the boot order of the virtual machines in the VPG, click **Define Boot Order**, otherwise go to the next step.

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

   Note: Up to five (5) virtual machines may boot on a host simultaneously. Following the boot, a 300 second (default) delay occurs until the next boot batch.

   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.

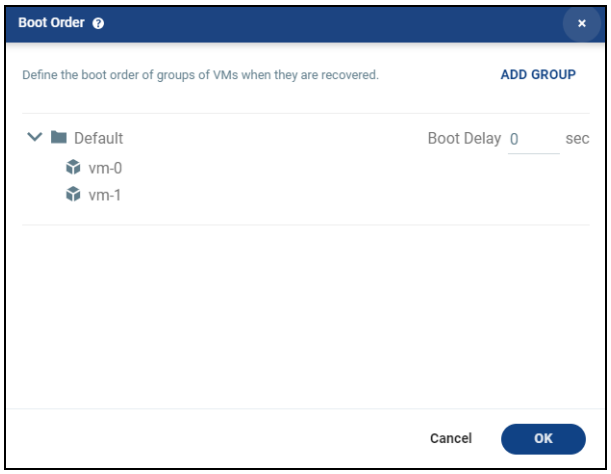

6. Click **ADD GROUP** to add a new group. Then, do the following:
   a) To change the name of a group, click the Pencil icon next to the group.
   b) To delete a group, click the delete icon on the right side. You cannot delete the Default group nor a group that contains a virtual machine.
   c) Drag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

   Click **NEXT**.

   The REPLICATION step is displayed.
7. Select the **Recovery Site**, and select **vCD** from the drop-down list. (If there is only one site available, it's displayed by default.

   The **REPLICATION** step is re-displayed, with additional fields relevant for vCD.

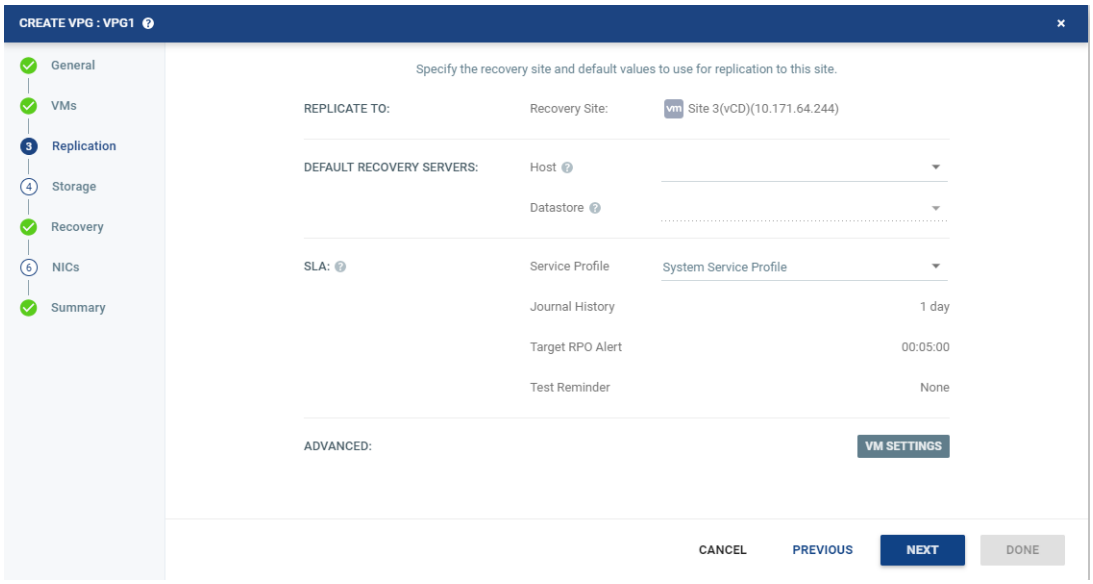

8. Select the **Recovery Org vDC** to use in the recovery site.

   Note: You cannot select a recovery site if any of the virtual machines you selected are already in VPGs that recover to that site.

   Note: At least one VRA needs to be installed on a host within a resource pool being used by a Recovery Org vDC.
9. In the **SLA** area, you define the Service Level Agreement for which this VPG is associated.
   - When **Zerto Cloud Manager** is used, select the **Service Profile** to use.
     The Service Profile determines the VPG SLA settings for the group. This applies predefined settings for the Journal History, Target RPO Alert and the Test Reminder. These settings apply to **every** virtual machine in the group.
   - When a **Custom** service profile is available, the VPG SLA settings are editable, and the **Advanced** button becomes available. When you change these settings, they apply to **every** virtual machine in the group.
10. Click **ADVANCED**. The Advanced Journal Settings dialog is displayed.

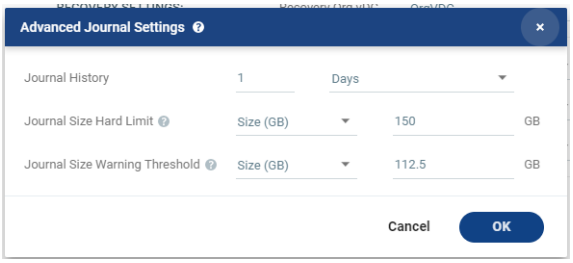

SETTING & DESCRIPTION	SELECT...
Journal History	
The time that all write commands are saved in the journal. The longer the information is saved in the journal, the more space is required for each journal in the VPG.	■ Number of **hours** from **1 to 23** ■ Number of **days** from **1 to 30**
Default Journal Storage (Hyper-V), or **Default Journal Datastore** (vSphere)	
The storage/datastore used for the journal data for each virtual machine in the VPG. **Note:** This field is **not** relevant when replicating to a **vCD** recovery site.	■ Select the storage/datastore accessible to the host. When you select a specific journal storage/datastore, the journals for each virtual machine in the VPG are stored in this storage/datastore, regardless of where the recovery storage/datastore is for each virtual machine. All protected virtual machines are recovered to the hosts that can access the specified journal storage/datastore.
Journal Size Hard Limit	

The maximum size that the journal can grow, either as a percentage or a fixed amount.
 	
The journal is always **thin-provisioned**.
 	
Note: The Journal Size Hard Limit applies independently **both** to the Journal History **and** also to the Scratch Journal Volume.
 	
For Example: If the Journal Size Hard Limit is configured to a maximum size of 160 GB limit, then during Failover Test, both the Journal History and the Scratch Journal Volume together can take up to 320 GB. Each one with a maximum size of 160 GB limit.
	■ **Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore.
	
	If Unlimited is selected, Size and Percentage options are **not** displayed.
	
	■ **Size (GB):** The maximum journal size in GB.
	  ■ The **minimum** journal size, set by Zerto, is **8GB** for Hyper-V and vSphere environments, and **10GB** for Microsoft Azure environments.
	
	■ **Percentage:** The percentage of the virtual machine volume size to which the journal can grow.
	  ■ This value can be configured to more than 100% of the protected VM's volume size.

SETTING & DESCRIPTION	SELECT...
Journal Size Warning Threshold	
The size of the journal that triggers a warning that the journal is nearing its hard limit.	■ **Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore. If Unlimited is selected, Size and Percentage options are **not** displayed. ■ **Size* (GB):** The size in GB that will generate a warning. ■ **Percentage*:** The percentage of the virtual machine volume size that will generate a warning. *The values of **Size** and **Percentage** must be **less** than the configured **Journal Size Hard Limit** so that the warning will be generated when needed. In addition to the warning threshold, Zerto will issue a message when the free space available for the journal is almost full.

11. Click **OK**.
    - **Target RPO Alert:** The maximum desired time between each automatic checkpoint write to the journal before an alert is issued.
    - **Test Reminder:** The amount of time in months recommended between each test, where you test the integrity of the VPG. A warning is issued if a test is not performed within this time frame.

    To change the replication settings per virtual machine, click **VM SETTINGS**.

    The Advanced VM Replication Settings window is displayed.

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

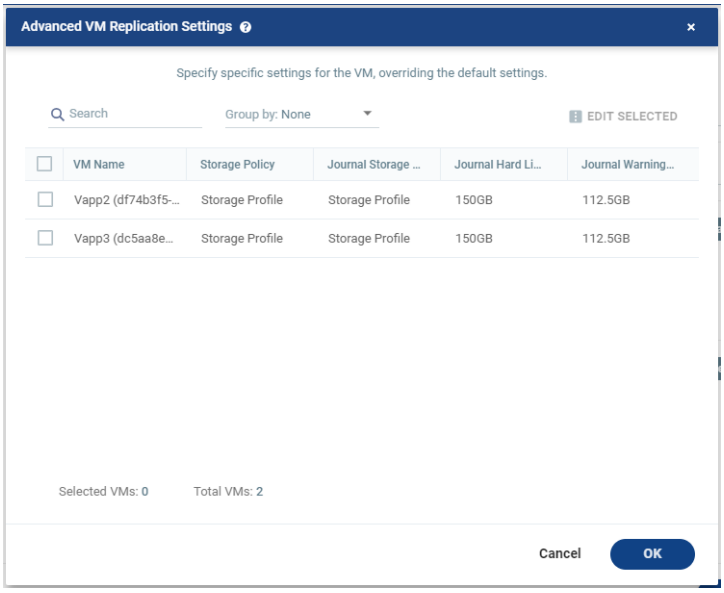

12. To edit information for a **single** VM, click the field **Storage Policy**, or **Journal Storage Policy**, and update the information.
    - **Storage Policy**: The Storage Policy in which the **VM configuration files** will reside. For considerations, see step 14.
    - **Journal Storage Policy**: The Storage Policy in which the **VM Journal files** will reside. For considerations, see step15.
13. To edit information for **several** virtual machines at the same time, select the virtual machines and click **EDIT SELECTED**.

The Edit VM window is displayed.

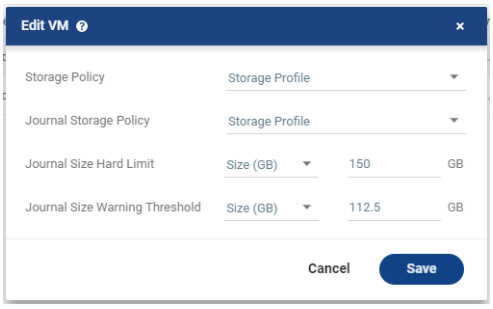

14. When selecting **Storage Policy**, consider the following:
    - Zerto will select a datastore from the selected Storage Policy in which to place these files, unless the datastore is excluded in the Configure Provider vDCs Dialog.
    - Zerto will try to determine a default Storage Policy according to:
        - A Storage Policy with the same name as the protected Storage Policy.
        - The default Orgvdc Storage Policy.

      If Zerto did not manage to determine a default Storage Policy, this field appears empty.
    - When you **click to edit**, a list of Storage Policies appear. These Storage Policies:
        - Were defined in VMware vCloud Director and are configured in the Orgvdc.
        - Have at least one Datastore that was not excluded as a Recovery Volume in the Configure Provider vDCs Dialog.
15. When selecting **Journal Storage Policy**, consider the following:
    - Zerto will select a datastore from the selected Storage Policy in which to place the Journal files, unless the datastore is excluded in the Configure Provider vDCs Dialog.
    - The default Journal Storage Policy is the same as the default VM Storage Policy.
    - If Zerto did not manage to determine a default Journal Storage Policy, this field appears empty.
    - When you **click to edit**, the option **Auto Select** appears, and a list of Storage Policies.
    - The list of Storage Policies associated with the Journal:
        - Were defined in VMware vCloud Director and are configured in the Orgvdc.
        - Have at least one Datastore that was not excluded as a **Journal** in the Configure Provider vDCs Dialog.
    - **Auto Select**: Selecting this means that the journal can be placed in any datastore visible to the host that Zerto selected for recovery, unless the datastore is excluded in the Configure Provider vDCs Dialog.

SETTING & DESCRIPTION	SELECT...
Recovery Host (not relevant when replicating to vCD)	
(Hyper-V) The cluster or host that will host the recovered virtual machine.	
(vSphere) The cluster, resource pool, or host that will host the recovered virtual machine. If the site is defined in Zerto Cloud Manager, only a resource pool can be specified and the resource pool must also have been defined in Zerto Cloud Manager. For details about Zerto Cloud Manager, see *Zerto Cloud Manager Administration Guide*. When a resource pool is specified, Zerto checks that the resource pool capacity is enough for all the virtual machines specified in the VPG	When a resource pool is specified, Zerto checks that the resource pool capacity is enough for all the virtual machines specified in the VPG. If a resource pool is specified and DRS is disabled for the site later on, all the resource pools are removed by VMware and recovery is to any one of the hosts in the recovery site with a VRA installed on it. All resource pool checks are made at the level of the VPG and do not take into account multiple VPGs using the same resource pool. If the resource pool CPU resources are defined as unlimited, the actual limit is inherited from the parent but if this inherited value is too small, failover, move, and failover test operations can fail, even without a warning alert being issued by Zerto Virtual Manager.
VM Recovery Datastore (vSphere) (not relevant when replicating to vCD)	

SETTING & DESCRIPTION	SELECT...
The datastore where the VMware metadata files for the virtual machine are stored, such as the VMX file.	If a cluster or resource pool is selected for the host, only datastores that are accessible by every ESX/ESXi host in the cluster or resource pool are displayed. This is also the datastore where RDM backing files for recovery volumes are located.
Recovery Storage (Hyper-V) (not relevant when replicating to vCD)	
The location where the metadata files for the virtual machine are stored, such as the VHDX file. I	If a cluster is selected for the host, only storage that are accessible by every host in the cluster are displayed.
Journal Size Hard Limit	
The maximum size that the journal can grow, either as a percentage or a fixed amount. ■ The journal is always **thin-provisioned**. ■ The Journal Size Hard Limit applies independently **both** to the Journal History **and** also to the Scratch Journal Volume. *For Example:* If the Journal Size Hard Limit is configured to a maximum size of 160 GB limit, then during Failover Test, both the Journal History and the Scratch Journal Volume together can take up to 320 GB. Each one with a maximum size of 160 GB limit.	**Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore. If Unlimited is selected, Size and Percentage options are **not** displayed. **Size (GB):** The maximum journal size in GB. ■ The **minimum** journal size, set by Zerto, is **8GB** for Hyper-V and vSphere environments, and **10GB** for Microsoft Azure environments. **Percentage:** The percentage of the virtual machine volume size to which the journal can grow. ■ This value can be configured to more than 100% of the protected VM's volume size.
Journal Size Warning Threshold	
The size of the journal that triggers a warning that the journal is nearing its hard limit.	**Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore. If Unlimited is selected, Size and Percentage options are **not** displayed. **Size* (GB):** The size in GB that will generate a warning. **Percentage*:** The percentage of the virtual machine volume size that will generate a warning. *The values of **Size** and **Percentage** must be **less** than the configured **Journal Size Hard Limit** so that the warning will be generated when needed. In addition to the warning threshold, Zerto will issue a message when the free space available for the journal is almost full.

SETTING & DESCRIPTION	SELECT...
Journal Storage (Hyper-V), or **Journal Datastore** (vSphere) (not relevant when replicating to vCD)	
The storage/datastore used for the journal data for each virtual machine in the VPG.	**(vSphere)** To change the default, specify a host and then select one of the datastores accessible by this host to be used as the journal datastore. When you select specific journal datastore, the journals for each virtual machine in the VPG are stored in this datastore, regardless of where the recovery datastores are for each virtual machine. In this case, all the protected virtual machines must be recovered to hosts that can access the specified journal datastore. **(Hyper-V)** To change the default, specify a host and then select the storage location accessible by this host to be used as the journal storage. When you select specific journal storage, the journals for each virtual machine in the VPG are stored in this storage, regardless of where the recovery storage is for each virtual machine. In this case, all the protected virtual machines must be recovered to hosts that can access the specified journal storage.

16. In the Advanced VM Replication Settings window, click **OK**.
17. Click **NEXT**.
    The **Storage** step is displayed.
    By default the storage used for the virtual machine definition is also used for the virtual machine data.
    For each virtual machine in the VPG, Zerto displays its storage-related information.

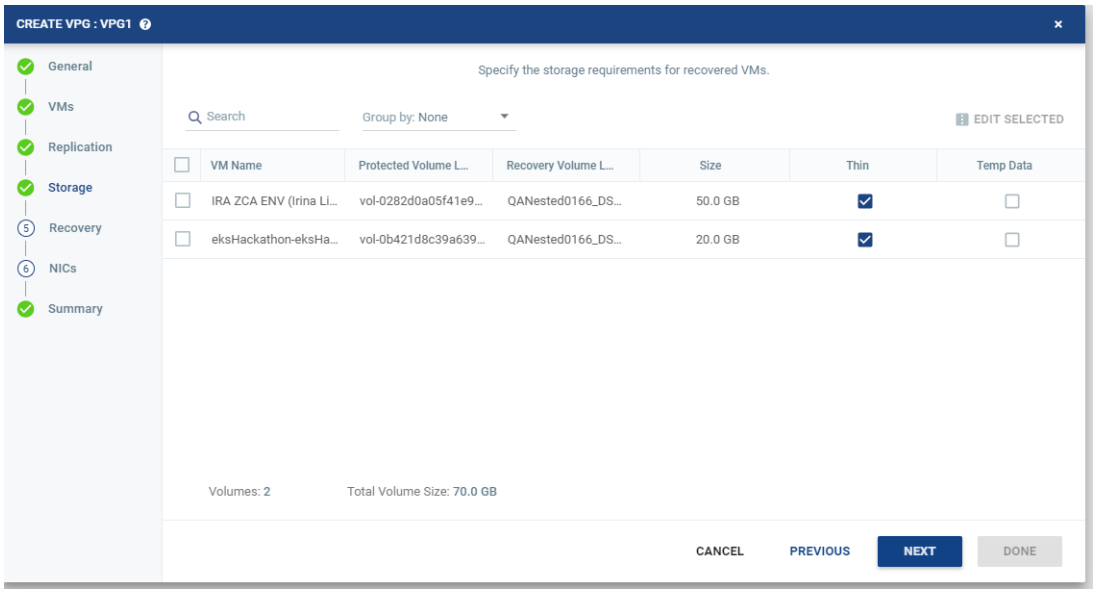

18. To edit storage information for one of the virtual machines volumes, select the volume/s and click **EDIT SELECTED.**

    - For **Hyper-V** recovery environments, the following window appears. For details, click here.

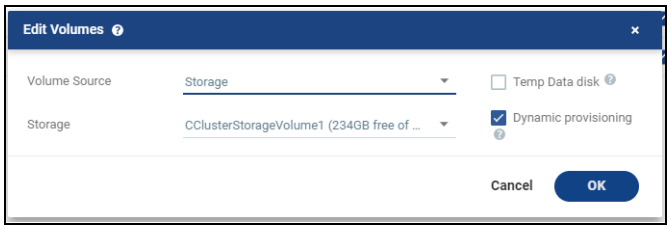

- For **vSphere** recovery environments, the following window appears. For details, click here.

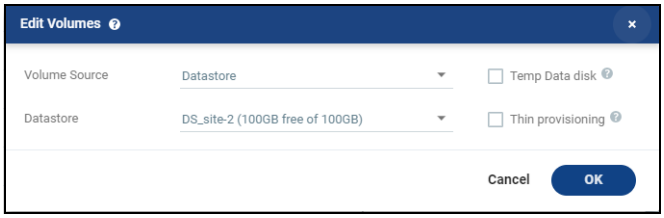

- For **vCD** recovery environments, the following window appears. For details, click here.

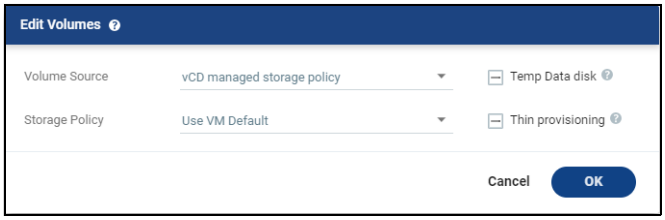

SETTING & DESCRIPTION	SELECT...
Volume Source	
■ (Hyper-V) Select a **Volume Source** for recovery from one of the drop-down options: ■ **Storage** ■ **Preseeded volume**	■ **Volume Source > Storage:** A new volume is used for replicated data. ■ From the **Storage** drop-down list, specify the storage to use to create disks for the replicated data. The storage specified for the replication must have at least the same amount of space as the protected volume and then an additional amount for the journal. The amount of additional space needed for the journal can be fixed by specifying a maximum size for the journal, or can be calculated as the average change rate for the virtual machines in the VPG, multiplied by the length of time specified for the journal history. For more details, see *Zerto Scale and Benchmarking Guidelines*, in the section **Estimating WAN Bandwidth for VMware and Hyper-V**.
	■ **Volume Source > Preseeded volume:** Whether to copy the protected data to a virtual disk in the recovery site. Zerto **recommends** using this option particularly for large disks so that the initial synchronization will be faster since a **Delta Sync** can be used to synchronize any changes written to the recovery site after the creation of the preseeded disk. When **not** using a preseeded disk, the initial synchronization phase must copy the whole disk over the WAN. When using a preseeded virtual disk, you select the storage and exact location, folder, and name of the preseeded disk. Zerto takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. Only disks with the same size as the protected disk can be selected when browsing for a preseeded disk. The storage where the preseeded disk is placed is also used as the recovery storage for the replicated data.

SETTING & DESCRIPTION	SELECT...
■ (vSphere) Select a **Volume Source** for recovery from one of the drop-down options: ■ **Datastore** ■ **RDM** ■ **Preseeded volume**	**Volume Source > Datastore:** A new volume is used for replicated data. 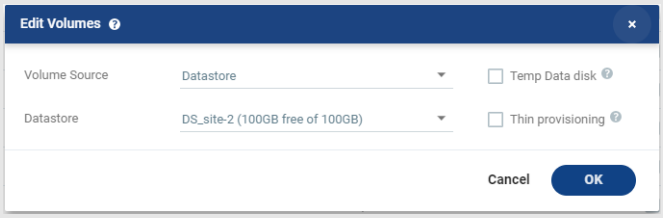 ■ Specify the **Datastore** to use to create disks for the replicated data. ■ If the **source disk is thin provisioned**, the default for the recovery volume is also thin provisioned. ■ The datastore specified for replication must have at least the same amount of space as the **protected volume** and an additional amount for the **journal**. ■ The amount of additional space needed for the journal can be fixed by specifying a maximum size for the journal, or can be calculated as the average change rate for the virtual machines in the VPG, multiplied by the length of time specified for the journal history. ■ Zerto supports the SCSI protocol. Only disks that support this protocol can be specified. Then, define the following: ■ **Datastore:** The Datastore where the preseeded disk is located. Only disks with the same size as the protected disk can be selected when browsing for a preseeded disk. For more details, see *Zerto Scale and Benchmarking Guidelines*, in the section **Estimating WAN Bandwidth for VMware and Hyper-V**.
	Volume Source > RDM: The VMware RDM (Raw Device Mapping) which will be used for the replication. 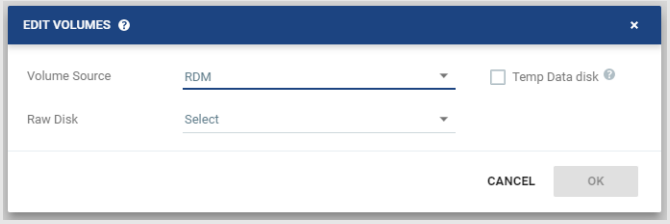 By default, **RDM is recovered as thin-provisioned VMDK** in the datastore specified in the **VM Recovery Datastore/Storage** field in the Edit VM dialog, and not to RDM. Only a raw disk with the **same size as the protected disk** can be selected from the list of available raw disks. Other raw disks with different sizes are not available for selection. The RDM is always stored in the recovery datastore, used for the virtual machine. The following **limitations** apply to protecting RDM disks: ■ RDM disks with an even number of blocks can replicate to RDM disks of the same size with an even number of blocks and to VMDKs. ■ RDM disks with an odd number of blocks can only replicate to RDM disks of the same size with an odd number of blocks and not to VMDKs. ■ You cannot define an RDM disk to be protected to a cloud service provider via a Zerto Cloud Connector nor if the virtual machine uses a BusLogic SCSI controller, nor when protecting or recovering virtual machines in an environment running vCenter Server 5.x with ESX/ESXi version 4.1 hosts.

SETTING & DESCRIPTION	SELECT...
■ (vSphere) **Volume Source** continued	**Volume Source > Preseeded volume**: Select this when you want to copy the protected data to a virtual disk in the recovery site. 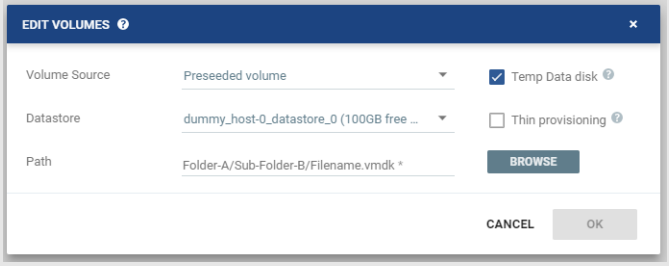 Consider the following, then proceed to define the Datastore and the Path: ■ Zerto **recommends** using this option particularly for **large disks** so that the initial synchronization is **faster** since a Delta Sync can be used to synchronize any changes written to the recovery site after the creation of the preseeded disk. ■ If a preseeded disk is **not** selected, the initial synchronization phase must copy the **whole disk** over the WAN. ■ If you use a preseeded virtual disk, you select the datastore and exact location, folder, and name of the preseeded disk, which cannot be an IDE disk. Zerto takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. ■ The datastore where the preseeded disk is placed is also used as the recovery datastore for the replicated data. ■ If the preseeded disk is **greater than 1TB on NFS storage**, the VPG creation might fail. This is a known VMware problem when the NFS client does not wait for sufficient time for the NFS storage array to initialize the virtual disk after the RPC parameter of the NFS client times out. The timeout default value is 10 seconds. See VMware documentation, http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1027919, which describes the configuration option to tune the RPC timeout parameter by using the command: `esxcfg-advcfg -s <Timeout> /NFS/SetAttrRPCTimeout` ■ If the protected disks are **non-default geometry**, configure the VPG using preseeded volumes. ■ If the protected disk is an **RDM disk**, it can be used to preseed to a recovery VMDK disk. Zerto makes sure that the VMDK disk size is a correct match for the RDM disk. ■ If the VPG is being defined for a Zerto Organization, ZORG, the location of the preseeded disk must be defined in the Zerto Cloud Manager. See *Zerto Cloud Manager Administration Guide*. Then, define the following: ■ **Datastore:** The Datastore where the preseeded disk is located. Only disks with the same size as the protected disk can be selected when browsing for a preseeded disk. ■ **Path:** The full path to the preseeded disk.

SETTING & DESCRIPTION	SELECT...
■ (vCD) Select a **Volume Source** for recovery from one of the drop-down options: ■ **vCD managed storage policy** ■ **Preseeded volume**	**Volume Source > vCD managed storage policy:** Zerto will select a datastore, from the list of available datastores, in the selected Storage Policy in which to place the Volume, unless the datastore is excluded in the Configure Provider vDCs Dialog. ■ If there are several valid datastores, the datastore with the most available space is selected. ■ Zerto recalculates the datastore available space for each volume sequentially, taking into consideration previously allocated volumes. **Volume Source > Preseeded volume:** Select this when you want to copy the protected data to a virtual disk in the recovery site. ■ Zerto **recommends** using this option particularly for **large disks** so that the initial synchronization is **faster** since a Delta Sync can be used to synchronize any changes written to the recovery site after the creation of the preseeded disk. ■ If a preseeded disk is **not** selected, the initial synchronization phase must copy the **whole disk** over the WAN. ■ The preseeded volume is a virtual disk (the VMDK flat file and descriptor) in the recovery site that has been prepared with a copy of the protected data. ■ Browse to the preseed folder configured for the customer and the disk name, of the preseeded disk. In order to use a preseeded VMDK, do the following: a) Create a folder in vCD to use for the preseeded disks in the datastore you want to use for the customer. b) Specify this datastore as a provider datastore for preseeded disks in the Configure Provider vDCs window, from the Advanced Settings window, as described in *Zerto Cloud Manager Administration Guide*. c) In the Zerto Cloud Manager specify the Preseed Folder Name for the ZORG, in the Manage ZORG tab. ■ Zerto searches for the preseeded folder in the available datastores in the Org vDCs specified in the vCD Cloud Resources for the ZORG in the Zerto Cloud Manager and takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. If the virtual machine has more than one preseeded disk, these disks must reside on the same datastore. ■ If the preseeded disk is **greater than 1TB on NFS storage**, the VPG creation might fail. This is a known VMware problem when the NFS client does not wait for sufficient time for the NFS storage array to initialize the virtual disk after the RPC parameter of the NFS client times out. The timeout default value is 10 seconds. See VMware documentation, http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1027919, which describes the configuration option to tune the RPC timeout parameter by using the command: **esxcfg-advcfg -s <Timeout> /NFS/SetAttrRPCTimeout** ■ If the VPG is being defined for a Zerto Organization, ZORG, the location of the preseeded disk must be defined in the Zerto Cloud Manager. See *Zerto Cloud Manager Administration Guide*. ■ Zerto supports the SCSI protocol. Only disks that support this protocol can be specified. Virtual machine RDMs in a vCenter Server are replicated as VMDKs in a vCD environment.

SETTING & DESCRIPTION	SELECT...
(vCD) continued	**Storage Policy**: Specify the Storage Policy for recovery from one of the options: ■ Storage Policy per volume is supported only in vCD supported versions, and when the selected Orgvdc is not configured for fast provisioning. ■ Zerto will select a datastore from the selected Storage Policy in which to place these files, unless the datastore is excluded in the Configure Provider vDCs Dialog. ■ The Storage Policies which appear in the drop-down list: ■ Include the **Use VM Default** option (default), which will apply the VM's storage policy to this volume. This is also the Storage Policy default value. ■ Were defined in VMware vCloud Director and are configured in the Orgvdc. ■ Have at least one Datastore that was not excluded as a Recovery Volume in the Configure Provider vDCs Dialog.
Temp Data disk	
If the virtual machine to be replicated includes a temp data disk as part of its configuration.	Specify a mirror disk for replication that is marked as a temp data disk. In this case, data is not replicated to the temp data disk after initial synchronization.
Thin provisioning (vSphere)	
If the recovery volumes are thin-provisioned or not.	If the source disk is thin provisioned, the default for the recovery volume is also thin provisioned. vCD only: Unless the Org vDC only supports thin-provisioned volumes
Dynamic provisioning (Hyper-V)	
If the recovery volumes are dynamic-provisioned or not.	If the source disk is dynamic provisioned, the default for the recovery volume is also dynamic provisioned. vCD only: Unless the Org vDC only supports dynamic-provisioned volumes

19. Click **OK**.
20. Click **NEXT**.

    The RECOVERY step is displayed. Recovery details include the scripts that should be run either at the start or end of a recovery operation.

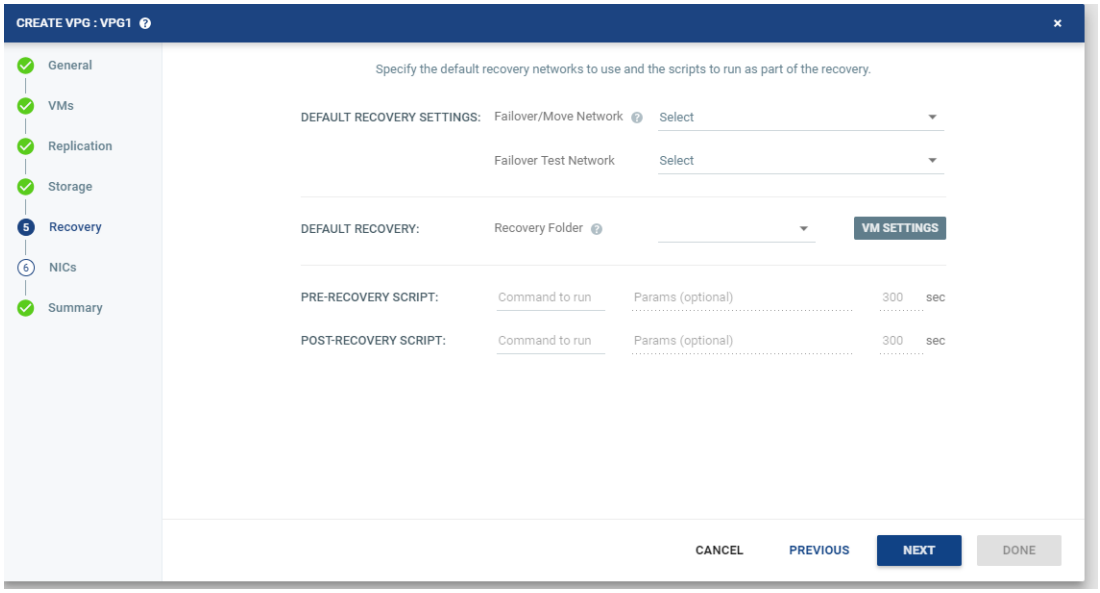

21. Select the default recovery settings.
    - **vCD Guest Customization:** When selected, VMware Guest OS Customization is enabled for the virtual machine in vCloud Director. Enabling guest customization means that the computer name and network settings configured for this virtual machine are applied to its Guest OS when the virtual machine is powered on. vCD Guest Customization must be selected to enable **re-IPing** the recovered virtual machines.
22. To run pre and post recovery scripts, enter the name of the script to run in the Command to run text box. You can then enter details about the script.
    - **Pre-recovery Script:** The information about a script that should run at the beginning of the recovery process.
    - **Post-recovery Script:** The information about a script that should run at the end of the recovery process.

    For both types of scripts, enter the following information:

TEXT BOX	DESCRIPTION
Command to run	The full path of the script. The script must be located on the same machine as the Zerto Virtual Manager for the recovery site.
Params	The parameters to pass to the script. Separate parameters with a space.
Timeout	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 the 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 time-out value is specified in **Site Settings > Performance and Throttling** tab.

NOTE:

Pre and post recovery scripts run in parallel. Therefore, ensure that the pre and post recovery scripts don't use common resources.

23. Click **NEXT**.

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

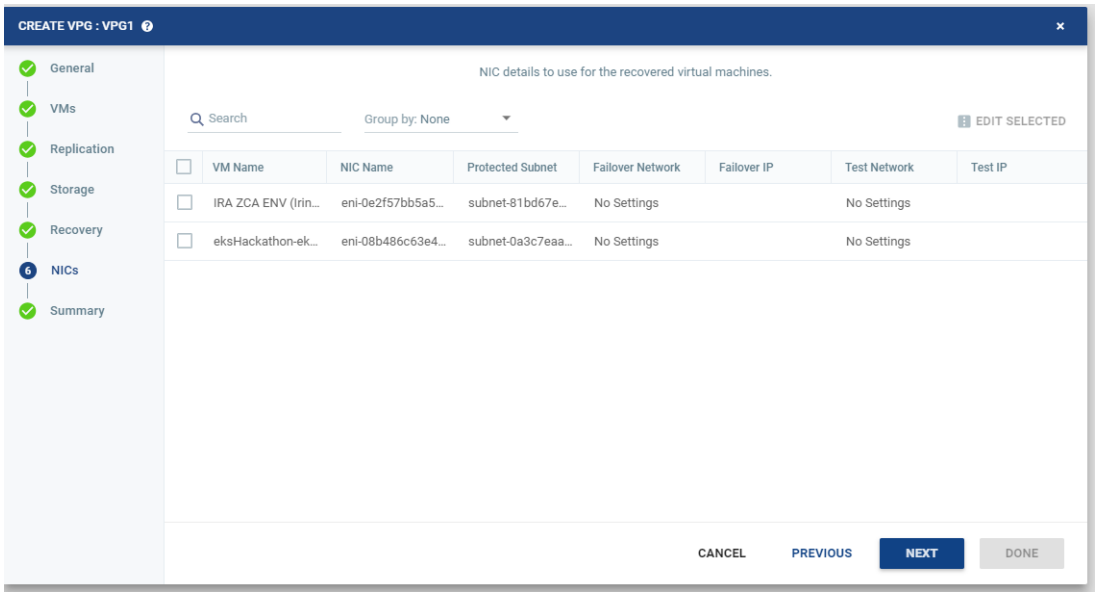

24. To edit information in one field, click the required field and update the information.
25. To edit information for several virtual machines at the same time, select the virtual machines, and click **EDIT SELECTED**.
26. Otherwise, go to step 28.

    The Edit VNIC window is displayed.

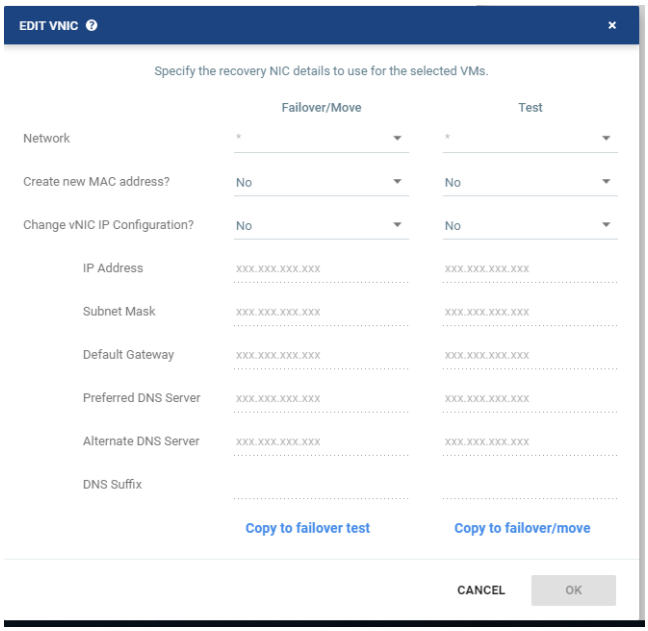

27. Specify the network details to use for the recovered virtual machines after a failover or move operation, in the Failover/Move column, and for the recovered virtual machines when testing replication, in the Test column.

    In each column, specify the following:
    - **Network:** The network to use for this virtual machine.
    - **Create New MAC Address?:** Whether the Media Access Control address (MAC address) used on the protected site should be replicated on the recovery site. The default is to use the same MAC address on both sites.
    - **Change vNIC IP Configuration?:** Which IP mode to use. Specify the IP address if you choose **static IP pool**. See the Zerto Interoperability Matrix for the list of operating systems for which Zerto supports Re-IPing.

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

- **Copy to failover test:** Select this to copy the settings in the Failover/Move column to the **Test** column.
- **Copy to failover/move:** Select this to copy the settings in the Test column to the **Failover/Move** column.

28. Click **OK**.
29. Click **NEXT**.

    The SUMMARY step is displayed. It shows the VPG configuration that you defined in previous steps.
30. Click **DONE**. The VPG is created.

For details of what happens after saving the VPG, see "What Happens After the VPG is Defined", on page 23.

The virtual machines in the VPG are protected as a vCD vApp in the recovery site. When recovering the VPG, reverse protection is configured back the virtual machines.

Replication From a Protected Site AWS to a Recovery Site Hyper-V

You can protect virtual machines to recovery Hyper-V hosts. The procedure is the same whether you intend to protect one virtual machine or multiple virtual machines.

When creating a VPG from AWS to Hyper-V all recovery operations bring up the recovered machines on Microsoft Hyper-V hosts in SCVMM.

When protecting virtual machines from AWS to Hyper-V, the following is required:

- Operating systems of the protected machines must be supported by Hyper-V. Refer to Hyper-V documentation for the list of supported operating systems.
- Virtual machine names cannot include any of the following special characters: `*?:<>/|\"`.

Zerto uses SCSI for Hyper-V SCVMM virtual machine disks.

To create a virtual protection group (VPG) to recover in Hyper-V:

1. In the Zerto User Interface, select **ACTIONS > CREATE VPG**.
   The GENERAL step of the Create VPG wizard is displayed.

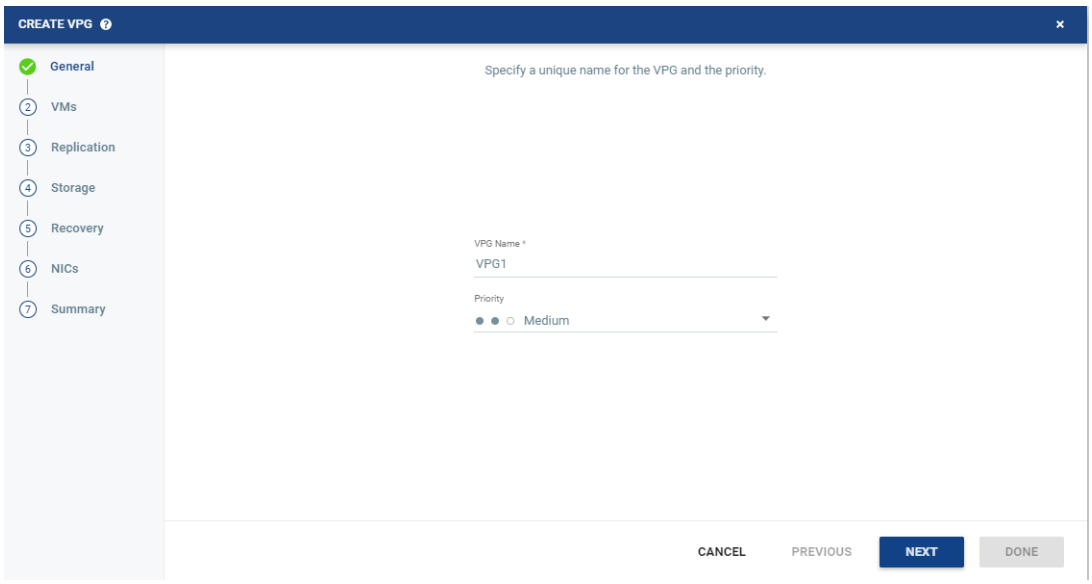

2. Specify the name of the VPG and the priority of the VPG.
   - **VPG Name:** The VPG name must be unique. The name cannot be more than 80 characters.
   - **Priority:** Determine the priority for transferring data from the protected site to the recovery site when there is limited bandwidth and more than one VPG is defined on the protected site.
     - **High Priority**: When there are updates to virtual machines protected in VPGs with different priorities, updates from the VPG with the highest priority are passed over the WAN first.
     - **Medium Priority:** Medium priority VPGs will only be able to use whatever bandwidth is left after the high priority VPGs have used it.
     - **Low Priority:** Low priority VPGs will use whatever bandwidth is left after the medium priority VPGs have use it.

   Updates to the protected virtual machines are always sent across the WAN before synchronization data, such as during a bitmap or delta sync.
   During synchronization, data from the VPG with the highest priority is passed over the WAN before data from medium and low priority VPGs.
3. Click **NEXT**.

The **VM**s step is displayed.

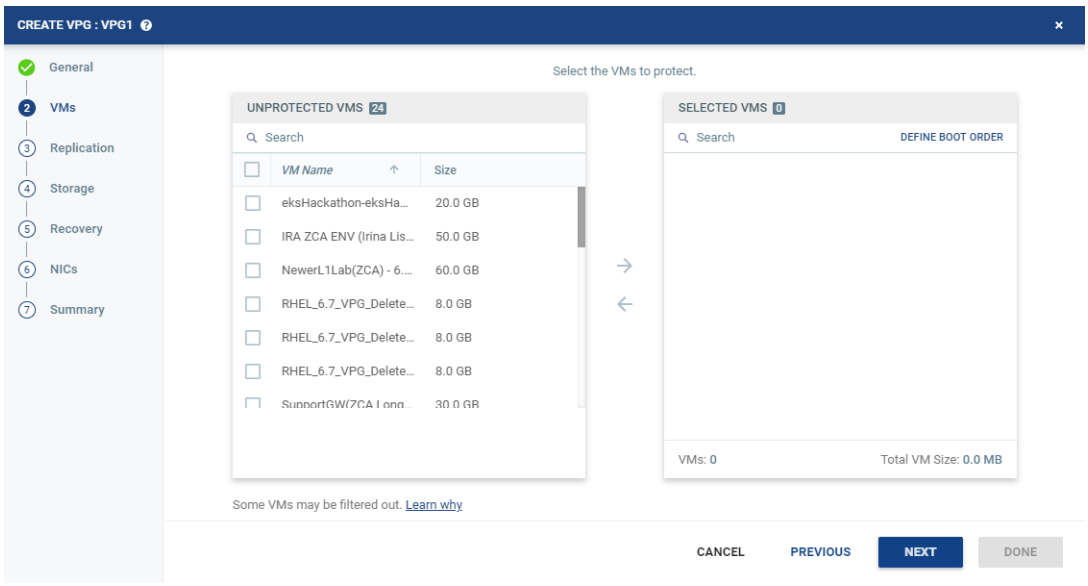

- The list of unprotected VMs includes VMs that were in the ZCA's **account and region** during installation.
- A VPG can include virtual machines that are already protected by another ZCA in the same account and region.
- The size of the disk is defined by data disks and the OS disk.
- AWS Instance Store disks (Temp disks) are not included in the list of unprotected VMs.

4. **Select** the VMs that will be part of this VPG and click the arrow to include these VMs in the VPG.
   - When using the **Search** field, you can use the wildcards; * or **?**
5. To define the boot order of the virtual machines in the VPG, click **Define Boot Order**, otherwise go to the next step.

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

   Note: Up to five (5) virtual machines may boot on a host simultaneously. Following the boot, a 300 second (default) delay occurs until the next boot batch.

   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.

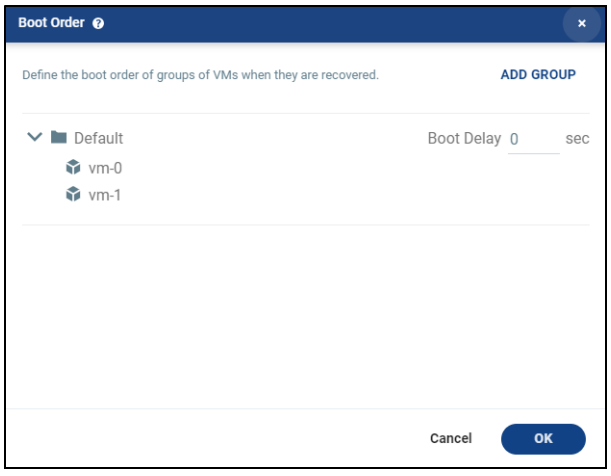

6. Click **ADD GROUP** to add a new group. Then, do the following:
   a) To change the name of a group, click the Pencil icon next to the group.
   b) To delete a group, click the delete icon on the right side. You cannot delete the Default group nor a group that contains a virtual machine.

c) Drag virtual machines to move them from one group to another.
d) Drag groups to change the order the groups are started, or, optionally, in **Boot Delay**, specify a time delay between starting up the virtual machines in the group and starting up the virtual machines in the next group.
For Example: Assume three groups, Default, Server, and Client, defined in this order. The boot delay defined for the Default group is 10, for the Server group is 100, and for the Client group 0. The virtual machines in the Default group are started together and after 10 seconds the virtual machines in the Server group are started. After 100 seconds the virtual machines in the Client group are started.
e) Click **OK** to save the boot order.

7. Click **NEXT**.

   The **REPLICATION** step is displayed.

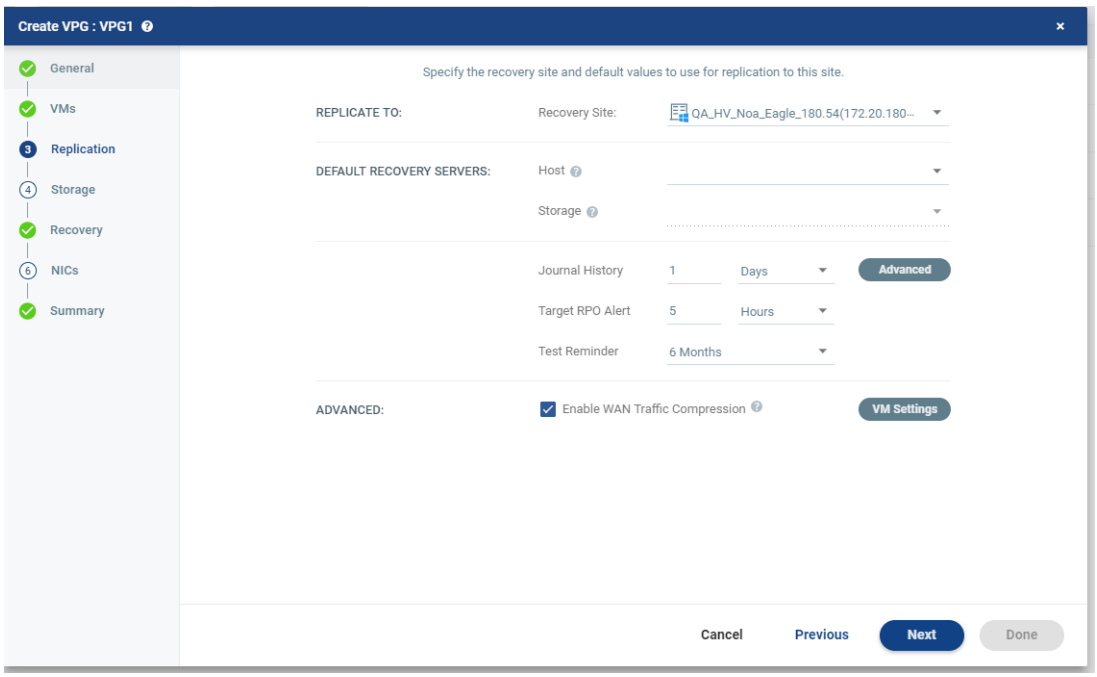

8. Specify the **Recovery Site**. This is the site to which you want to recover the virtual machines.
   - If the protected site is paired with **more than one** recovery site, the Recovery Site is left empty. Proceed to step 9.
   - If the protected site is paired with **only one** recovery site, and that recovery site is defined in **Zerto Cloud Manager**, the Recovery Site field is automatically filled in, as well as the defaults set for the SLA and Advanced settings. Proceed to step 10.
   - If the protected site is paired with **only one** recovery site, the Recovery Site field is automatically filled in, as well as the defaults set for the SLA and Advanced settings. Proceed to step 11.
9. Select the **Recovery Site** from the drop-down list.

   After specifying the Recovery Site, other fields are displayed including the Host and Storage to use for replication.
   - If the selected virtual machines are **already in VPGs** that recover to that site, select another recovery site. Proceed to step 11.
10. (**If the site is defined in** Zerto Cloud Manager) **ZORG:** Specify the name used by the **cloud service provider** to identify you as a Zerto Organization, ZORG.

    Note: If the site is defined in Zerto Cloud Manager, vSphere Standard edition **cannot** be used.

    For details about Zerto Cloud Manager, refer to *Zerto Cloud Manager Administration Guide*.

11. After selecting the recovery site, define the following parameters:

PARAMETER	FIELD	DESCRIPTION
Default Recovery Services	**Host**	The default **cluster**, **resource pool** or **host** in the recovery site that handles the replicated data. ■ If the site is defined in **Zerto Cloud Manager**, only a resource pool can be specified and the resource pool must also have been specified as a resource in Zerto Cloud Manager. ■ Zerto checks that the resource pool capacity is enough for any virtual machine specified in the VPG. ■ All resource pool checks are made at the VPG level and do not take into account multiple VPGs using the same resource pool. ■ If the resource pool CPU resources are specified as unlimited, the actual limit is inherited from the parent. If the inherited value is too small, failover, move, and failover test operations can fail, even without a warning alert being issued by Zerto Virtual Manager. ■ If a resource pool is specified and DRS is disabled for the site later on, all the resource pools are removed by VMware and recovery will be to any one of the hosts in the recovery site with a VRA installed on it.
	Storage	The default **cluster**, **resource pool** or **host** in the recovery site that handles the replicated data. If the site is defined in Zerto Cloud Manager, only a resource pool can be specified and the resource pool must also have been specified as a resource in Zerto Cloud Manager.
SLA	**Service Profile**	(Only if the site is defined in **Zerto Cloud Manager**) **Specify** the name of the service profile. The service profile determines the VPG SLA settings for the group, and is applied to **every virtual machine in the group**. To change the VPG SLA settings, select the **Custom** Service Profile. To change VPG SLA for individual virtual machines, proceed to step 12.
	Journal History	Specify the time-frame in which information is saved in the journal for the group, and which will apply to **every** virtual machine in the group. The longer the information is saved in the journal, the more space is required for each journal in the VPG. Select one of the following time-frames: ■ Number of **hours**: From **1** to **24** ■ Number of **days**: From **2** to **30** For **additional** journal-related fields, click **ADVANCED**, and proceed to step 12.
	Target RPO Alert	The maximum desired time between each automatic checkpoint write to the journal before an alert is issued. Default value is X hours.
	Test Reminder	The time recommended between testing the integrity of the VPG. A warning is issued if a test is not done within this time frame.
Advanced	**VM SETTINGS**	To change the replication settings **per virtual machine**, click **VM SETTINGS** button, and proceed to step 12.

12. For additional journal-related fields, click **ADVANCED**.
    The Advanced Journal Settings window is displayed.

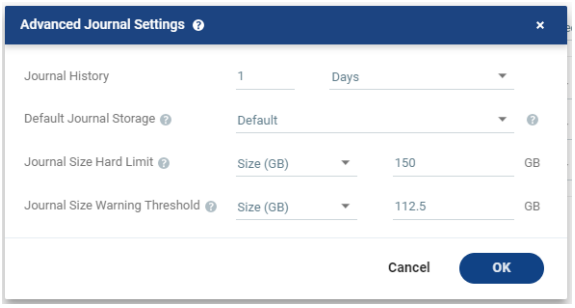

SETTING & DESCRIPTION	SELECT...
Journal History	
The time that all write commands are saved in the journal. The longer the information is saved in the journal, the more space is required for each journal in the VPG.	■ Number of **hours** from **1 to 23** ■ Number of **days** from **1 to 30**
Default Journal Storage (Hyper-V), or **Default Journal Datastore** (vSphere)	
The storage/datastore used for the journal data for each virtual machine in the VPG. **Note:** This field is **not** relevant when replicating to a **vCD** recovery site.	■ Select the storage/datastore accessible to the host. When you select a specific journal storage/datastore, the journals for each virtual machine in the VPG are stored in this storage/datastore, regardless of where the recovery storage/datastore is for each virtual machine. All protected virtual machines are recovered to the hosts that can access the specified journal storage/datastore.
Journal Size Hard Limit	

The maximum size that the journal can grow, either as a percentage or a fixed amount.
 	
The journal is always **thin-provisioned**.
 	
Note: The Journal Size Hard Limit applies independently **both** to the Journal History **and** also to the Scratch Journal Volume.
 	
For Example: If the Journal Size Hard Limit is configured to a maximum size of 160 GB limit, then during Failover Test, both the Journal History and the Scratch Journal Volume together can take up to 320 GB. Each one with a maximum size of 160 GB limit.
	■ **Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore.
	
	If Unlimited is selected, Size and Percentage options are **not** displayed.
	
	■ **Size (GB):** The maximum journal size in GB.
	  ■ The **minimum** journal size, set by Zerto, is **8GB** for Hyper-V and vSphere environments, and **10GB** for Microsoft Azure environments.
	
	■ **Percentage:** The percentage of the virtual machine volume size to which the journal can grow.
	  ■ This value can be configured to more than 100% of the protected VM's volume size.

SETTING & DESCRIPTION	SELECT...
Journal Size Warning Threshold	
The size of the journal that triggers a warning that the journal is nearing its hard limit.	■ **Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore. If Unlimited is selected, Size and Percentage options are **not** displayed. ■ **Size* (GB):** The size in GB that will generate a warning. ■ **Percentage*:** The percentage of the virtual machine volume size that will generate a warning. *The values of **Size** and **Percentage** must be **less** than the configured **Journal Size Hard Limit** so that the warning will be generated when needed. In addition to the warning threshold, Zerto will issue a message when the free space available for the journal is almost full.

13. Click **OK**.
14. To change the replication settings per virtual machine, click **VM SETTINGS**.
    The Advanced VM Replication Settings window is displayed.

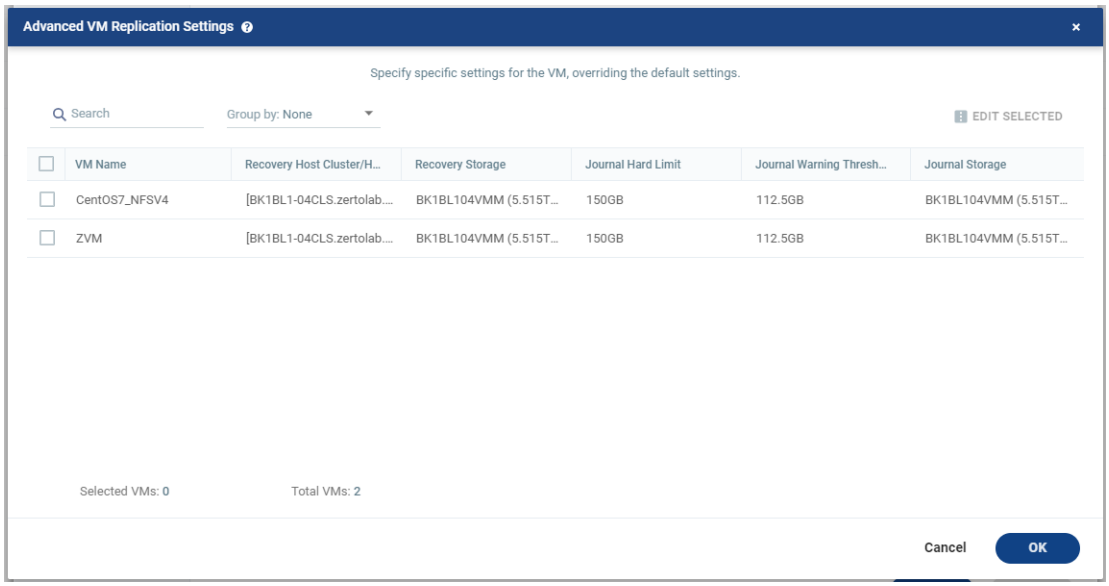

    In this window, you can edit the values of one or more of the virtual machines in the VPG.
15. To edit information in one field, **click the field** and update the information.
16. To edit information for several virtual machines at the same time, select the virtual machines and click **EDIT SELECTED**.

The Edit VM window is displayed.

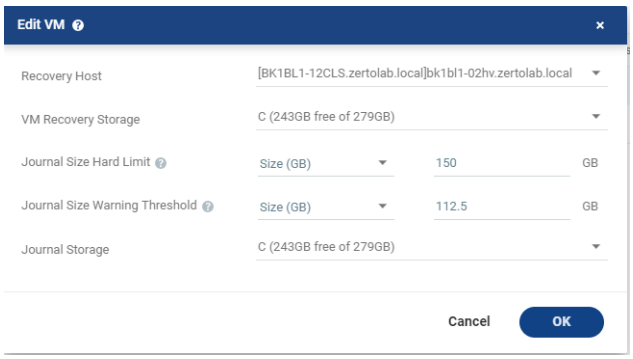

SETTING & DESCRIPTION	SELECT...
Recovery Host (not relevant when replicating to vCD)	
(Hyper-V) The cluster or host that will host the recovered virtual machine.	
(vSphere) The cluster, resource pool, or host that will host the recovered virtual machine. If the site is defined in Zerto Cloud Manager, only a resource pool can be specified and the resource pool must also have been defined in Zerto Cloud Manager. For details about Zerto Cloud Manager, see *Zerto Cloud Manager Administration Guide*. When a resource pool is specified, Zerto checks that the resource pool capacity is enough for all the virtual machines specified in the VPG	When a resource pool is specified, Zerto checks that the resource pool capacity is enough for all the virtual machines specified in the VPG. If a resource pool is specified and DRS is disabled for the site later on, all the resource pools are removed by VMware and recovery is to any one of the hosts in the recovery site with a VRA installed on it. All resource pool checks are made at the level of the VPG and do not take into account multiple VPGs using the same resource pool. If the resource pool CPU resources are defined as unlimited, the actual limit is inherited from the parent but if this inherited value is too small, failover, move, and failover test operations can fail, even without a warning alert being issued by Zerto Virtual Manager.
VM Recovery Datastore (vSphere) (not relevant when replicating to vCD)	
The datastore where the VMware metadata files for the virtual machine are stored, such as the VMX file.	If a cluster or resource pool is selected for the host, only datastores that are accessible by every ESX/ESXi host in the cluster or resource pool are displayed. This is also the datastore where RDM backing files for recovery volumes are located.
Recovery Storage (Hyper-V) (not relevant when replicating to vCD)	
The location where the metadata files for the virtual machine are stored, such as the VHDX file. I	If a cluster is selected for the host, only storage that are accessible by every host in the cluster are displayed.

SETTING & DESCRIPTION	SELECT...
Journal Size Hard Limit	
The maximum size that the journal can grow, either as a percentage or a fixed amount. ■ The journal is always **thin-provisioned**. ■ The Journal Size Hard Limit applies independently **both** to the Journal History **and** also to the Scratch Journal Volume. *For Example:* If the Journal Size Hard Limit is configured to a maximum size of 160 GB limit, then during Failover Test, both the Journal History and the Scratch Journal Volume together can take up to 320 GB. Each one with a maximum size of 160 GB limit.	**Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore. If Unlimited is selected, Size and Percentage options are **not** displayed.
	Size (GB): The maximum journal size in GB. ■ The **minimum** journal size, set by Zerto, is **8GB** for Hyper-V and vSphere environments, and **10GB** for Microsoft Azure environments.
	Percentage: The percentage of the virtual machine volume size to which the journal can grow. ■ This value can be configured to more than 100% of the protected VM's volume size.
Journal Size Warning Threshold	
The size of the journal that triggers a warning that the journal is nearing its hard limit.	**Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/datastore. If Unlimited is selected, Size and Percentage options are **not** displayed.
	Size* (GB): The size in GB that will generate a warning.
	Percentage*: The percentage of the virtual machine volume size that will generate a warning.
	***The values of **Size** and **Percentage** must be **less** than the configured **Journal Size Hard Limit** so that the warning will be generated when needed. In addition to the warning threshold, Zerto will issue a message when the free space available for the journal is almost full.
Journal Storage (Hyper-V), or **Journal Datastore** (vSphere) (not relevant when replicating to vCD)	
The storage/datastore used for the journal data for each virtual machine in the VPG.	**(vSphere)** To change the default, specify a host and then select one of the datastores accessible by this host to be used as the journal datastore. When you select specific journal datastore, the journals for each virtual machine in the VPG are stored in this datastore, regardless of where the recovery datastores are for each virtual machine. In this case, all the protected virtual machines must be recovered to hosts that can access the specified journal datastore. **(Hyper-V)** To change the default, specify a host and then select the storage location accessible by this host to be used as the journal storage. When you select specific journal storage, the journals for each virtual machine in the VPG are stored in this storage, regardless of where the recovery storage is for each virtual machine. In this case, all the protected virtual machines must be recovered to hosts that can access the specified journal storage.

17. Click **OK**.
18. In the Advanced VM Replication Settings window, click **OK**.

19. Click **NEXT**.

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

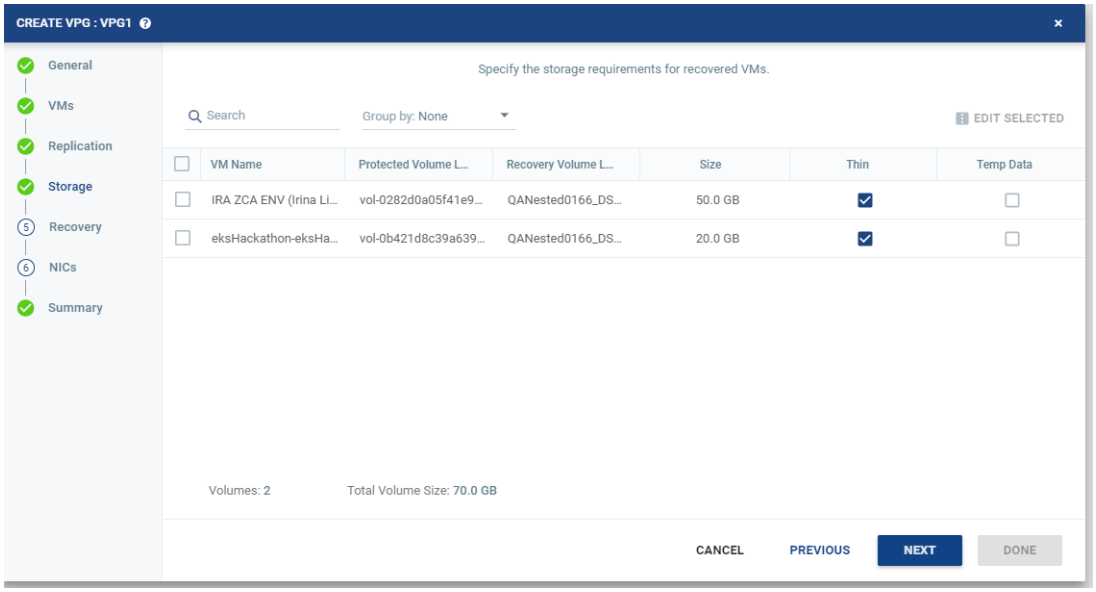

    You can define **Thin** provisioning and **Temp Data** in this window, **or** you can alternatively define them when you separately select and **edit each VMs volume**.

IMPORTANT:

Changing the VPG recovery volume from thin-provisioned to thick-provisioned or vice versa, results in **volume initial synchronization**.

See the following considerations regarding Thin provisioning:

- Unless the user **explicitly** requests Thin provisioning, provisioning type is the same type as provisioning in the **source VM**.
- If the **source** disk is Thin provisioned, the default for the **recovery** volume is also Thin provisioned.
- If the user uses preseed disks, Zerto maintains the provisioning types of the disks, so they can have other provisioning types.

PRESEED	PROVISIONING IN THE RECOVERY VM
Not selected	■ User **can** select Thin provisioning
Selected	■ User **cannot** select Thin provisioning ■ Provisioning is the **same** as defined in **source VMs**

20. To define whether the recovery volumes are thin-provisioned or not, select the **Thin** checkbox.
21. If the virtual machine to be replicated includes a temp data disk as part of its configuration, select the **Temp Data** checkbox to mark the recovery disk for this disk as a temp data disk. In this case, data is not replicated to the temp data disk after initial synchronization.

22. To edit storage information for one of the virtual machines' volume location, first select the virtual machine, then click **EDIT SELECTED**. The Edit Volumes window is displayed.

    - For **Hyper-V** recovery environments, the following window appears. For details, click here.

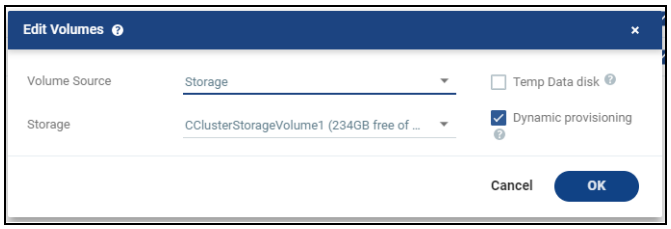

- For **vSphere** recovery environments, the following window appears. For details, click here.

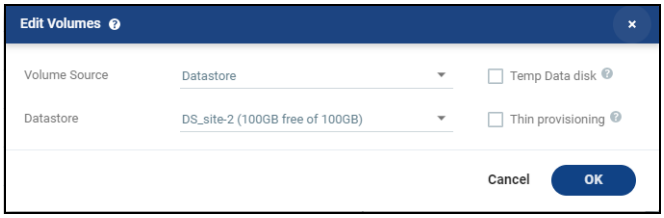

- For **vCD** recovery environments, the following window appears. For details, click here.

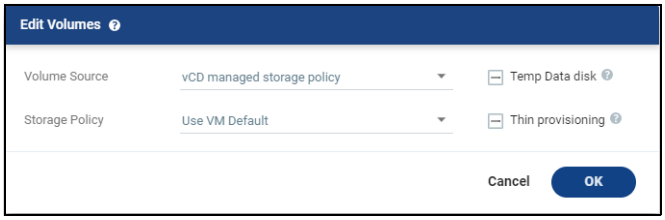

SETTING & DESCRIPTION	SELECT...
Volume Source	
■ (Hyper-V) Select a **Volume Source** for recovery from one of the drop-down options: ■ **Storage** ■ **Preseeded volume**	■ **Volume Source > Storage:** A new volume is used for replicated data. ■ From the **Storage** drop-down list, specify the storage to use to create disks for the replicated data. The storage specified for the replication must have at least the same amount of space as the protected volume and then an additional amount for the journal. The amount of additional space needed for the journal can be fixed by specifying a maximum size for the journal, or can be calculated as the average change rate for the virtual machines in the VPG, multiplied by the length of time specified for the journal history. For more details, see *Zerto Scale and Benchmarking Guidelines*, in the section **Estimating WAN Bandwidth for VMware and Hyper-V**.
	■ **Volume Source > Preseeded volume:** Whether to copy the protected data to a virtual disk in the recovery site. Zerto **recommends** using this option particularly for large disks so that the initial synchronization will be faster since a **Delta Sync** can be used to synchronize any changes written to the recovery site after the creation of the preseeded disk. When **not** using a preseeded disk, the initial synchronization phase must copy the whole disk over the WAN. When using a preseeded virtual disk, you select the storage and exact location, folder, and name of the preseeded disk. Zerto takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. Only disks with the same size as the protected disk can be selected when browsing for a preseeded disk. The storage where the preseeded disk is placed is also used as the recovery storage for the replicated data.

SETTING & DESCRIPTION	SELECT...
■ (vSphere) Select a **Volume Source** for recovery from one of the drop-down options: ■ **Datastore** ■ **RDM** ■ **Preseeded volume**	**Volume Source > Datastore:** A new volume is used for replicated data. 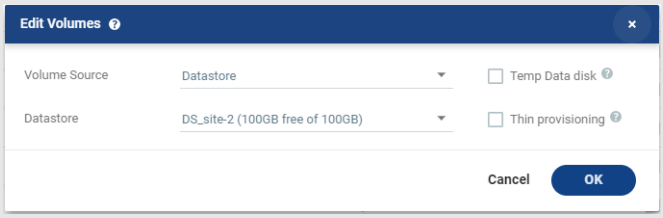 ■ Specify the **Datastore** to use to create disks for the replicated data. ■ If the **source disk is thin provisioned**, the default for the recovery volume is also thin provisioned. ■ The datastore specified for replication must have at least the same amount of space as the **protected volume** and an additional amount for the **journal**. ■ The amount of additional space needed for the journal can be fixed by specifying a maximum size for the journal, or can be calculated as the average change rate for the virtual machines in the VPG, multiplied by the length of time specified for the journal history. ■ Zerto supports the SCSI protocol. Only disks that support this protocol can be specified. Then, define the following: ■ **Datastore:** The Datastore where the preseeded disk is located. Only disks with the same size as the protected disk can be selected when browsing for a preseeded disk. For more details, see *Zerto Scale and Benchmarking Guidelines*, in the section **Estimating WAN Bandwidth for VMware and Hyper-V**.
	Volume Source > RDM: The VMware RDM (Raw Device Mapping) which will be used for the replication. 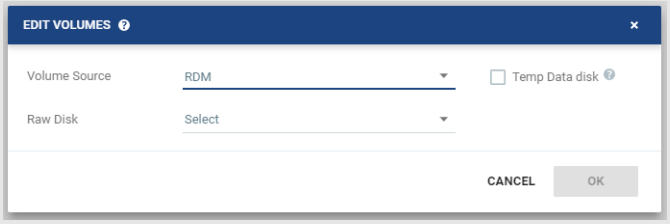 By default, **RDM is recovered as thin-provisioned VMDK** in the datastore specified in the **VM Recovery Datastore/Storage** field in the Edit VM dialog, and not to RDM. Only a raw disk with the **same size as the protected disk** can be selected from the list of available raw disks. Other raw disks with different sizes are not available for selection. The RDM is always stored in the recovery datastore, used for the virtual machine. The following **limitations** apply to protecting RDM disks: ■ RDM disks with an even number of blocks can replicate to RDM disks of the same size with an even number of blocks and to VMDKs. ■ RDM disks with an odd number of blocks can only replicate to RDM disks of the same size with an odd number of blocks and not to VMDKs. ■ You cannot define an RDM disk to be protected to a cloud service provider via a Zerto Cloud Connector nor if the virtual machine uses a BusLogic SCSI controller, nor when protecting or recovering virtual machines in an environment running vCenter Server 5.x with ESX/ESXi version 4.1 hosts.

SETTING & DESCRIPTION	SELECT...
■ (vSphere) **Volume Source** continued	**Volume Source > Preseeded volume**: Select this when you want to copy the protected data to a virtual disk in the recovery site. 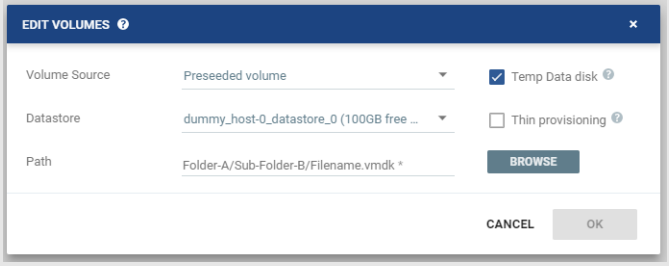 Consider the following, then proceed to define the Datastore and the Path: ■ Zerto **recommends** using this option particularly for **large disks** so that the initial synchronization is **faster** since a Delta Sync can be used to synchronize any changes written to the recovery site after the creation of the preseeded disk. ■ If a preseeded disk is **not** selected, the initial synchronization phase must copy the **whole disk** over the WAN. ■ If you use a preseeded virtual disk, you select the datastore and exact location, folder, and name of the preseeded disk, which cannot be an IDE disk. Zerto takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. ■ The datastore where the preseeded disk is placed is also used as the recovery datastore for the replicated data. ■ If the preseeded disk is **greater than 1TB on NFS storage**, the VPG creation might fail. This is a known VMware problem when the NFS client does not wait for sufficient time for the NFS storage array to initialize the virtual disk after the RPC parameter of the NFS client times out. The timeout default value is 10 seconds. See VMware documentation, http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1027919, which describes the configuration option to tune the RPC timeout parameter by using the command: `esxcfg-advcfg -s <Timeout> /NFS/SetAttrRPCTimeout` ■ If the protected disks are **non-default geometry**, configure the VPG using preseeded volumes. ■ If the protected disk is an **RDM disk**, it can be used to preseed to a recovery VMDK disk. Zerto makes sure that the VMDK disk size is a correct match for the RDM disk. ■ If the VPG is being defined for a Zerto Organization, ZORG, the location of the preseeded disk must be defined in the Zerto Cloud Manager. See *Zerto Cloud Manager Administration Guide*. Then, define the following: ■ **Datastore:** The Datastore where the preseeded disk is located. Only disks with the same size as the protected disk can be selected when browsing for a preseeded disk. ■ **Path:** The full path to the preseeded disk.

SETTING & DESCRIPTION	SELECT...
■ (vCD) Select a **Volume Source** for recovery from one of the drop-down options: ■ **vCD managed storage policy** ■ **Preseeded volume**	**Volume Source > vCD managed storage policy:** Zerto will select a datastore, from the list of available datastores, in the selected Storage Policy in which to place the Volume, unless the datastore is excluded in the Configure Provider vDCs Dialog. ■ If there are several valid datastores, the datastore with the most available space is selected. ■ Zerto recalculates the datastore available space for each volume sequentially, taking into consideration previously allocated volumes. **Volume Source > Preseeded volume:** Select this when you want to copy the protected data to a virtual disk in the recovery site. ■ Zerto **recommends** using this option particularly for **large disks** so that the initial synchronization is **faster** since a Delta Sync can be used to synchronize any changes written to the recovery site after the creation of the preseeded disk. ■ If a preseeded disk is **not** selected, the initial synchronization phase must copy the **whole disk** over the WAN. ■ The preseeded volume is a virtual disk (the VMDK flat file and descriptor) in the recovery site that has been prepared with a copy of the protected data. ■ Browse to the preseed folder configured for the customer and the disk name, of the preseeded disk. In order to use a preseeded VMDK, do the following: a) Create a folder in vCD to use for the preseeded disks in the datastore you want to use for the customer. b) Specify this datastore as a provider datastore for preseeded disks in the Configure Provider vDCs window, from the Advanced Settings window, as described in *Zerto Cloud Manager Administration Guide*. c) In the Zerto Cloud Manager specify the Preseed Folder Name for the ZORG, in the Manage ZORG tab. ■ Zerto searches for the preseeded folder in the available datastores in the Org vDCs specified in the vCD Cloud Resources for the ZORG in the Zerto Cloud Manager and takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA. If the virtual machine has more than one preseeded disk, these disks must reside on the same datastore. ■ If the preseeded disk is **greater than 1TB on NFS storage**, the VPG creation might fail. This is a known VMware problem when the NFS client does not wait for sufficient time for the NFS storage array to initialize the virtual disk after the RPC parameter of the NFS client times out. The timeout default value is 10 seconds. See VMware documentation, http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1027919, which describes the configuration option to tune the RPC timeout parameter by using the command: **esxcfg-advcfg -s <Timeout> /NFS/SetAttrRPCTimeout** ■ If the VPG is being defined for a Zerto Organization, ZORG, the location of the preseeded disk must be defined in the Zerto Cloud Manager. See *Zerto Cloud Manager Administration Guide*. ■ Zerto supports the SCSI protocol. Only disks that support this protocol can be specified. Virtual machine RDMs in a vCenter Server are replicated as VMDKs in a vCD environment.

SETTING & DESCRIPTION	SELECT...
(vCD) continued	**Storage Policy**: Specify the Storage Policy for recovery from one of the options: ■ Storage Policy per volume is supported only in vCD supported versions, and when the selected Orgvdc is not configured for fast provisioning. ■ Zerto will select a datastore from the selected Storage Policy in which to place these files, unless the datastore is excluded in the Configure Provider vDCs Dialog. ■ The Storage Policies which appear in the drop-down list: ■ Include the **Use VM Default** option (default), which will apply the VM's storage policy to this volume. This is also the Storage Policy default value. ■ Were defined in VMware vCloud Director and are configured in the Orgvdc. ■ Have at least one Datastore that was not excluded as a Recovery Volume in the Configure Provider vDCs Dialog.
Temp Data disk	
If the virtual machine to be replicated includes a temp data disk as part of its configuration.	Specify a mirror disk for replication that is marked as a temp data disk. In this case, data is not replicated to the temp data disk after initial synchronization.
Thin provisioning (vSphere)	
If the recovery volumes are thin-provisioned or not.	If the source disk is thin provisioned, the default for the recovery volume is also thin provisioned. vCD only: Unless the Org vDC only supports thin-provisioned volumes
Dynamic provisioning (Hyper-V)	
If the recovery volumes are dynamic-provisioned or not.	If the source disk is dynamic provisioned, the default for the recovery volume is also dynamic provisioned. vCD only: Unless the Org vDC only supports dynamic-provisioned volumes

23. Click **OK**.
24. Click **NEXT**.

    The RECOVERY step is displayed. Recovery details include the networks to use for failover, move, and for testing failover, and whether scripts should run as part of the recovery operation.

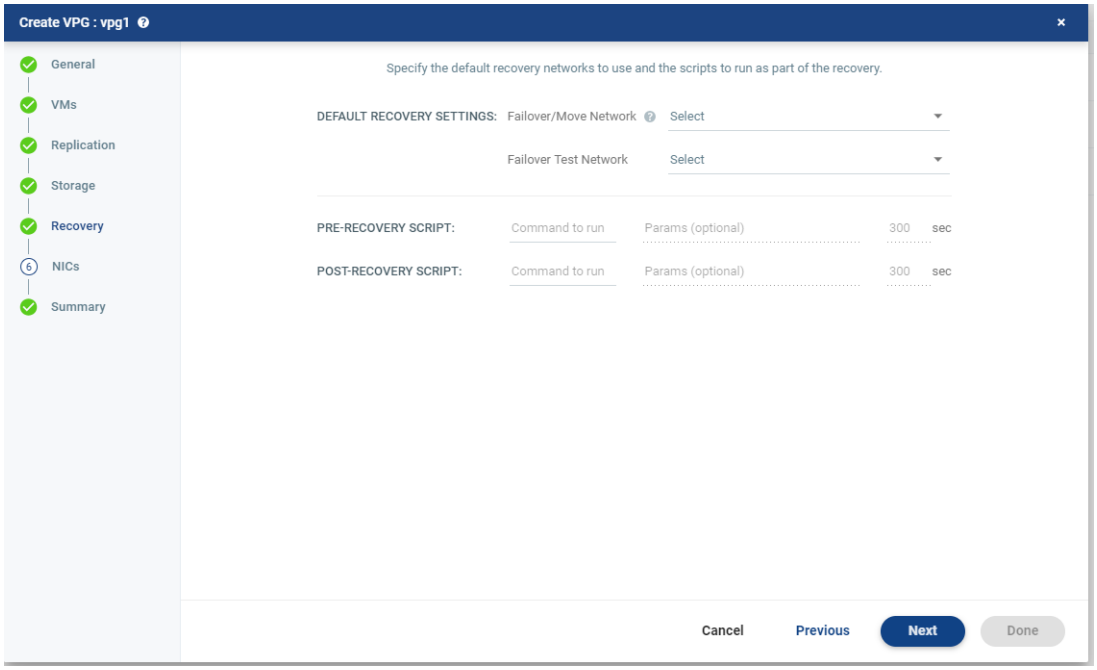

25. Select the default recovery settings. There are applied to every virtual machine in the VPG.
    - **Failover/Move Network:** The network to use during a failover or move operation in which the recovered virtual machines will run.
    - **Failover Test Network:** The network to use when testing the failover of virtual machines in the recovery site. Zerto recommends using a fenced-out network so as not to impact the production network at this site.
26. To run pre and post recovery scripts, enter the name of the script to run in the Command to run text box. You can then enter details about the script.
    - **Pre-recovery Script:** The information about a script that should run at the beginning of the recovery process.
    - **Post-recovery Script:** The information about a script that should run at the end of the recovery process.

    For both types of scripts, enter the following information:

TEXT BOX	DESCRIPTION
Command to run	The full path of the script. The script must be located on the same machine as the Zerto Virtual Manager for the recovery site.
Params	The parameters to pass to the script. Separate parameters with a space.
Timeout	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 the 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 time-out value is specified in **Site Settings > Performance and Throttling** tab.

NOTE:

Pre and post recovery scripts run in parallel. Therefore, ensure that the pre and post recovery scripts don't use common resources.

27. Click **NEXT**.

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

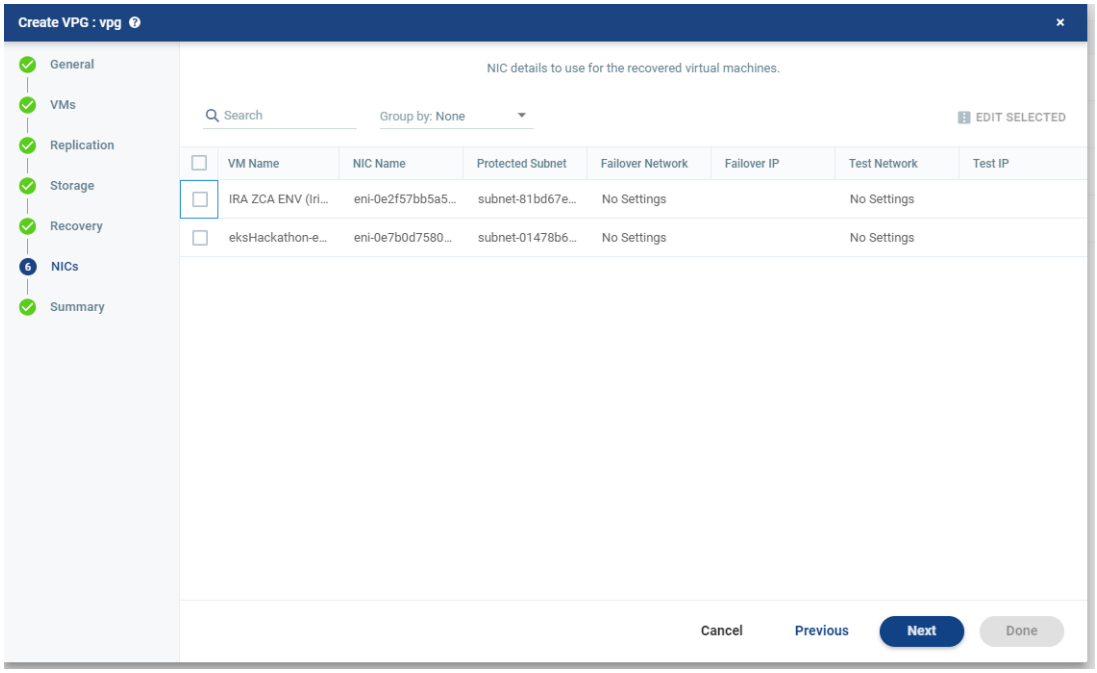

28. To edit information in one field, click the field and update the information.
29. To edit information for several virtual machines at the same time, select the virtual machines and click **EDIT SELECTED**.
30. Otherwise, go to step 33.

    The Edit VNIC window is displayed.

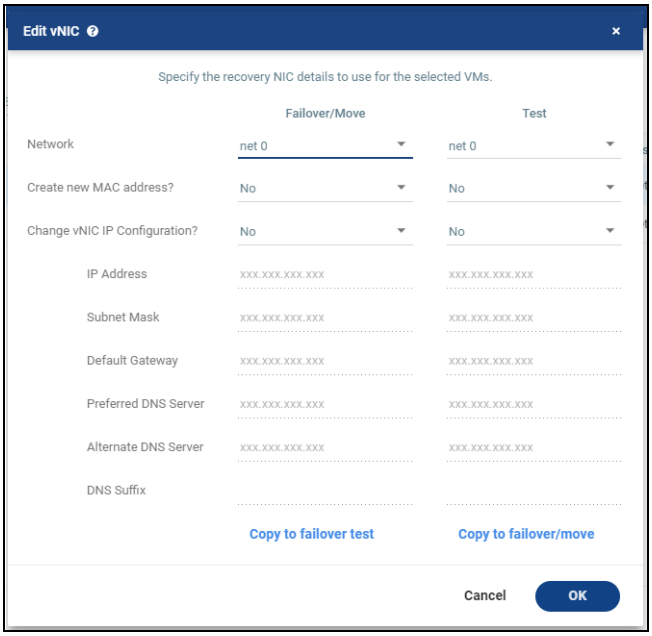

Specify the network details to use for the recovered virtual machines after a failover or move operation, in the Failover/Move column, and for the recovered virtual machines when testing replication, in the Test column.

31. In each column, **specify** the following:
    - **Network:** The network to use for this virtual machine.
    - **Create new MAC address?:** Whether the Media Access Control address (MAC address) used on the protected site should be replicated on the recovery site. The default is to use the same MAC address on both sites. Note that if you

check this option, to create a new MAC address, and the current IP address is not specified, the protected virtual machine static IP address might not be used for the recovered virtual machine.

- **Change vNIC IP Configuration?**: Whether or not to keep the default virtual NIC (vNIC) IP configuration. The vNIC IP is only changed after recovery for virtual machines with VMware Tools running.
  See the Zerto Interoperability Matrix for the list of operating systems for which Zerto supports Re-IPing.
  - To change the vNIC IP, in the Failover/Move or Test column, select **Yes**. If you select to use a static IP connection, set the **IP address**, **subnet mask**, and **default gateway**.
  - Optionally, change the preferred and **alternate DNS server IPs** and the DNS suffix.
  - If you leave the DNS server and suffix entries empty, or select to use DHCP, the IP configuration and DNS server configurations are assigned automatically, to match the protected virtual machine. You can change the DNS suffix.
  - If the virtual machine has multiple NICs but is configured to only have a single default gateway, fill in a **0** for each octet in the Default gateway field for the NICs with no default gateway.

  During a failover, move, or test failover, if the recovered virtual machine is assigned a different IP than the original IP, then after the virtual machine has started, it is automatically rebooted so that it starts up with the correct IP. If the same network is used for both production and test failovers, Zerto recommends changing the IP address for the virtual machines started for the test, so that there is no IP clash between the test machines and the production machines.
- **Copy to failover test:** Select this to copy the settings in the Failover/Move column to the **Test** column.
- **Copy to failover/move:** Select this to copy the settings in the Test column to the **Failover/Move** column.

32. Click **OK**.
33. Click **NEXT**.

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

For details of what happens after saving the VPG, see "What Happens After the VPG is Defined", on page 35.

Replication From a Protected Site AWS to a Recovery Site Azure

You can protect virtual machines from AWS to Microsoft Azure. The procedure is the same whether you intend to protect one virtual machine or multiple virtual machines.

When creating a VPG from AWS to Azure, the data is stored in a storage account and all replicated data from protected virtual machines to Azure is encrypted in the storage account. All recovery operations bring up the recovered machines in resource groups in Azure.

See also:

- "Requirements for Microsoft Azure Environments", on page 94
- "VPGs Recovering to Azure Standard Storage and Premium Managed Disks", on page 94
- "Converting Premium Virtual Machines for Protection", on page 94
- "Protecting From Amazon Web Services - To a Microsoft Azure Recovery Site", on page 97

Requirements for Microsoft Azure Environments

Before protecting your virtual machines, review Zerto Prerequisites and Requirements for Microsoft Azure Environments.

VPGs Recovering to Azure Standard Storage and Premium Managed Disks

Azure recovery of the VPG can be configured to **Premium Managed** or **Standard Storage**.

Based on the VPG configuration, VMs can be recovered to either Standard storage, Premium Managed disks or both.

- **Recovering VMs to the Standard Storage Account:**

  A VPG that all its virtual machines are recovered to the Standard storage account, as defined during the ZCA installation, can be **Reverse Protected**. For more information about Reverse Protection, see the *Reverse Protection for a Moved VPG* section.
- **Recovering VMs to Premium Managed Disks**:

  A VPG that all its virtual machines are recovered to Premium Managed disks goes into **Recovered** state. To protect a virtual machine with Premium Managed disks, see the *Converting Premium Virtual Machines for Protection* section.
- **Recovering VMs to Both Standard Storage and Premium Managed Disks:**

  In the case where a VPG has virtual machines recovering to both Standard storage and Premium Managed disks, the recovered VPG goes into a **Needs Configuration** state. In this scenario, you can remove the virtual machines that are recovered to Premium Managed disks in order to protect the virtual machines recovered to Standard storage.

  To remove virtual machines recovered to Premium Managed disks:
  a) Open the **Edit VPG** window > **VMs** tab.
  b) Remove the VMs that are in Premium Managed disks. (These VMs appear as grayed out.)
  c) Click **Edit Save**; the VPG is updated.

Converting Premium Virtual Machines for Protection

The **Premium to Standard Conversion** tool enables replicating and failing over Premium VMs in Azure. The tool clones the Premium Managed VMs and creates the same VM with Standard Storage disks, which the ZCA will then be able to protect.

The Premium to Standard Conversion tool is installed on the Azure ZCA machine as part of the ZVM installation.

The tool requires the user to enter input parameters which are then parsed and validated. Once the input parameters are validated, the application starts running. The conversion process then begins for each of the Premium VMs received in the input parameters.

GLOSSARY:

Premium VM: A Premium VM is when the original VM has all its disks on **Premium Managed storage**.

Note: We do not support conversion of VMs with disks on Premium Storage account.

Standard VM: A Standard VM is a copy made of the Premium VM, where the disks are copied from Premium Managed storage to the **ZCA's Standard Storage account**.

IMPORTANT:

Note that running the Premium to Standard Conversion tool will cause your Premium VM to be powered off.

To convert Premium VMs for protection:

1. From the Azure ZCA machine, locate the **Premium to Standard Conversion tool** folder from under the main ZVM installer folder.
2. Download and extract the Convert Premium VM tool files.
3. Copy the name of the tool (exe. file).
4. Open the **Command Prompt** window and paste the name of the tool (exe. file) to the command.
5. There are two ways to insert parameters; either as flag arguments or as a path to a .json file that contains the input parameters.

   The following parameters need to be entered:
   - **User Name & Password**: These are the same credentials used in the ZCA installation.
   - **Region Id**: The region in which the VM to be converted resides.
   - **Subscription Id**: The subscription in which the VM to be converted resides.
   - **VM Identifiers**: The list of VMs that need to be converted. If more than one VM needs to be converted, the VM Identifiers should be separated with commas.
     The VM identifier should be copied from the Premium VM properties (**Properties** tab > **Resource ID**.)
   - **Storage Account Name**: The storage account as defined for the ZCA.
   - **Container Name**: The name of the container in the Standard Storage Account where the new volumes will be stored.

NOTE:

You can view argument examples within the Command Prompt window by entering the help argument (e.g. '-h')

FLAG ARGUMENTS RUNNING EXAMPLE:

```
-u username@microsoft.com -p 12345678 -s f241xxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx -r "West Europe" --vms
/subscriptions/f241xxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx/resourceGroups/resourceGroupName/providers/Micros
oft.Compute'virtualMachines/Win2016Vm,subcriptions/f241xxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx/resourceGroups
/resourceGroupName/providers/Microsoft.Compute/virtualMachines/WindowsS2016 -sa storageaccountname
-c containername
```

.JSON FILE EXAMPLE:

Path to a .json file argument:

ConvertPremiumVm C:\jsonFolder\inputFile.json

(If the .json file within the folder contains the .exe file of the application, you can just insert the name of the .json file, without the application name)

.json file content:

```
{
"UserName": "name@gmail.com",
"Password": "123456",
  "SubscriptionId": "f241xxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
"RegionId": "West Europe",
  "VMsIdentifiers": [

"/subscriptions/f241xxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx/resourceGroups/namerglocal/providers/Microsoft.Compute/virtualMachines/PremiumVmTest",

"/subscriptions/f241xxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx/resourceGroups/namerglocal/providers/Microsoft.Compute/virtualMachines/PremiumVmTest1"
  ],
  "StorageAccountName" : "namelocal1",
  "ContainerName" : "vhds"
}
```

NOTE:

Do not perform any actions on the VMs that are converting until the conversion process is completed.

6. The information is parsed and the log files are printed and saved to the **Premium to Standard Conversion** tool folder. The conversion process is performed for each VM sequentially. The following occurs:
   - The Premium VM is turned off (if necessary).
   - The VM's volumes are copied into the Standard Storage account as defined in the Storage Account Name parameter. The copy is done sequentially.
   - The Premium VM NICs are detached and receives a NIC with a single IP, regardless of its original NIC/IP configuration.
   - A new VM is created in the Standard Storage account, with the identical NICs as defined in the original VM.
   - The name of the new VM is "<original_name>-Standard" with the new standard volumes attached to it.
7. Create a VPG with the new Standard Storage VMs to start protecting them. See *Protecting From a Microsoft Azure Site - To a Microsoft Azure Recovery Site*.

NOTE:

Upon failing over the Standard VM, only one IP address is assigned to the recovered VM. Any additional network configurations should be entered manually.

Undo Process

If there is a failure in one of the phases during the conversion process, the system will automatically rollback to the previous state.

If at least one volume cannot be copied, all volumes that have already been copied to the Standard VM are deleted. The Premium VM is then powered on again (only if the VM was powered on in its original state).

In some cases the system won't succeed to automatically rollback, for example, if the tool's application crashes. You can open the log entries to see which undo operations are left to execute. In this case, you need to manually rollback to the previous state, as required at each stage, depending where the failure occurred.

To manually rollback:

1. Detach the NICs from the Standard VM and attach it to the powered down Premium VM.
2. Delete the new primary NIC that was created for the Premium VM.
3. Remove disks from the Standard Storage account.
4. Power on the Premium VM (if necessary).

Protecting From Amazon Web Services - To a Microsoft Azure Recovery Site

To create a virtual protection group (VPG) to recover in Azure:

1. In the Zerto User Interface, select **ACTIONS > CREATE VPG**.
   The GENERAL step of the Create VPG wizard is displayed.

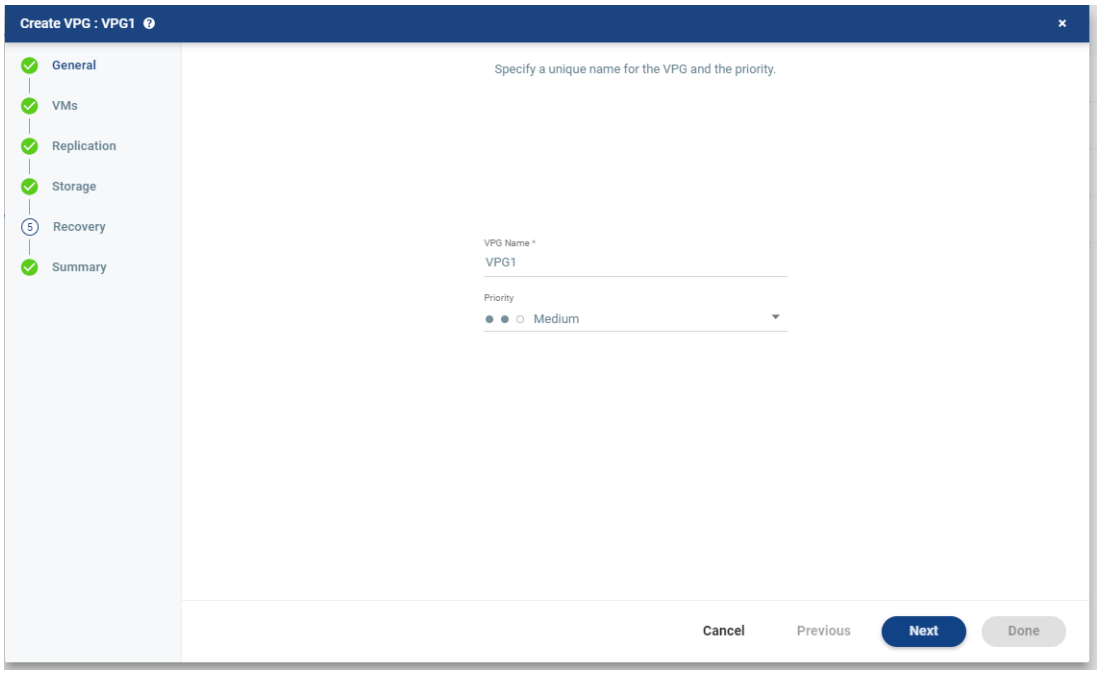

2. Specify the name of the VPG and the priority of the VPG.
   - **VPG Name:** The VPG name must be unique. The name cannot be more than 80 characters.
   - **Priority:** Determine the priority for transferring data from the protected site to the recovery site when there is limited bandwidth and more than one VPG is defined on the protected site.
     - **High Priority**: When there are updates to virtual machines protected in VPGs with different priorities, updates from the VPG with the highest priority are passed over the WAN first.
     - **Medium Priority:** Medium priority VPGs will only be able to use whatever bandwidth is left after the high priority VPGs have used it.
     - **Low Priority:** Low priority VPGs will use whatever bandwidth is left after the medium priority VPGs have use it.

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

3. Click **NEXT**.

   The VMs step is displayed.

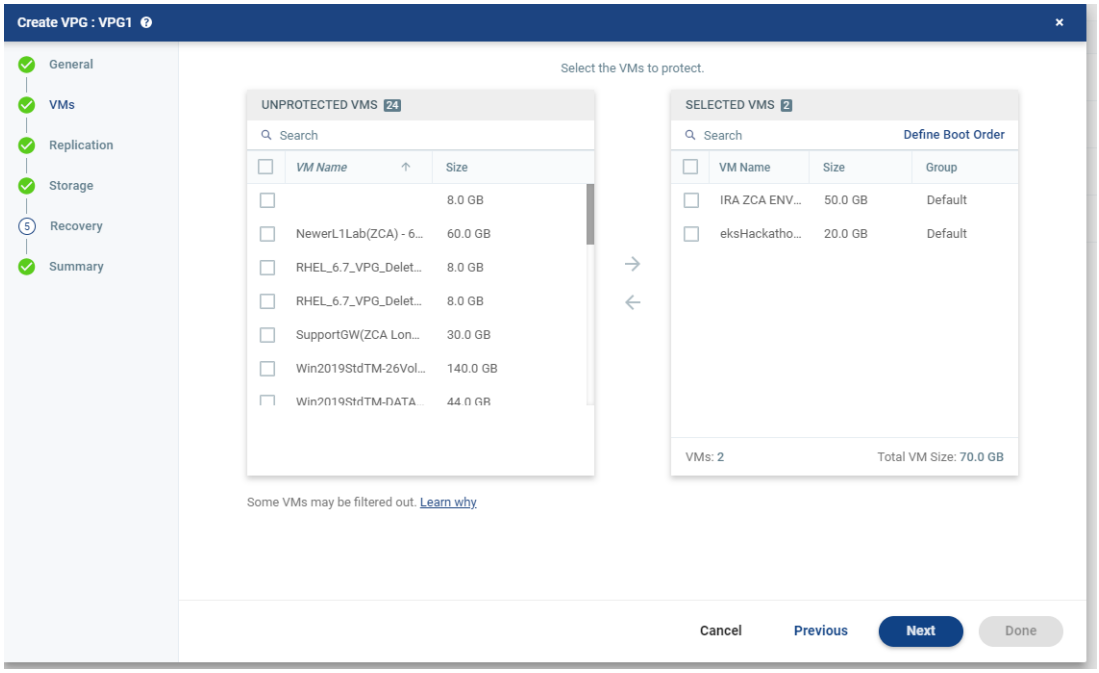

- The list of unprotected VMs includes VMs that were in the ZCA's **account and region** during installation.
- A VPG can include virtual machines that are already protected by another ZCA in the same account and region.
- The size of the disk is defined by data disks and the OS disk.
- AWS Instance Store disks (Temp disks) are not included in the list of unprotected VMs.

4. From the **Unprotected VMs** list, select the VMs to include in this VPG and click the arrow to move them to the **Selected VMs**.
   - When using the **Search** field, you can use the wildcards; ***** or **?**
5. Click **NEXT**.

   The **REPLICATION** step is displayed.
6. Select the **Recovery Site** to which you want to recover the virtual machines.

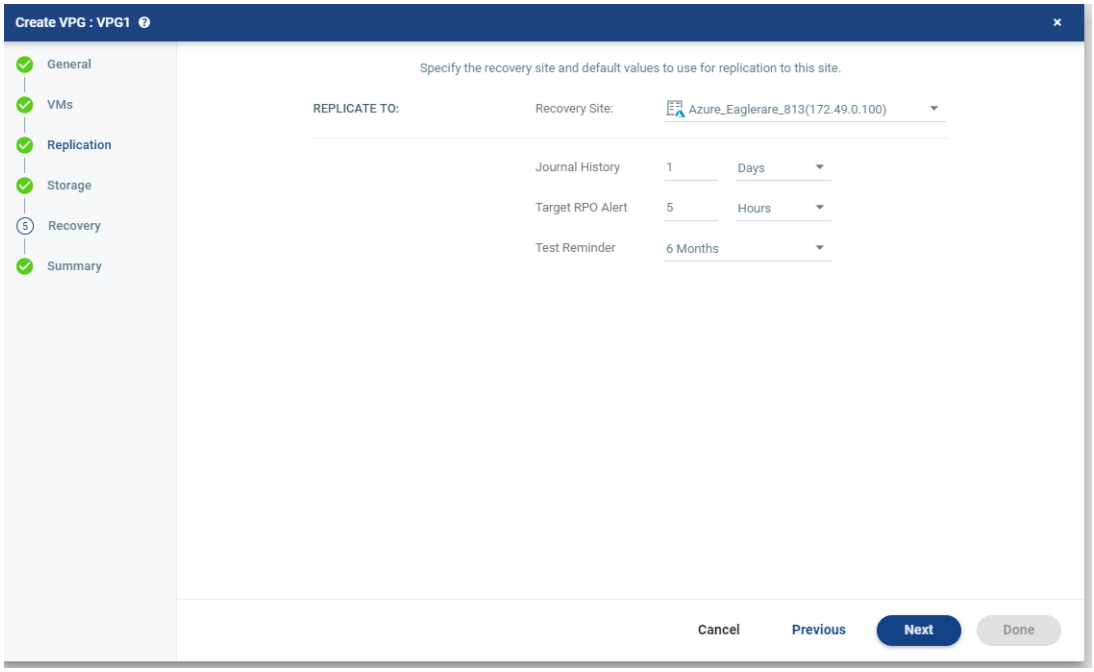

7. After selecting the recovery site, define the following parameters:

PARAMETER	FIELD	DESCRIPTION
SLA	**Journal History**	Specify the time-frame in which information is saved in the journal for the group, and which will apply to **every** virtual machine in the group. The longer the information is saved in the journal, the more space is required for each journal in the VPG. Select one of the following time-frames: ■ Number of **hours**: From **1** to **24** ■ Number of **days**: From **2** to **30** For **additional** journal-related fields, click **ADVANCED**, and proceed to step 12.
	Target RPO Alert	The maximum desired time between each automatic checkpoint write to the journal before an alert is issued. Default value is X hours.
	Test Reminder	The time recommended between testing the integrity of the VPG. A warning is issued if a test is not done within this time frame.

8. Click **NEXT**.

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

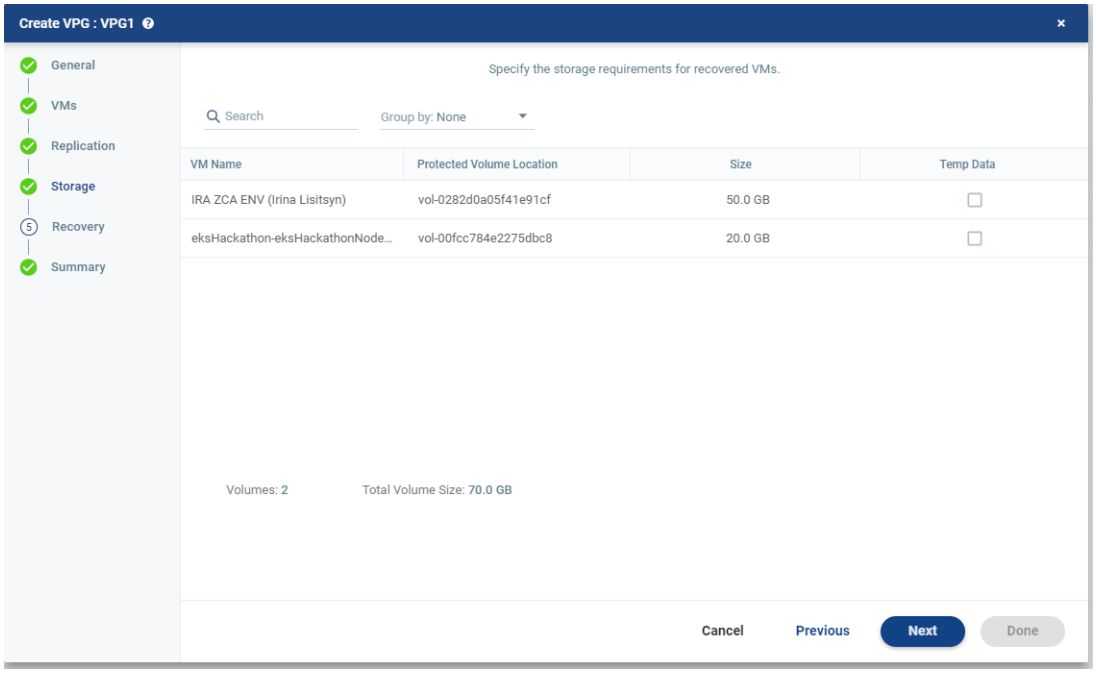

9. Specify whether the protected volume is a temp data disk.

   Temp Data: If the virtual machine to be replicated includes a temp data disk as part of its configuration, mark the recovery disk for this disk as a temp data disk. In this case, data is not replicated to the temp data disk after initial synchronization.

10. Click **NEXT**.

    The **RECOVERY** step is displayed. Recovery details include the networks to use for failover, move, and for testing failover, and whether scripts should run as part of the recovery operation.

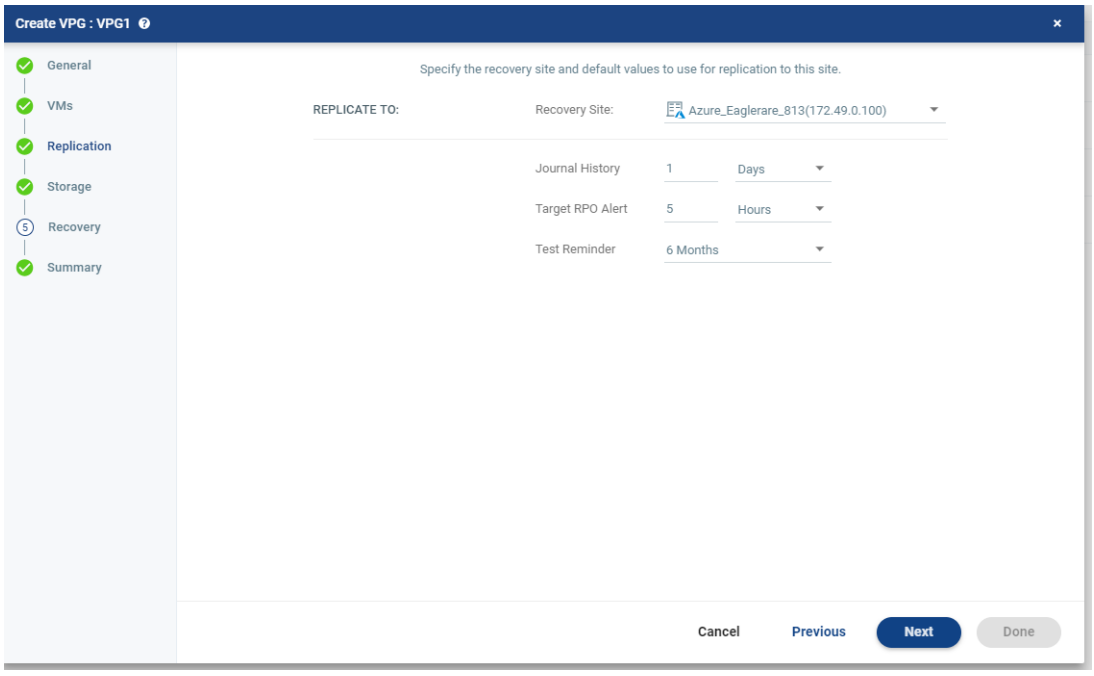

11. Select the default recovery settings. There are applied to every virtual machine in the VPG.

    - **VNet:** The virtual network dedicated to your Azure subscription.
    - **Subnet**: The subnet or the VNet network.

- **Network Security Group:** The Azure network security to be associated with the virtual machines in this VPG. You can associate one network security group with the virtual machines. The NIC will be associated with the network security group defined at the virtual machine level.
- **Recovery Disk Type:** Select the Azure recovery storage type to which the entire VPG will be recovered to; Premium Managed or Standard Storage. The Virtual Machine Series and Virtual Machine Size fields are updated with the relevant options based on the selected Recovery Disk Type.

  Note: To protect **Premium Managed** disks, see **Converting Premium Virtual Machines for Protection**.
- **Virtual Machine Series:** The virtual machine series from which to select the size. Azure virtual machine series are optimized for different types of applications. The default is set to **DSv2**. You can choose the virtual machine series appropriate for the application being protected in the VPG.
- **Virtual Machine Size:** The virtual machine size, within the virtual machine series, to assign to recovered virtual machines. Different sizes within a virtual machine series vary, for example in a number of cores, RAM, and local storage size. The default is set to **Standard_DS1_v2**. You can choose the virtual machine size appropriate for the application being protected in the VPG. The price per virtual machine is related to the virtual machine configuration.

12. For additional settings, click **ADVANCED VM SETTINGS**.

    The **Advanced VM Settings** window is displayed, which shows the recovery network settings for failover and move for virtual machines in the VPG. You can view the recovery network settings for failover tests by clicking **TEST**.

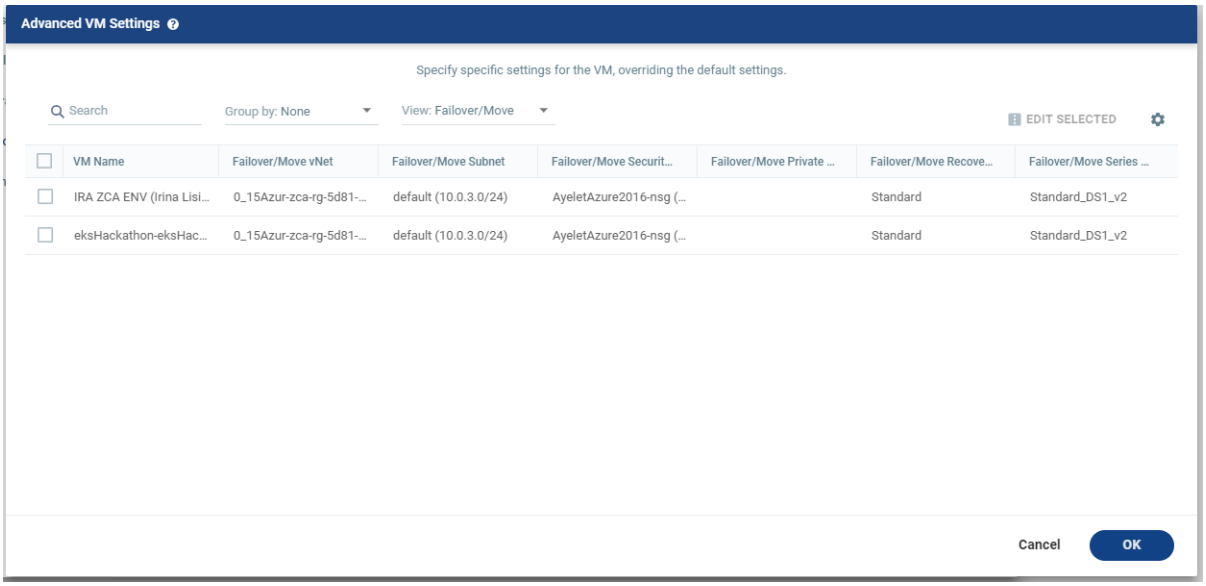

13. To edit information in one field, click the field and update the information.
14. To edit information for several virtual machines at the same time, select the virtual machines and click **EDIT SELECTED**.

    The **Edit VM Settings** window is displayed.

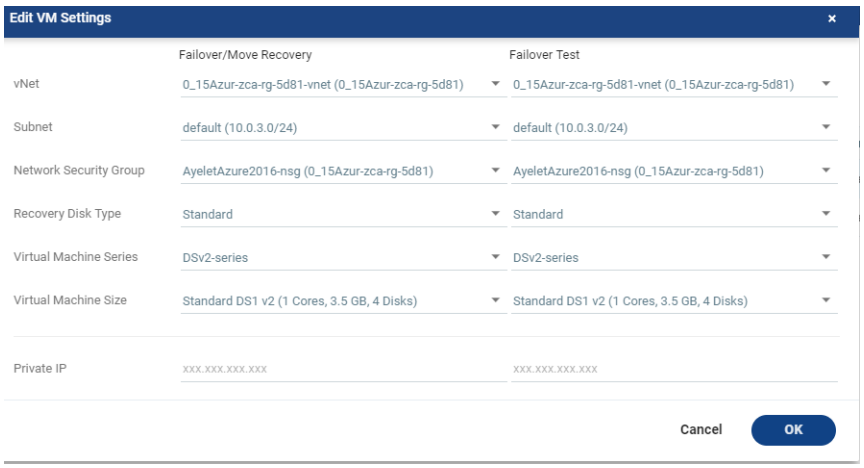

15. Update the values for **VNet**, **Subnet**, **Network Security Group**, **VM Series**, **Series Size**, and **Private IP** as necessary.

    Note: Only private IPs specified for Windows machines are assigned during a recovery operation. For Linux machines, the IP is assigned from the specified subnet range.

    Clearing the values in the **Private IP** field results in an IP being automatically assigned from the subnet range during a recovery operation.

    Refer to the Zerto Virtual Replication Interoperability Matrix for the list of operating systems for which Zerto supports re-IP.
16. Click **OK** twice to return to the main page of the **RECOVERY** step.
17. To run pre and post recovery scripts, enter the name of the script to run in the Command to run text box. You can then enter details about the script.
    - **Pre-recovery Script:** The information about a script that should run at the beginning of the recovery process.
    - **Post-recovery Script:** The information about a script that should run at the end of the recovery process.

    For both types of scripts, enter the following information:

TEXT BOX	DESCRIPTION
Command to run	The full path of the script. The script must be located on the same machine as the Zerto Virtual Manager for the recovery site.
Params	The parameters to pass to the script. Separate parameters with a space.
Timeout	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 the 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 time-out value is specified in **Site Settings > Performance and Throttling** tab.

NOTE:

Pre and post recovery scripts run in parallel. Therefore, ensure that the pre and post recovery scripts don't use common resources.

18. Click **NEXT**.

    The **SUMMARY** step is displayed. It shows a summary of the VPG configuration you defined in the previous steps.
19. Click **DONE**.

For details of what happens after saving the VPG, see "What Happens After the VPG is Defined", on page 35.

CHAPTER 6: MONITORING ZERTO VIRTUAL MANAGER

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 section:

- "The DASHBOARD Tab", below
- "Monitoring VPGs – The VPGs Tab", on page 105
- "Monitoring a Single VPG", on page 107
- "Monitoring Tasks", on page 111
- "Monitoring Protected Virtual Machines – The VMs Tab", on page 112

The following site monitoring option is described in this section:

- "Monitoring Peer Sites - The SITES Tab", on page 114

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.

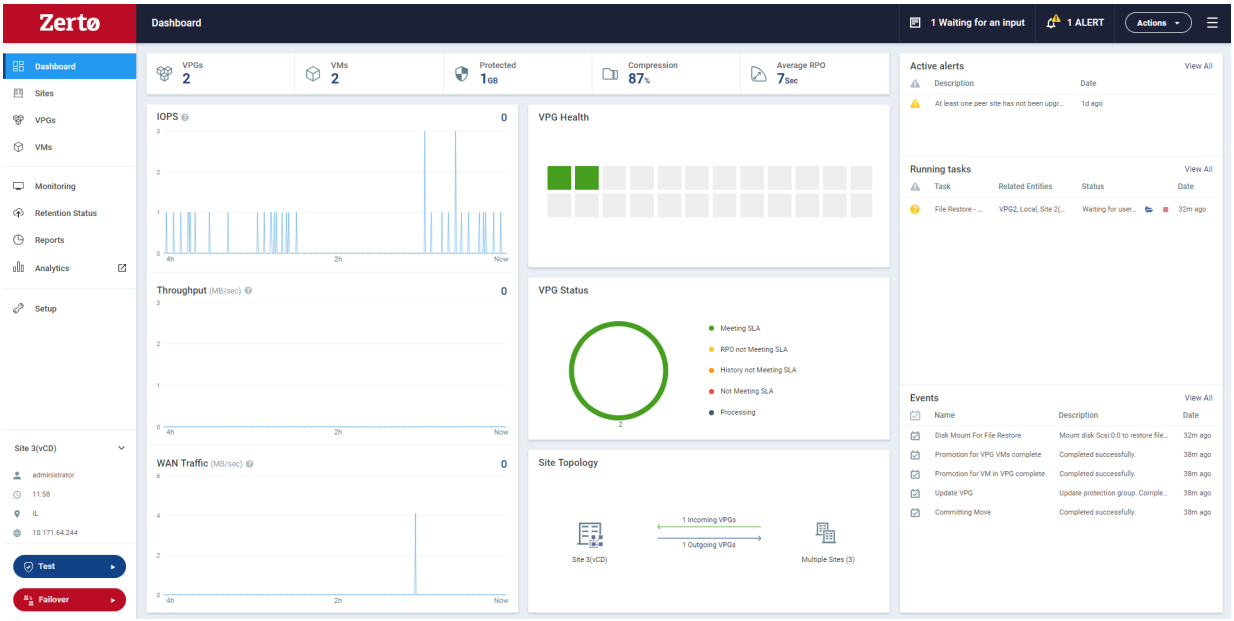

The following information is displayed:

VPG HEALTH

The VPGs being recovered 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 the site 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 (I/O per Second): The total amount of I/O write operations generated by the protected virtual machines on the viewed site. (The IOPS graph is displayed for VPGs replicating from on-premise sites).

Throughput (MB per second): The total amount of uncompressed data written by the virtual machines protected to the recovery sites.

During synchronization processes (such as bitmap sync, initial sync and delta sync) this value will also consist of the uncompressed data read from the protected disks.

WAN Traffic (MB per second): The total amount of compressed data transferred between the viewed site and all its recovery sites.

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.

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.

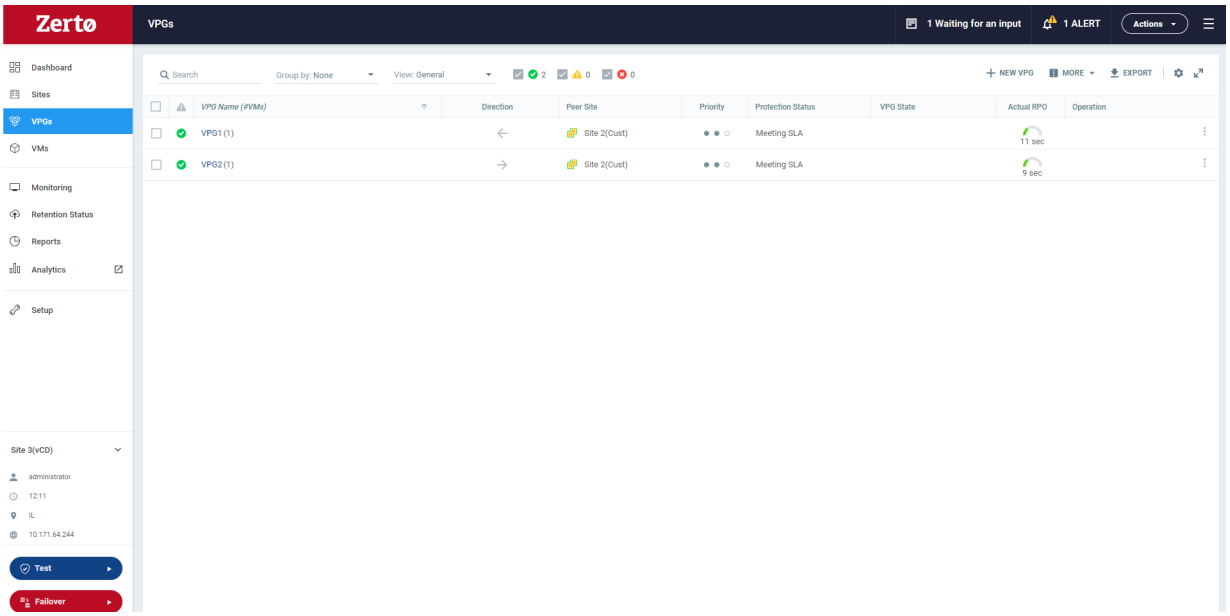

You can create a query using the view field to display VPG information in a list. You can filter the VPGs that will be displayed according to their status by checking the checkboxes alongside the VPG status icons (☑ 3 ☑ 1 ☑ 1). The query can be customized by adding and removing filters.

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.

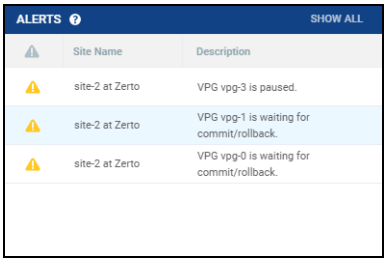

- **VPG Name (#VMs):** The name of the VPG. The name is a link: Click the VPG name to drill-down to more specific details about the VPG that are displayed in a dynamic tab. The number of VMs protected in the VPG is displayed in parentheses.
- **Direction:** The direction of the replication, from this 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 **IO** 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 - RETENTION STATUS

The following information is displayed in the RETENTION STATUS view:

- **Alert status indicator:** The color indicates the status of the VPG. Clicking 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.

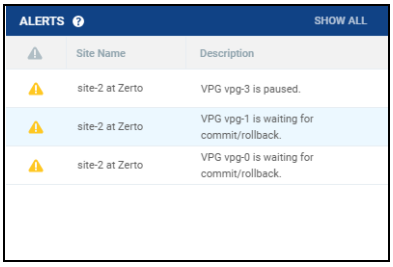

- **VPG Name (#VMs):** The name of the VPG. The name is a link: Click the VPG name to drill-down to more specific details about the VPG that are displayed in a dynamic tab.
- **Retention Policy:** Whether the VPG is protected against a disaster only with the ability to recover to a point in time up to 30 days before the disaster, or protection is extended to include retention sets of the virtual machines, going back for a maximum of one year.
- **Retention Policy Status:** The status of the retention set.
- **Repository Name:** The name of the repository where the jobs are stored.
- **Restore Point Range:** The restore points for the retention job out of the total retention jobs run for the VPG.
- **Retention Policy Scheduling:** The schedule for the retention process.

Additional Fields and Options

In the GENERAL, PERFORMANCE, and RETENTION 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.

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.

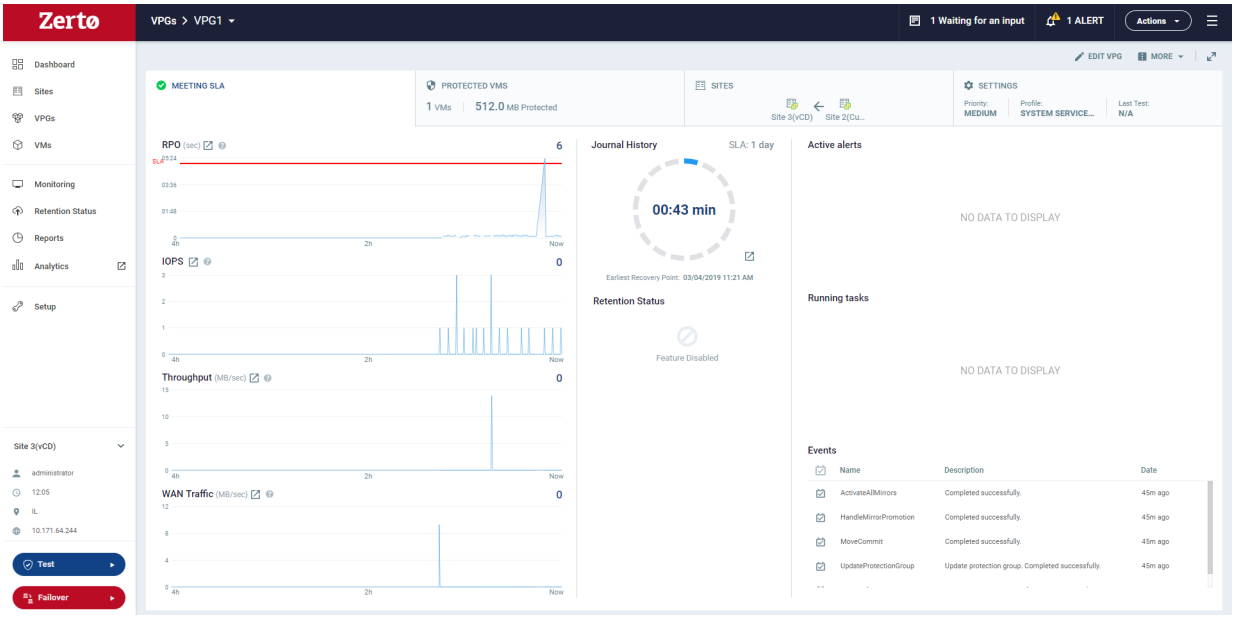

General Tab

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: The current Recovery Point Objective (RPO) of the Virtual Protection Group (VPG).

IOPS (I/O per second): The total amount of I/O write operations generated by the protected virtual machines comprising the Virtual Protection Group (VPG).

Throughput (MB per Second): The total amount of uncompressed data written by the protected virtual machines comprising the Virtual Protection Group (VPG).

During synchronization processes (such as bitmap sync, initial sync and delta sync) this value will also consist of the uncompressed data read from the protected disks..

WAN Traffic (MB per second): The total amount of compressed data transferred between the protected and recovery sites of the Virtual Protection Group (VPG).

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.

RETENTION STATUS

If Long Term Retention is enabled, the following 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 retention sets of the virtual machines, going back for a maximum of one year.
- **Retention Status:** The status of the retention set.
- **Repository:** The name of the repository where the retention sets are stored.
- **Restore Point Range:** The restore points for the retention sets out of the total retention sets run for the VPG.
- **Scheduling:** The retention schedule.

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

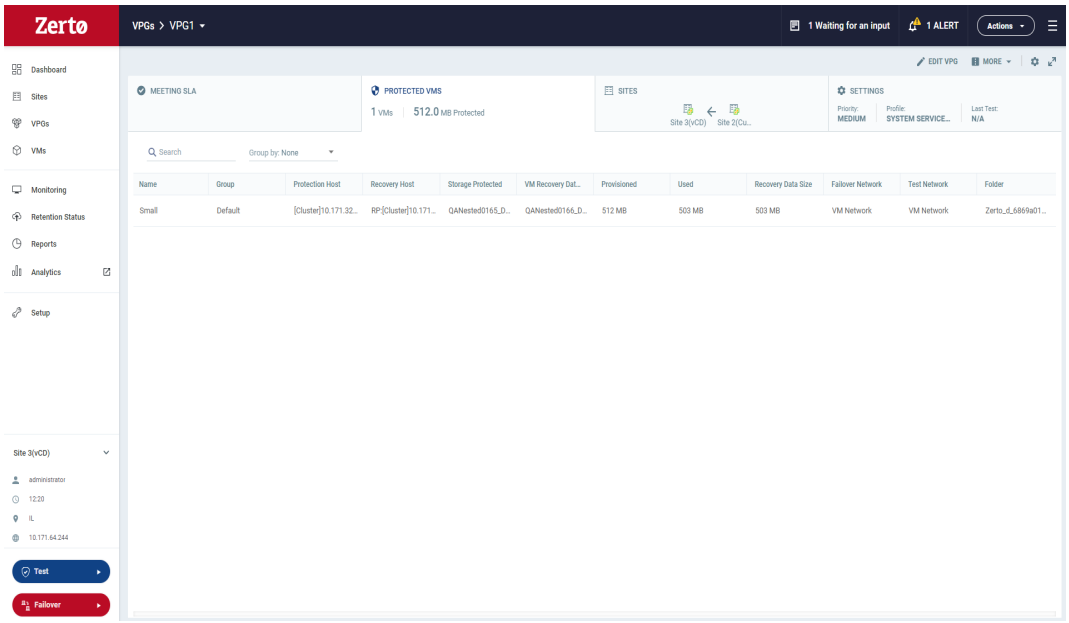

The PROTECTED VMs tab shows details about the protected virtual machines:

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

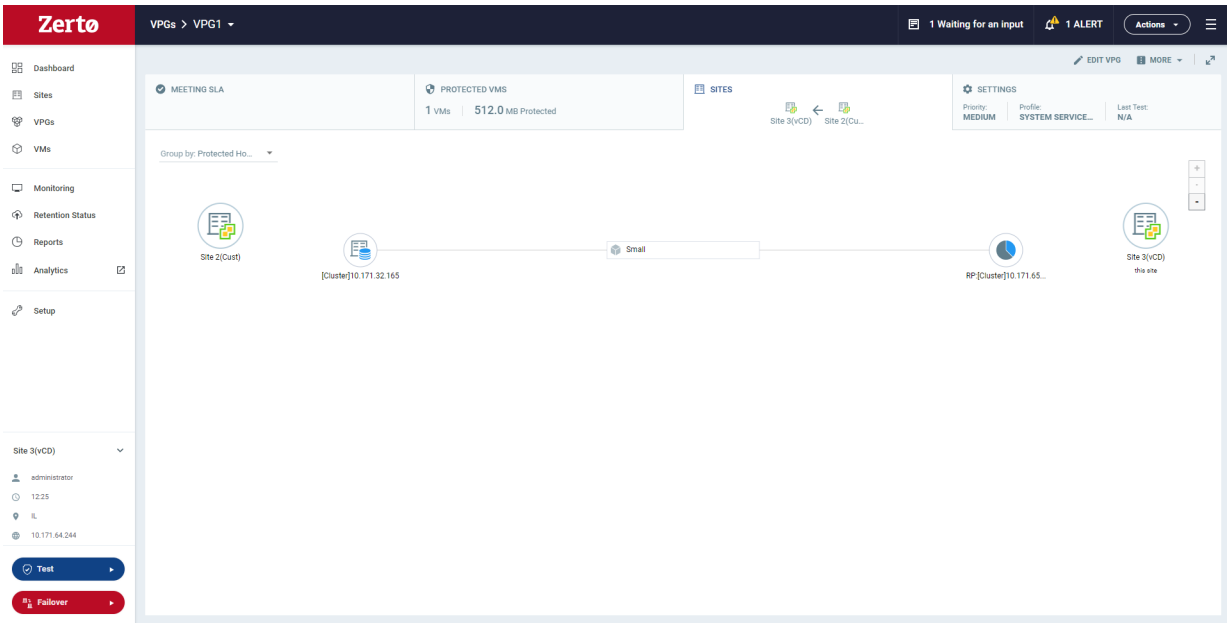

The SITES tab shows the topology of the VPG, including both the protected and recovery sites.

SETTINGS Tab

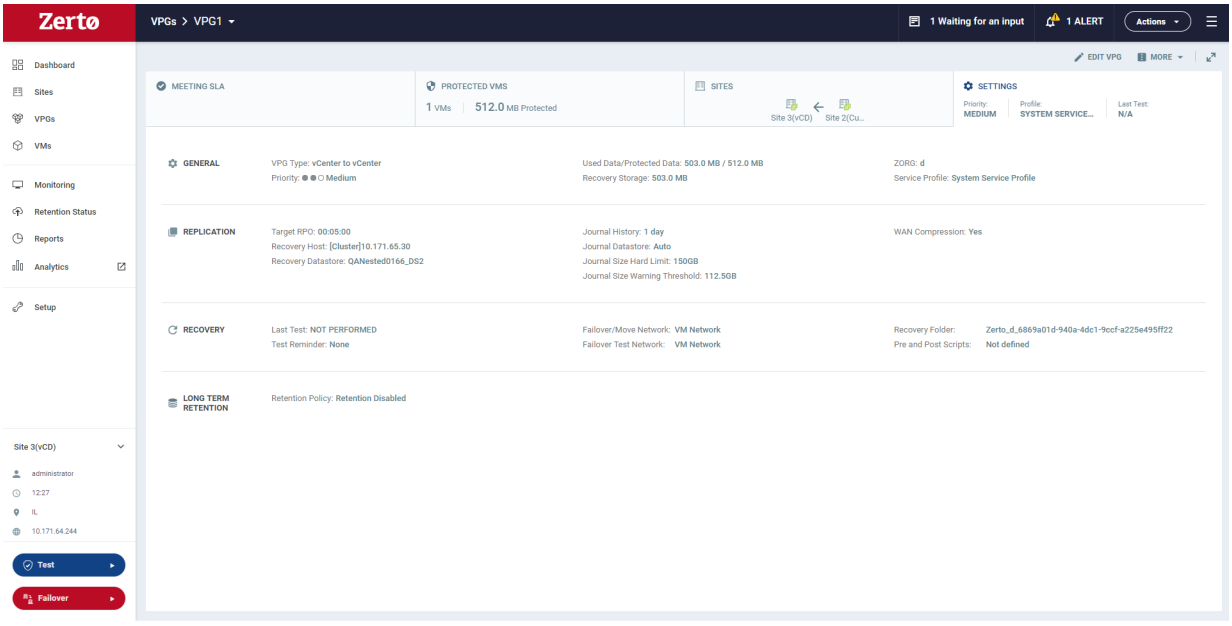

The SETTINGS tab shows details about the VPG settings, divided into general, replication, recovery, and retention categories.

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.

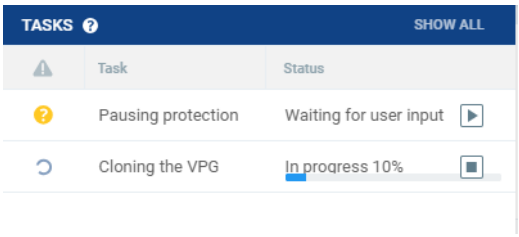

The following information is displayed for each task:

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.

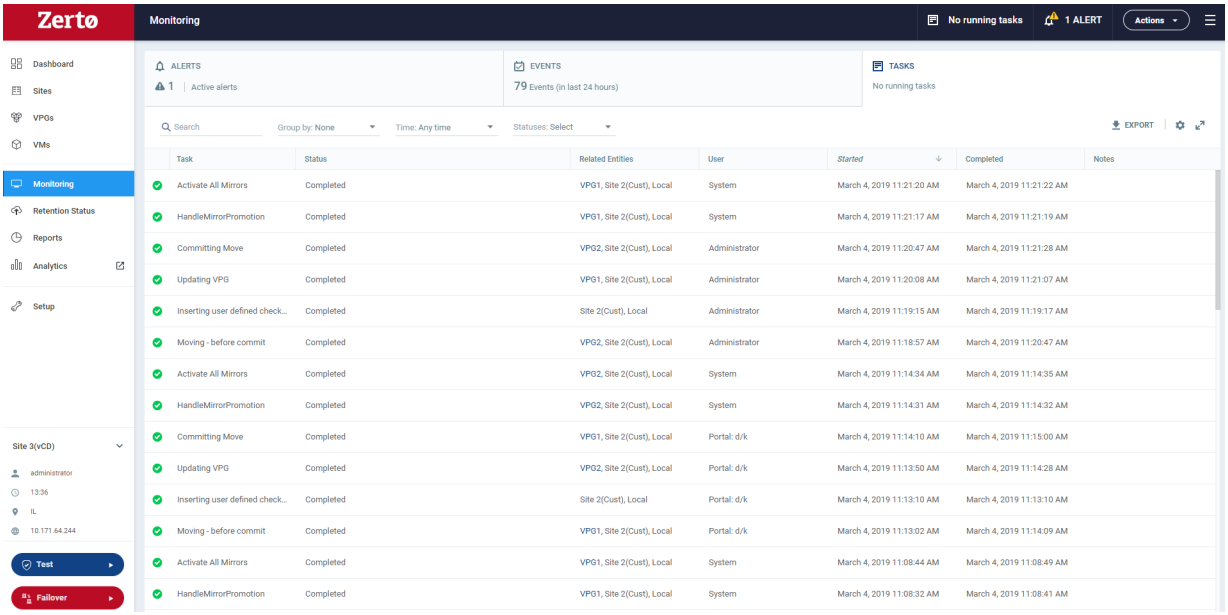

The following information is displayed for each task:

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.

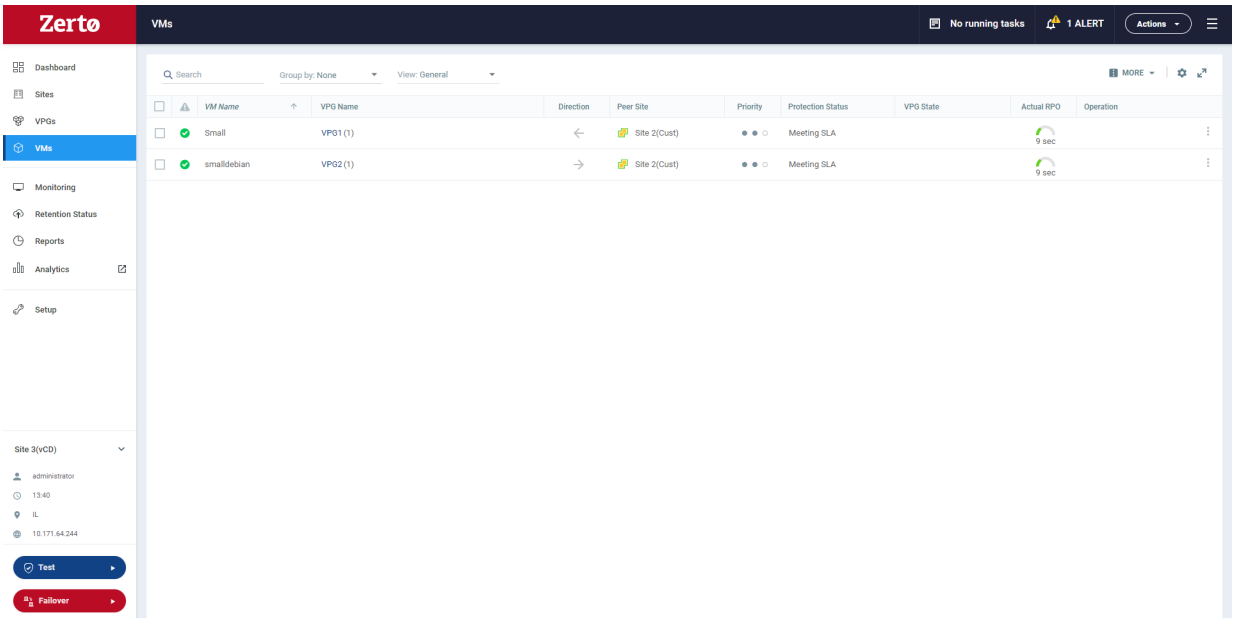

You can filter information in columns via the filter icon next to each column title. You can also sort the list by each column.

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.

RETENTION STATUS View

The following information is displayed in the RETENTION STATUS view:

- **Alert status indicator:** The color indicates the status of the VPG. Clicking 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.

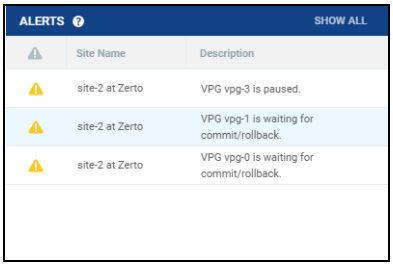

- **VPG Name (#VMs):** The name of the VPG. The name is a link: Click the VPG name to drill-down to more specific details about the VPG that are displayed in a dynamic tab.
- **Retention Policy:** Whether the VPG is protected against a disaster only with the ability to recover to a point in time up to 30 days before the disaster, or protection is extended to include retention sets of the virtual machines, going back for a maximum of one year.
- **Retention Policy Status:** The status of the retention set.
- **Repository Name:** The name of the repository where the jobs are stored.
- **Restore Point Range:** The restore points for the retention job out of the total retention jobs run for the VPG.
- **Retention Policy Scheduling:** The schedule for the retention process.

Additional Fields

In the GENERAL, PERFORMANCE, and RETENTION 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.

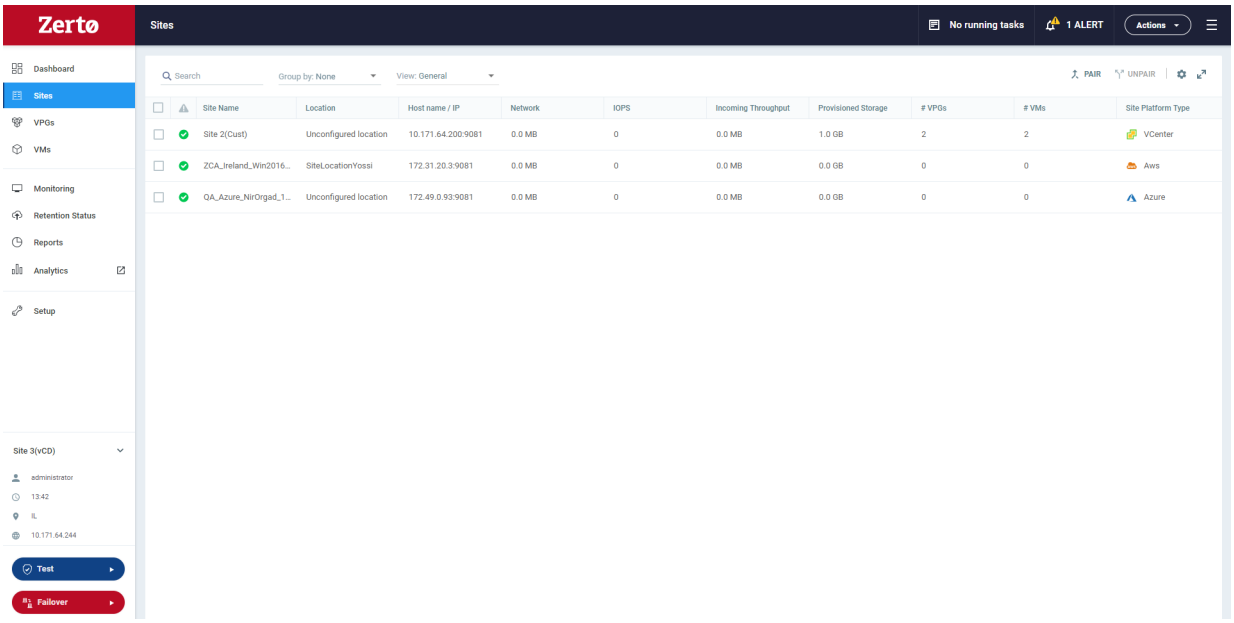

You can filter information in columns via the filter icon next to each column title. You can also sort the list by each column.

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 version installed at this site.

CHAPTER 7: 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 section:

- "Editing a VPG", below
- "Adding Virtual Machines to a VPG", on page 117
- "Removing Virtual Machines from a VPG", on page 118
- "Pausing and Resuming the Protection of a VPG", on page 118
- "Forcing the Synchronization of a VPG", on page 119
- "Deleting a VPG", on page 119
- "Ensuring Application Consistency – Checkpoints", on page 120
- "Running Scripts Before or After Recovering a VPG", on page 121
- "Exporting and Importing VPG Definitions", on page 124
- "VPG Statuses and Synchronization Triggers", on page 125

Monitoring VPGs and the VMs that are protected is described in "Monitoring Zerto Virtual Manager", on page 103.

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 retention process is running.

After modifying the VPG, the setting is updated.

While the VPG settings are 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 settings are 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**.

   Note: The Edit VPG wizard is displayed, enabling editing the VPG, including adding and removing virtual machines from the VPG. 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 117

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.

Adding Virtual Machines to a VPG

You can add virtual machines that are not already included in a VPG, to an existing VPG. A virtual machine can be protected in a maximum of three existing VPGs, provided that the VPGs are recovered to different sites.

Note: You cannot edit the VPG to add a virtual machine while a retention process is running.

To add a virtual machine to an existing VPG:

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

   The **Edit VPG** wizard is displayed, enabling you to edit the VPG, including adding and removing virtual machines from the VPG.
2. In the **VMs** step, select the virtual machines to add and click the arrow pointing right to include these machines in the VPG. A VPG can include virtual machines that are not yet protected and virtual machines that are already protected. You can view protected virtual machines by clicking **Select VMs** in the Advanced (Multi Target) area.

   Virtual machines protected in the maximum number of VPGs are not displayed in the Select VMs dialog.
3. If you want to define the boot order of the VPGs, click **DEFINE BOOT ORDER**.
4. Configure the settings for the new virtual machines in the same way that you configured the other virtual machines in the VPG, when you created the VPG.
5. Click **DONE**.

The virtual machines are added to the VPG. This process may take a few minutes. While the VPG definition is being updated, you cannot perform any operation on the VPG, such as adding a checkpoint, editing its properties, or recovering it.

After the VPG definition has been updated, the protected and recovery sites are then synchronized. During the synchronization period, the Protection Status displayed in the VPGs tab of the Zerto User Interface is: Meeting SLA n/m VMs where **n** is the number of virtual machines that were originally in the VPG, and **m** is the total number of virtual machines in the VPG, including the virtual machines that are currently being synced. While the virtual machines that were added are being synced, the VPG can be failed over but the failover only includes the original virtual machines in the VPG.

When the sync process for a virtual machine is complete, Zerto Virtual Manager tags the first checkpoint that includes a new virtual machine with: VM 'XXX' is fully synced where XXX is the name of the virtual machine that was synced.

When you perform a recovery operation using one of these checkpoints, or any later checkpoint, all the virtual machines that have completed syncing will be recovered.

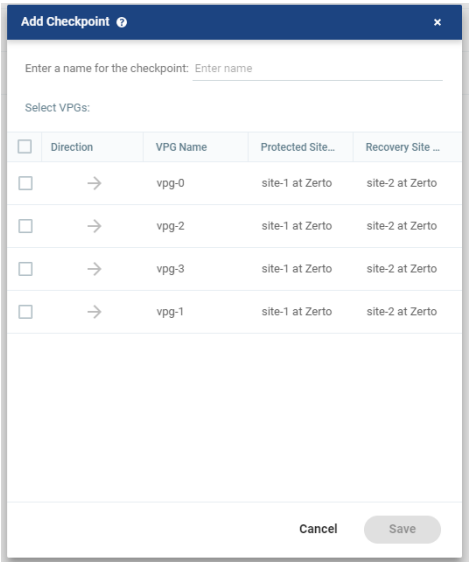

If the virtual machine is added to a VPG replicating to a resource pool in VMware vSphere environments, Zerto checks that the additional virtual machine doesn't exceed the resource pool capacity, such that the sum of the virtual machine reservation is less than or equal to the resource pool CPU and storage settings.

Removing Virtual Machines from a VPG

If a user removes a virtual machine from a VPG, the checkpoints of the VPG are retained.

Note: Once a virtual machine is removed, it is no longer possible to recover it.

Pausing and Resuming 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 retention process 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 retention process 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, and it's temp discs, start 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 amount of time which is 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. When using Long Term Retention, any retention sets stored for the VPG are not deleted and the virtual machines in this VPG can be restored.

Note: You cannot delete a VPG while a retention process 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 120.

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.

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 use a **script** to place the application in a quiesced mode, such as Oracle Hot Backup mode, and execute the Zerto **PowerShell** cmdlet **Set-Checkpoint**, then release the quiesced mode. For more information about Zerto PowerShell cmdlets, see *Zerto 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. 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 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 marked as part of a Failover test are not diluted..

See also Adding a Checkpoint to Identify a Key Point

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.

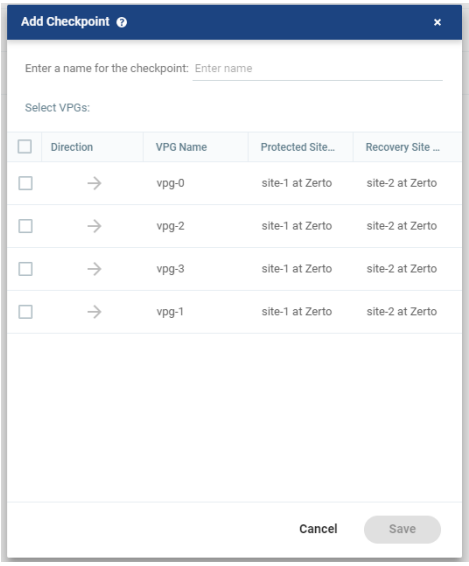

   A list of VPGs is displayed with the requested VPG selected. You can select more VPGs to add the same checkpoint to, for example, when something is happening at your site that affects multiple VPGs.

   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 149, or actually performing a failover, as described in "Managing Failover Live to AWS", on page 168, you can choose the checkpoint as the point to recover to.

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).
- **%ZertoVCenterIP%**:The IP address of the hypervisor manager, VMware vCenter Server or Microsoft SCVMM, where the VPG is recovered.
- **%ZertoVCenterPort%**: The port used by the Zerto Virtual Manager to communicate with the hypervisor manager, VMware vCenter Server or Microsoft SCVMM.
- **%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:

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 124.

Example Scripts

IMPORTANT: The scripts are provided by example only and are not supported under any Zerto support program or service.

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, for example, exporting the settings before uninstalling a version of Zerto and importing the settings after reinstalling Zerto.

Note: Zerto regularly exports settings to the **Zerto_Installation_Folder**`\ZVR\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 > Zerto Diagnostics**.
   The Zerto Diagnostics menu dialog is displayed.

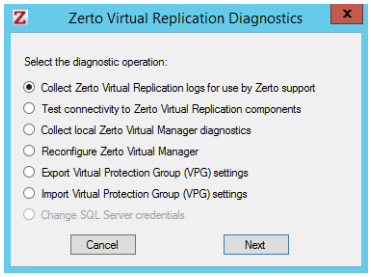

2. Select the Export Protection Group Settings option and click **Next**.

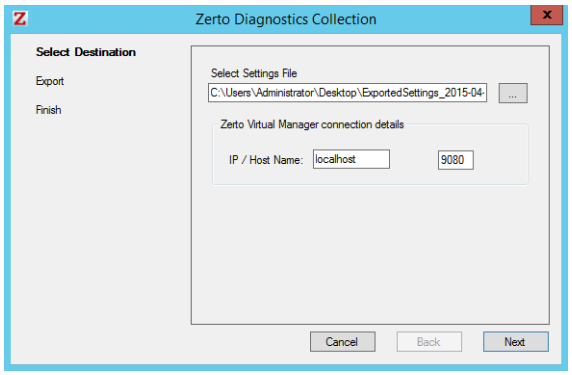

3. Select the destination for the file to contain exported settings and specify the Zerto Virtual Manager IP address and port where the VPGs are protecting virtual machines.
4. Click **Next**.
   The list of exported VPGs is displayed.

5. Click **Done**.

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

To import VPG settings:

1. Click **Start > Programs > Zerto > Zerto Diagnostics**.
   The Zerto Diagnostics menu dialog is displayed.
2. Select the Import Protection Group Settings option.
3. Click **Next**.

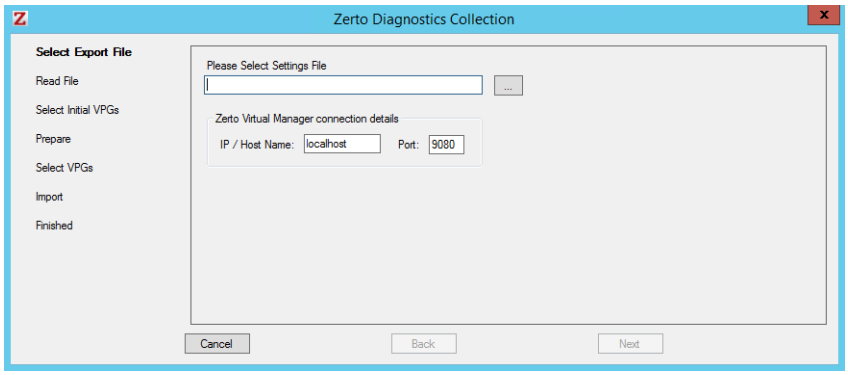

4. Select the file previously exported and enter the Zerto Virtual Manager IP address and port specified when exporting the VPGs.
5. Click **Next**.
   The list of exported VPGs is displayed.

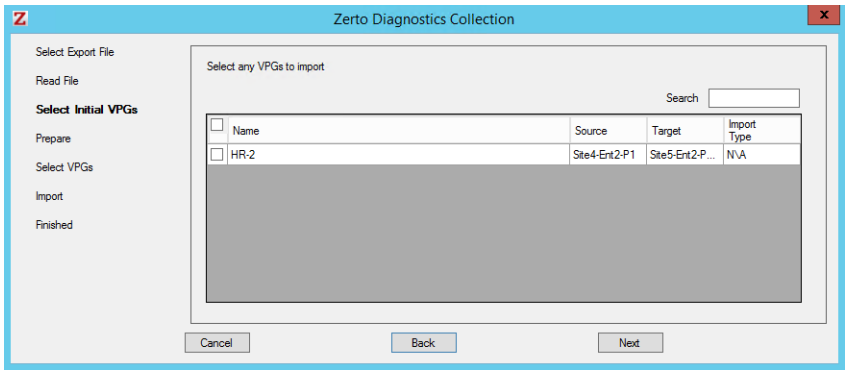

6. Select the VPGs to import. Only VPGs with names that are not already defined can be imported. VPGs in the import files with the same name as an existing VPG are disabled.
7. Click **Next**.
   The list of imported VPGs is displayed. If the VPG could not be imported, the reason for the failure is specified.
   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 126
- "VPG Synchronization Triggers", on page 130

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 or Based on Alerts	—	
	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	A retention process is running.
Bitmap Syncing	A change tracking mechanism of the protected machines during a disconnected state or when a VRA buffer is full. In these situations, Zerto 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 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 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 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 143. **Note:** 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.
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	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. **Note:** 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.
Delta syncing a volume	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. **Note:** 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.
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	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. **Note:** 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.

SUBSTATUS	DESCRIPTION
Initial Sync	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. **Note:** 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.
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.

SUBSTATUS	DESCRIPTION
Volume Full Syncing	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. **Note:** 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.
Volume Initial Sync	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`. **Note:** 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 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.
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 120.
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.

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 119.
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.

TRIGGER	DESCRIPTION
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 8: MANAGING ZERTO VIRTUAL MANAGER

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

The following topics are described in this section:

- "Check Connectivity Between Zerto Components", below
- "Reconfiguring the Zerto Virtual Manager Setup", on page 133
- "Reconfiguring the Microsoft SQL Server Database Used by the Zerto Virtual Manager", on page 134
- "Replacing the SSL Certificate", on page 135
- "Pair to Another Site and Unpairing Sites", on page 135

Check Connectivity Between Zerto 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 > Zerto Diagnostics.

   The Zerto Diagnostics menu dialog is displayed.

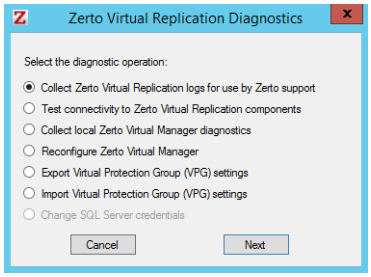

2. Select the Test Connectivity to Zerto components option and click **Next**.

   The IP Connectivity dialog is displayed.

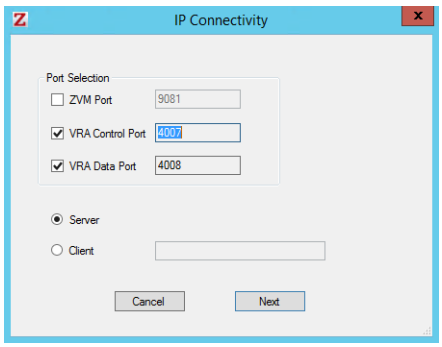

   You can use this dialog to check the following:
   - TCP communication between the Zerto Virtual Managers (ZVMs) on the protected and recovery sites. The default port, specified during installation, is 9081.
   - Communication between VRAs on the protected and recovery sites, via the control port and 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**.

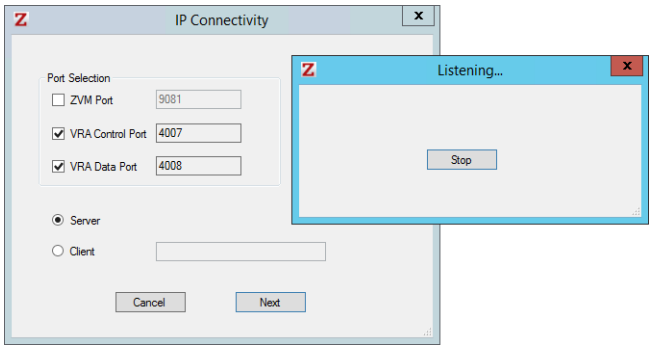

   The Client options tests the client; on completion a result dialog is displayed.

6. Click **Stop** (server test) or **OK** (client test) to return to the Zerto 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 Diagnostics utility.

To reconfigure the Zerto Virtual Manager:

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

   The Zerto Diagnostics menu dialog is displayed.

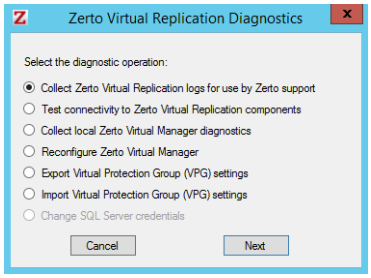

2. Select the Reconfigure Zerto Virtual Manager option and click **Next**.

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

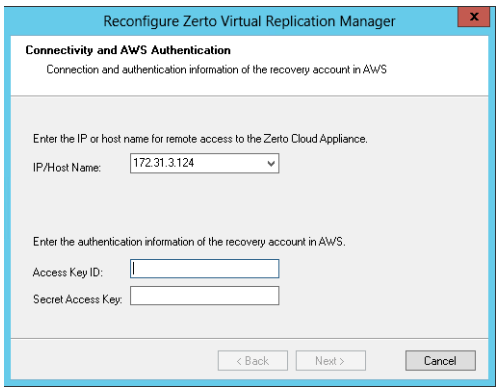

3. Change the IP/host name, and access key ID and secret access key, if necessary.

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, 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 > Zerto Diagnostics**.

   The Zerto Diagnostics menu dialog is displayed.

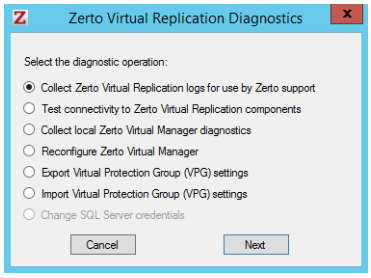

2. Select the Change SQL Server Credentials option and click **Next**.

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

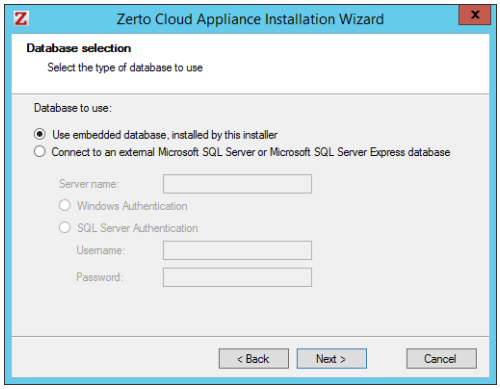

   Server Name: The domain name and server instance to connect to, with the format <server_name>\<instance_name> or <Server_IP>\<instance_name>.

   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 133 and select a new SSL certificate when the dialog for Zerto Virtual Manager setup is displayed:

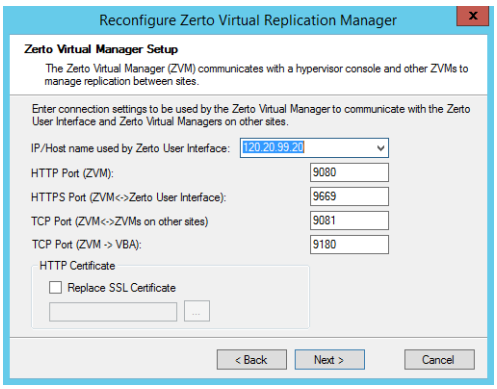

HTTP Certificate: Check Replace SSL Certificate and browse for a replacement certificate.

Pair to Another Site and Unpairing Sites

See the following sections:

- "Pair to Another Site", below
- "Unpairing Sites", on page 136

Pair to Another Site

You can pair to any site where Zerto is installed.

Zerto can be installed at multiple sites and each of these sites can be paired to any other site on which Zerto has been installed. Virtual machines that are protected on one site can be recovered to any paired site.

To pair to a site:

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

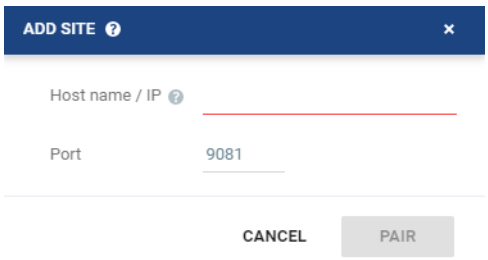

2. Specify the following:
   Host name/IP: IP address or fully qualified DNS host name of the **remote** site Zerto Virtual Manager to pair to.

Port: The TCP port communication between the sites. Enter the port that was specified during 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 site is connected to the Zerto Virtual Manager at the remote site.

   After the pairing completes the content of the SITES tab changes to include summary information about the paired 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. For vSphere, Hyper-V and Azure platforms, you can select to keep disks to use for preseeding if the VMs are re-protected. If you select this option, the disks are not removed from the recovery site.
4. To unpair, click **CONTINUE**.

   The sites are no longer paired. If there are VPGs on either site, they are 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, depending if you selected to keep disks to use for preseeding.

   The locations of the saved target disks are specified in the **Events** tab in the ZVM application on the **Recovery** site.

CHAPTER 9: OVERVIEW OF DISASTER RECOVERY OPERATIONS

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

- "The Failover Test Operation", below
- "The Move Operation", on page 138
- "The Failover Operation", on page 138
- "The Restore File Operation", on page 139

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 149.

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.

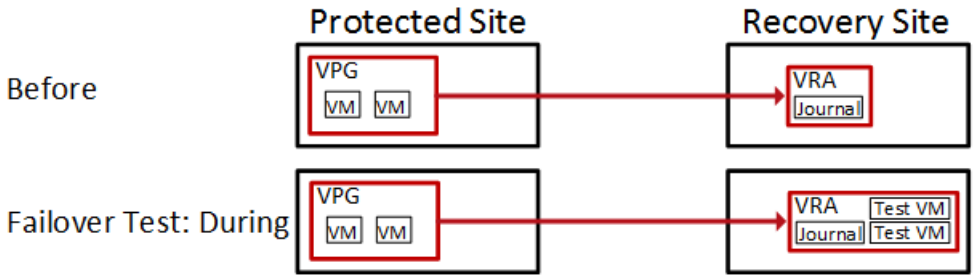

The Move Operation

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 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 157.

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

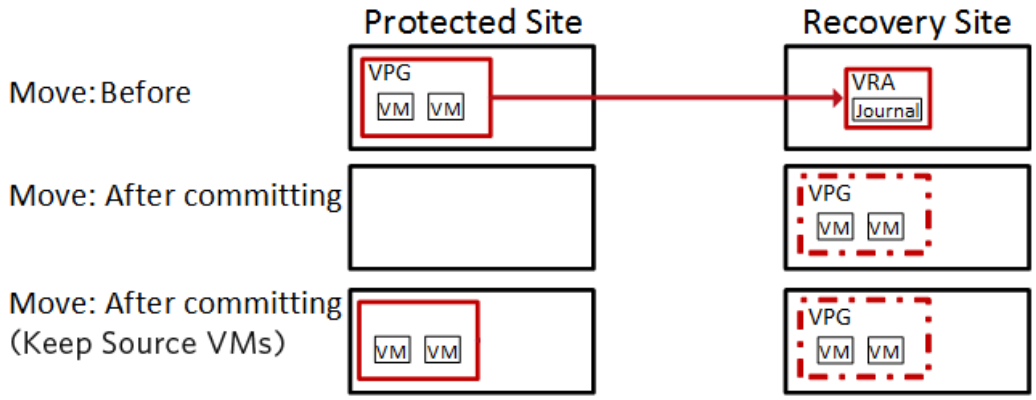

Note: The Move operation leaves the VPG in a Recovered state.

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 makes sure that virtual machines in AWS are recovered to this specified point-in-time. For details, see "Managing Failover Live to AWS", on page 168.

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.

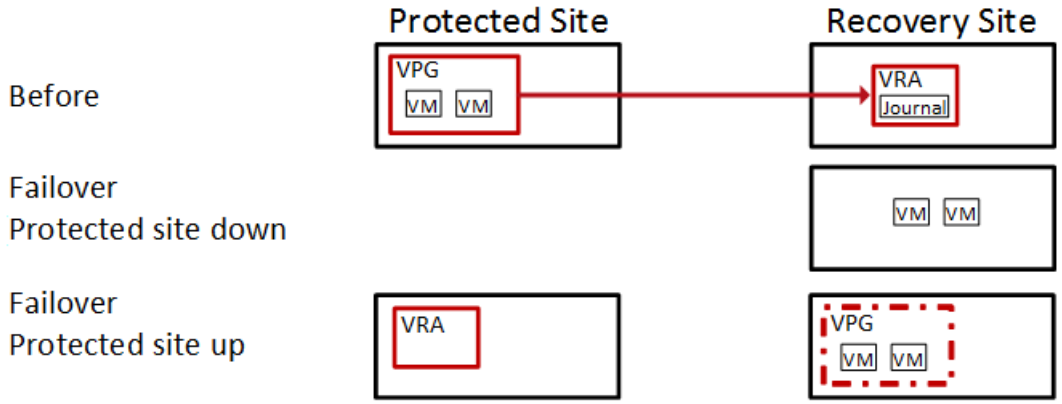

Note: The Failover operation leaves the VPG in a Recovered state.

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 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 180.

CHAPTER 10: 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 section:

- "Site Settings", below
- "Seeing What is Licensed", on page 145
- "Submitting a Support Case", on page 146
- "Submitting a Feature Request", on page 147
- "About the Zerto Version", on page 148

Site Settings

The Site Settings dialog enables configuring various site settings. These include the 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 ☰ and select Site Settings.
   The Site Setting window is displayed.

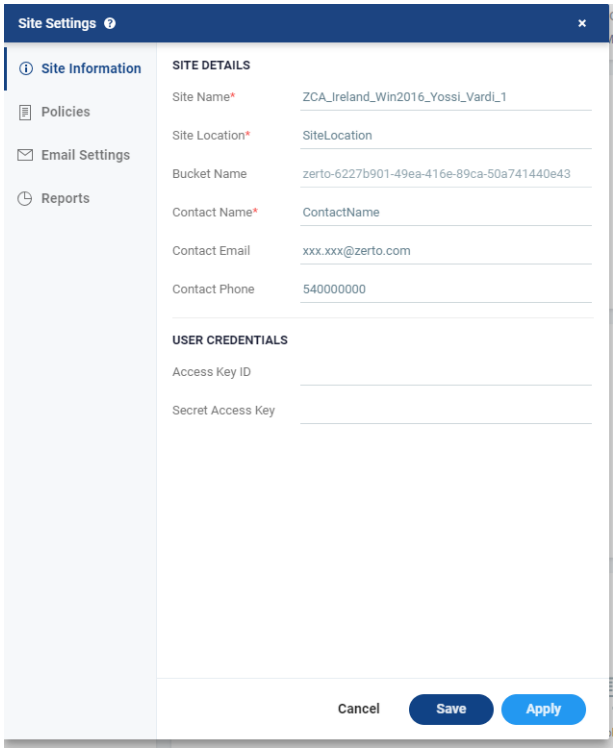

2. Make any required changes to the settings, click **SAVE** and then **APPLY**. The following setting can be defined:
   - "Editing Information About a Site", below
   - "Defining Site Policies", on page 143
   - "Configuring Email Settings", on page 144
   - "Defining the Resource Report Sampling Period", on page 145

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.

To update information about the local site:

1. In the Zerto User Interface (top right), click **SETTING** ( ) and select **Site Settings**.
   The Site Settings dialog is displayed.
   - In **vSphere** and **Hyper-V** sites, the following fields are displayed:

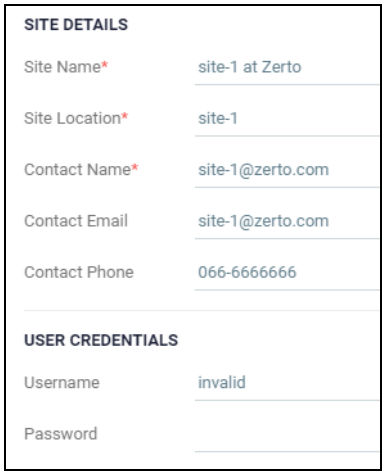

   - In **AWS** sites, the following fields are displayed:

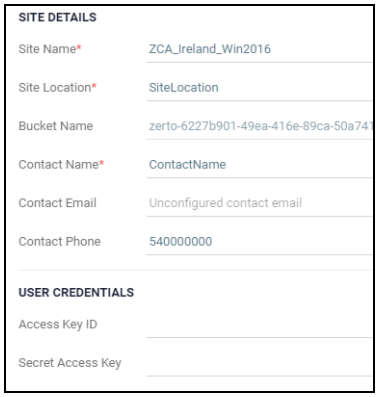

   - In **Microsoft Azure** sites, the following fields are displayed:

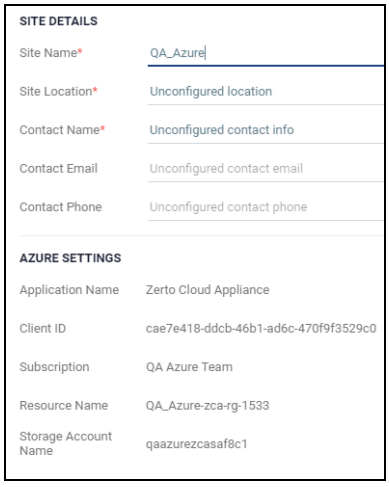

2. Define general information about the site.
   - (Mandatory) **Site Name:** The name used to identify the site.
   - (Mandatory) **Site Location:** Information such as the address of the site or a significant name with which to identify the site location.
   - (AWS environments only) **Bucket Name:** The name of the bucket that was created when Zerto was installed. This cannot be changed.
   - (Mandatory) **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. (On premise environments) To change the **User Credentials** to access the vCenter or Hyper-V SCVMM server from the Zerto Virtual Manager:
   - **User Name**: The administrator name used to access the vCenter or SCVMM server. The name can be entered using either of the following formats:
     - **username**
     - **domain\username**
   - **Password:** The password used to access the vCenter or SCVMM server for the given user name. To ensure security, after saving the settings, the password field is cleared.
4. (Azure environments only) **Azure Settings**:
   - **Application Name:** The name used to access Azure.
   - **Client ID**: A unique identifier that is associated with the access name.
   - **Subscription:** The subscription associated with the user.
   - **Resource Name:** The name of the resource the user created.
   - **Storage Account Name:** The name of the storage account created or selected for this site during installation.
5. Click **SAVE**.

Defining Site Policies

You can set default recovery and replication policies.

Configuring Disaster Recovery Policies

To configure disaster recovery policies:

1. Click **Policies**.

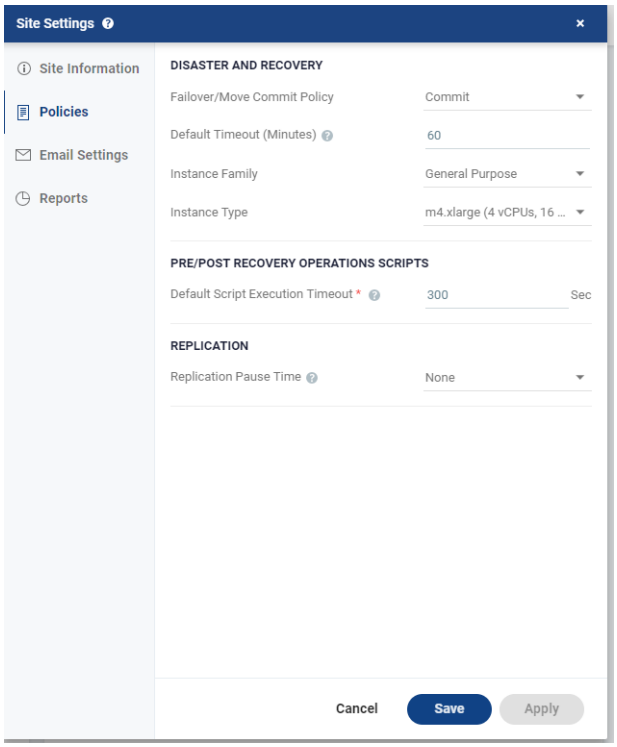

2. Choose the Failover/Move Commit Policy to use during a failover or move operation, described in "Initiating Failover Live", on page 169 and "Moving Protected Virtual Machines to a Remote Site", on page 159 respectively. The following options are available:
   - **None:** The failover or move operation 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. Specify the **Default Script Execution Timeout** or a script running before or after the failover, move, or test failover. A value of zero indicates that there is no time out. Values can be between 300 to 6000 seconds.

7. 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 176. The value set here is applied to existing and new VPGs.
8. Click **APPLY** or **SAVE**.

Configuring Email Settings

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

Email Settings

To configure email settings:

1. Click **Email Settings**.

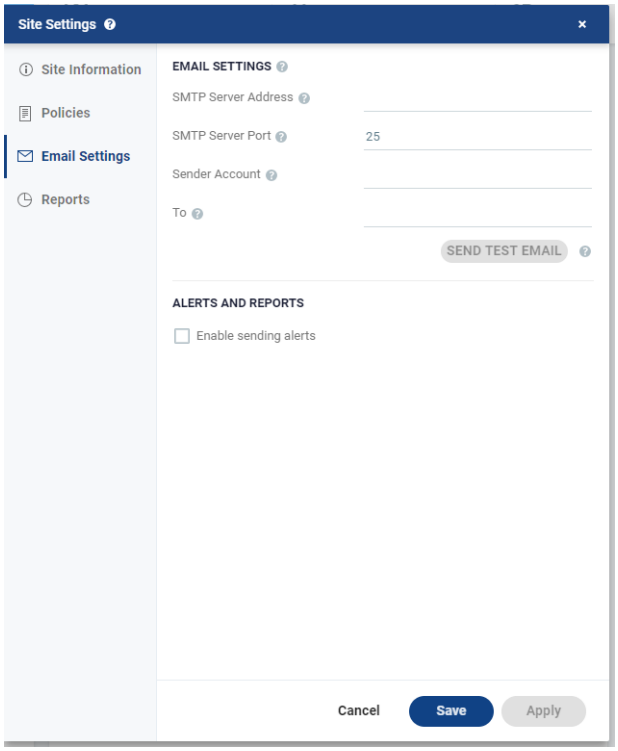

2. Specify the **SMTP server Address.** The Zerto Virtual Manager must be able to reach this address.
3. If the **SMTP Server Port** was changed from the default, 25, specify the port number.
4. Specify a valid email address for the email sender name in the **Sender Account** field.
5. Specify a valid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

Enabling Sending Alerts

You can enable sending alerts to an email address.

Select **Enable sending alerts** to enable this option.

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**.

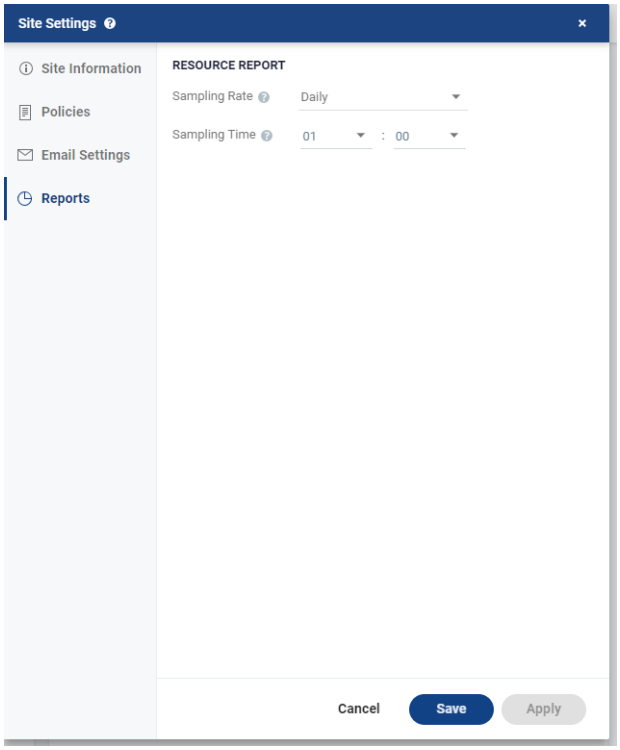

2. Choose the Sampling Rate.
3. Choose the Sampling Time.

   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 Reports", on page 192.

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 by clicking ☰ in the top right of the header and selecting **License**.

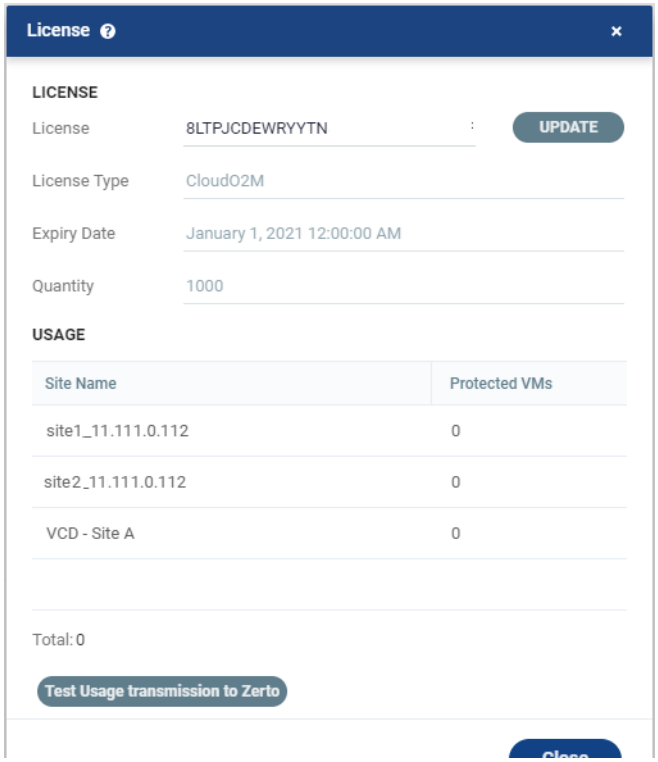

The Zerto license includes the following details:

- **License:** The license key itself.
- **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.
- **Usage:** The sites using the license and the 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 with this key.

To update a license key:

1. In the Zerto User Interface, in the top right of the header click and select **License**.
2. Enter a valid license key and click **APPLY** or **SAVE**.
   The license is updated on the local site and the paired remote sites.

Submitting a Support Case

Support cases can be opened directly in the Zerto User Interface.

Creating a support case in the Zerto User Interface simplifies the submission process since much of the information that is required when entering a case using the Zerto Support Portal, such as the version and build numbers, is automatically added to the case when it is submitted via the Zerto User Interface.

In addition, when the case is submitted, a snapshot of the current environment is also attached to the case. 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 case quickly and, whenever possible, without the need to request more information from you.

Note: The clocks on the machines where Zerto 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 case can fail.

To open a support case:

1. In the Zerto User Interface, click ☰ in the top right of the header and select Open a Case.
   The Open Support Case window for the site opens.

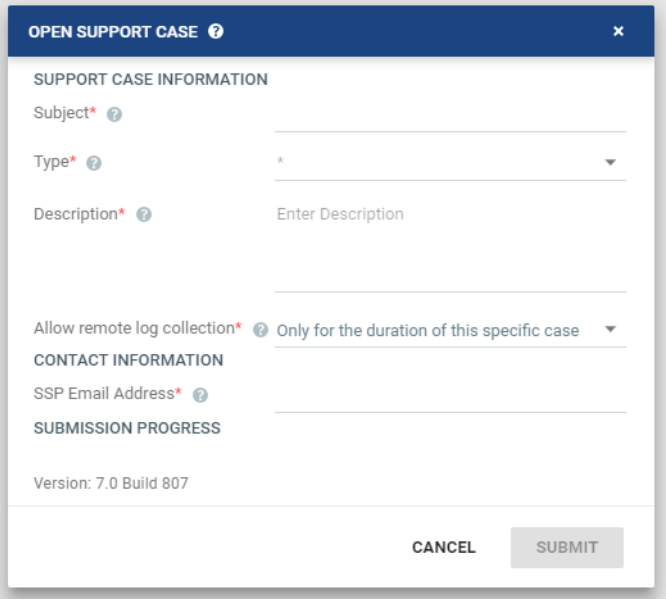

2. Specify the case details:
   - **Subject:** The subject of the support case.
   - **Type:** The type of case being opened. Available options are:
     - Problem
     - Question
   - **Description:** A description of the problem or question in addition to the information supplied in the subject.
   - **Allow remote log collection:** How many logs is Zerto allowed to collect. Available options are:
     - Only for this case
     - For the next 30 days
     - Never
   - **SSP Email Address:** A valid email address registered with Zerto, with permission to open cases.
3. Click **SUBMIT**.

The case is processed and its progress is displayed. If the email address is not valid, the case is rejected. Once the case submission starts, it cannot be canceled.

Submitting a Feature Request

From the Zerto User Interface, you can access the Feature Requests page in myZerto in order to submit a feature request.

To submit a feature request:

1. In the Zerto User Interface, on the top right of the header, click ☰.
2. Click **Submit Feature Request**; the Feature Requests page in myZerto opens.

About the Zerto Version

In the About window, you can do the following:

- View the version of Zerto being run.
- Enable or disable the **Zerto CALLHOME feature**. The Zerto CALLHOME feature enables support notification and analytics for the following purposes:
  - To improve Zerto.
  - To send notifications to the user when a new Zerto version is available, or when new hypervisor versions are supported by Zerto.
- Enable or disable Zerto Virtual Manager to send data to the SaaS platform for monitoring purposes, using the Zerto Mobile App. This action is done by licensed Zerto Virtual Manager users.

When clicking About, the following options appear:

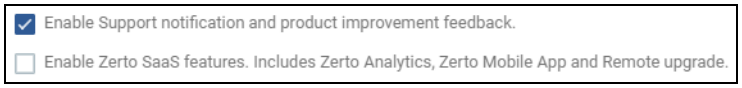

To perform these actions, do the following:

1. In the Zerto User Interface, in the top right of the header, click ☰, and then click **About**.

   The version and build of Zerto installed in the site are displayed.
2. The Zerto CALLHOME feature is selected by default; **Enable Support notification and product improvement feedback**.

   Note: This option is grayed out for Microsoft Azure and AWS.

   If the user deselects Enable Support notification and analytics, a warning appears notifying the user that deselecting this option will stop Zerto from sending notifications when new Zerto updates are available, or when new hosts are supported.
3. If you want Zerto to send information to our Online Services and Zerto Mobile App, and enable remote upgrade, select **Enable Zerto SaaS features. Include Zerto Analytics, Zerto Mobile App and Remote upgrade.** This is selected by default.

   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.
   - If the user deselects Enable Online Services and Zerto Mobile, a warning appears notifying the user that deselecting this option will stop Zerto from sending information to Online Services and to the Zerto Mobile Application, rendering these services inoperable for the entire installation.

CHAPTER 11: 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 section describes how to test VPG recovery.

The following topics are described in this section:

- "The Test Failover Process", below
- "Starting and Stopping Failover Tests" on page 150
- "Viewing Test Results", on page 154
- "Live Disaster Recovery Testing", on page 154

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 retention process 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 Failover Live During a Test", on page 175.

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 specific VMs in a VPG, 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 click

.
The Failover Test wizard is displayed.

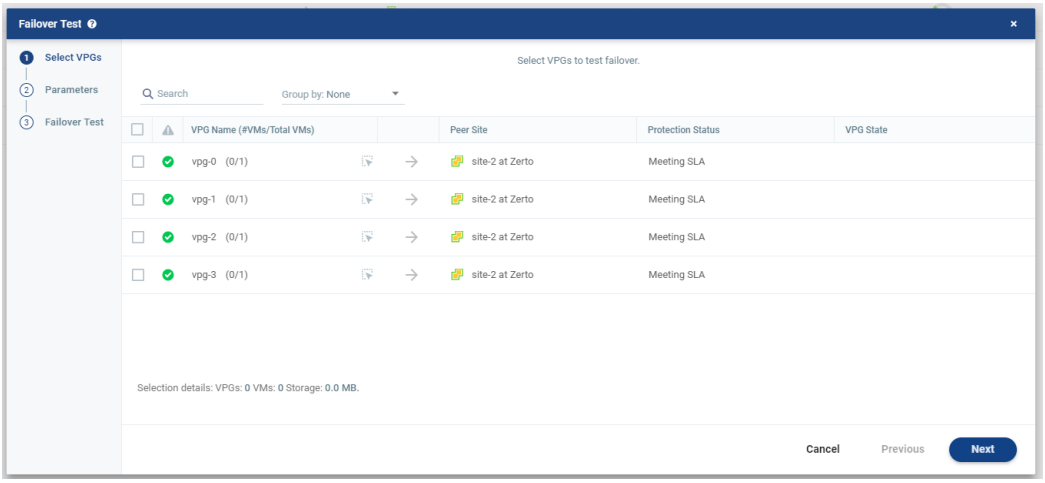

2. Select the VPGs to test. By default, all VPGs are listed.
   a) To select specific VMs in a VPG, click the icon next to each VPG to get a list of VMs. The Select VMs to Failover dialog is displayed. By default, all VMs are selected.

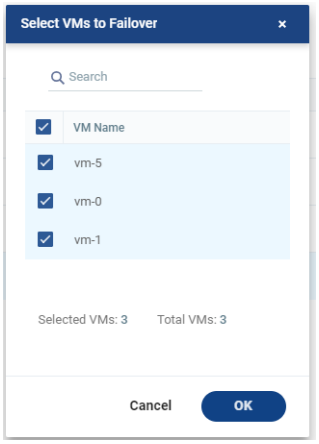

   b) Select the VMs to test.

   Note: Selecting specific VMs in a VPG to failover is not supported when replicating from a **vCD site**.

   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 **PARAMETERS** step is displayed.

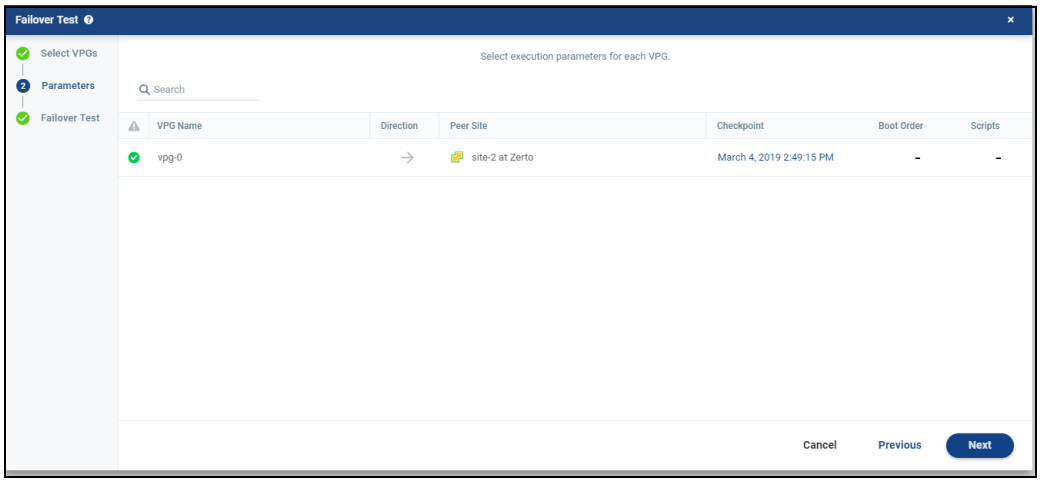

You can select the checkpoint to use for the recovery and see if a boot order and scripts are defined for the VPG.

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.

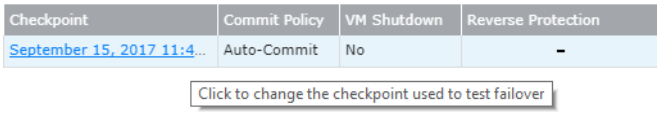

   A window appears, displaying a list of the VPGs' checkpoints.

   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. 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.
6. Click **OK**. A warning appears informing the user if the selected VMs were not protected with the selected checkpoint.
7. 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.

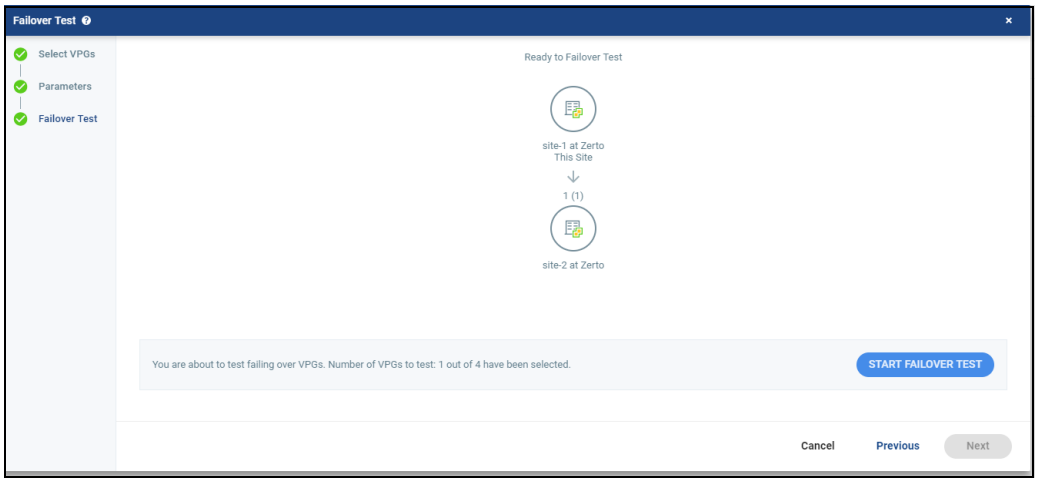

8. To start the test, click **START FAILOVER TEST**.

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 is **m4.xlarge** . 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 143. 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 Failover Live During a Test", on page 175.
- 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.

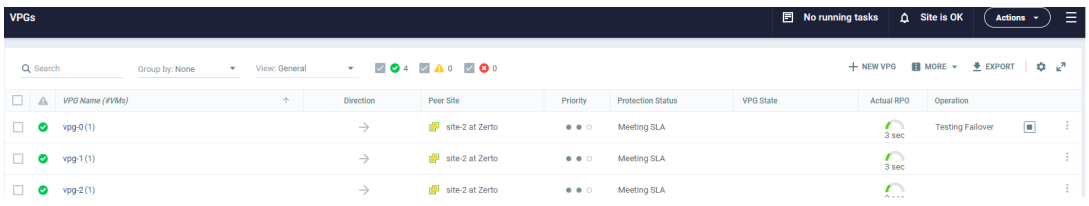

- In the Zerto User Interface, click the **VPGs** tab, and then click the name of a VPG you are testing. A dynamic tab is created displaying the specific VPG details including the status of the failover test.

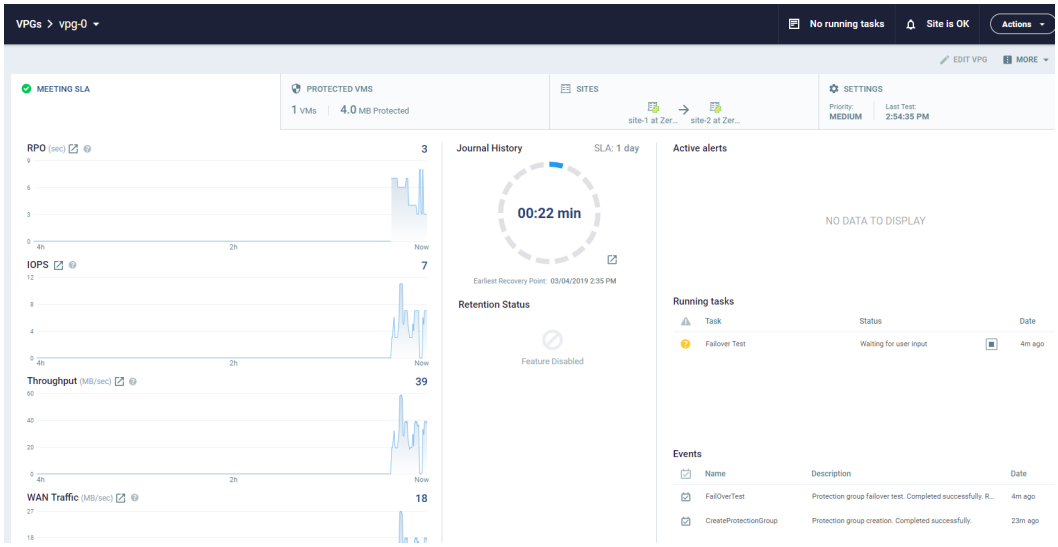

To stop a failover test:

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

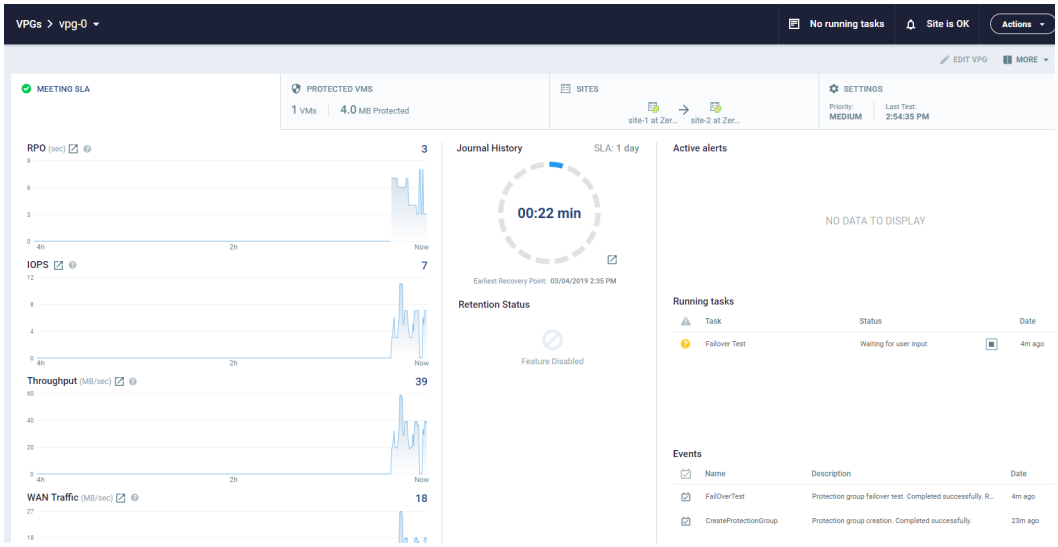

You can also stop the test via the TASKS popup dialog in the status bar, or by selecting **MONITORING > TASKS**.

The Stop Test dialog is displayed.

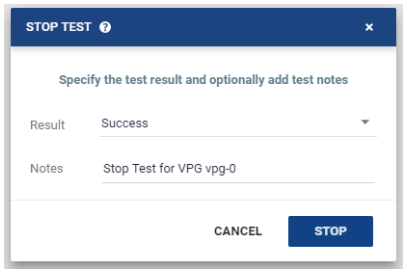

2. In the Result field specify whether the test succeeded or failed.
3. Optionally, in the Notes field, add a description of the test. For example, specify where external files that describe the tests performed are saved. Notes are limited to 255 characters.

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.

5.

Viewing Test Results

After stopping a test, you can see the test results in Zerto reports. For more information, see "Zerto Reports", on page 192.

Live Disaster Recovery Testing

This section describes how to use the basic Zerto 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:

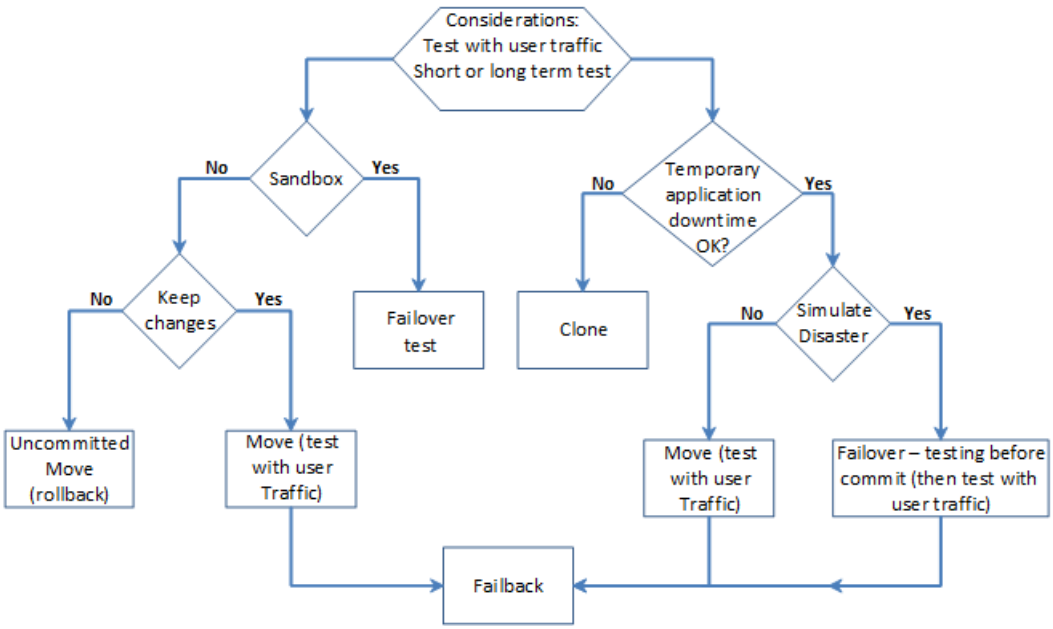

During any live test, Zerto recommends that you only maintain one working version of the same virtual machine. Thus, the first step in any test, except for a Failover Test, is to make sure that the protected virtual machines are shut down before starting to

test recovered machines. During a Zerto Move operation, the first step Zerto performs, is to shut down the protected machines, to ensure data integrity. However, a Zerto 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 150. 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 159. 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 12: MIGRATING A VPG TO AWS

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

- "The Move Process", below
- "Moving Protected Virtual Machines to a Remote Site", on page 159

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 retention process 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 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 Live to AWS", on page 168.

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 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 is m5.xlarge. 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 143. 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.

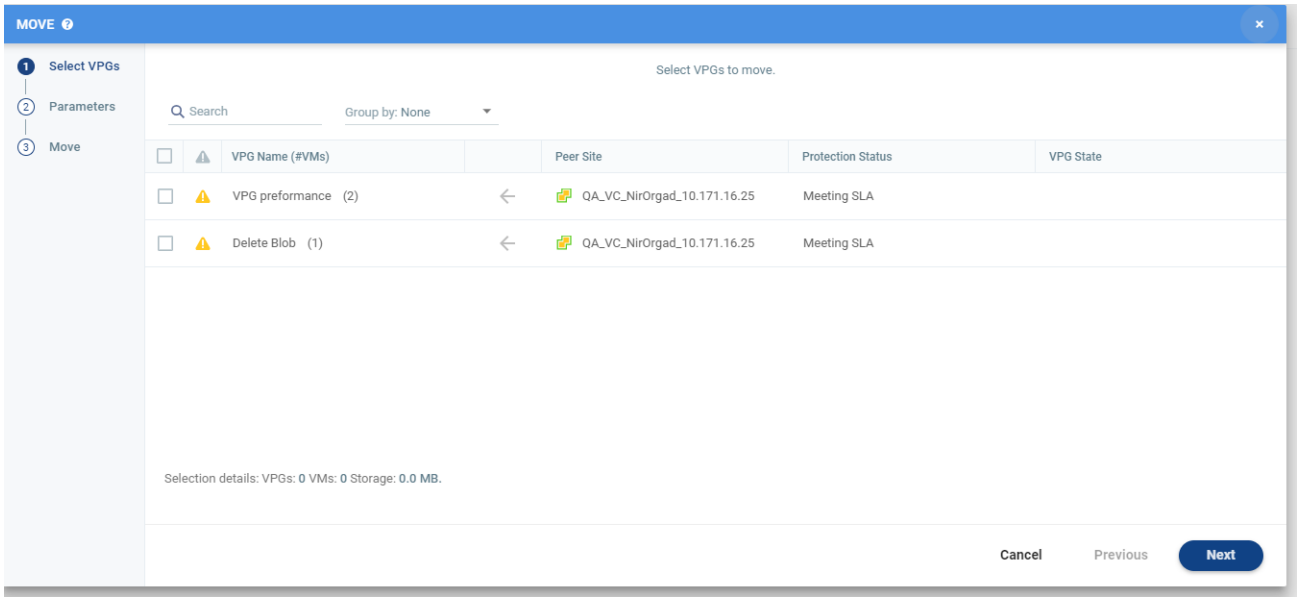

2. Select the VPGs to move.
   At the bottom, the selection details show the amount of data and the total number of virtual machines selected.
   The Direction arrow shows the direction of the process: from the protected site to the peer, recovery, site.
3. Click **NEXT**. The PARAMETERS step is displayed.

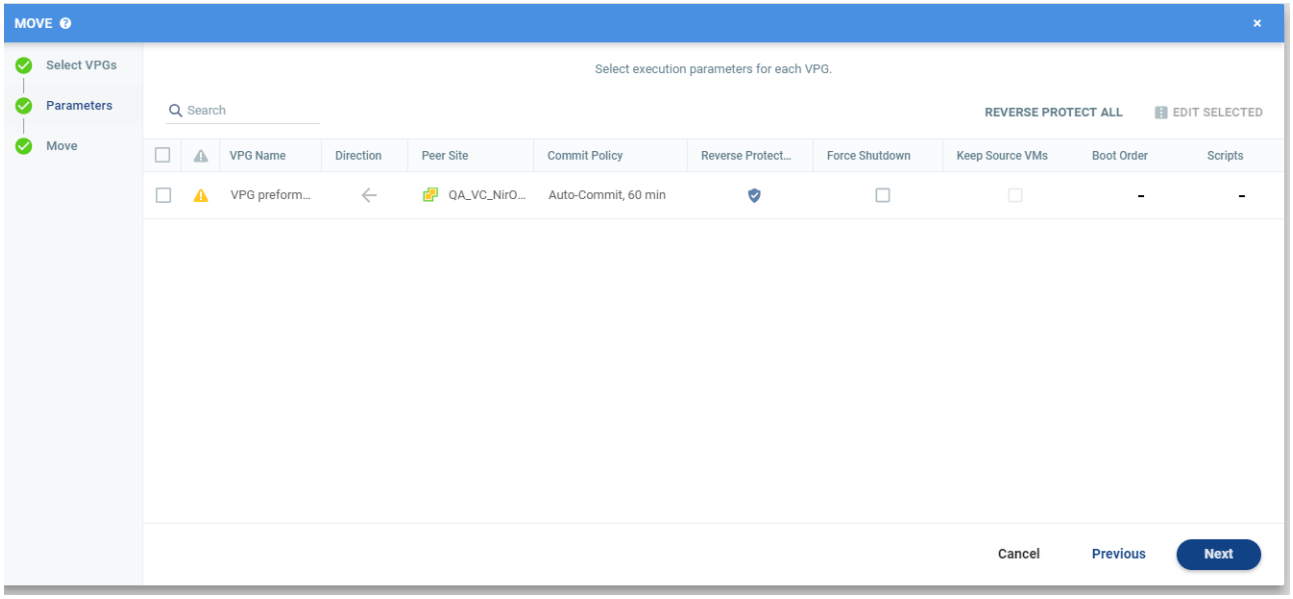

   You can change the following values to use for the recovery:
   - **Commit Policy**
   - **Force Shutdown** policy
   - **Keep Source VMs** settings

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**, with **no** testing, select **Auto-Commit** and **0** minutes.
   b) Select **None** if you do **not** want an automatic commit or rollback. You must manually commit or roll back.
   c) To test **before** committing or rolling back, specify an **amount of time** to test the recovered machines, in minutes. This is the amount of time that the commit or rollback operation is **delayed**, 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.

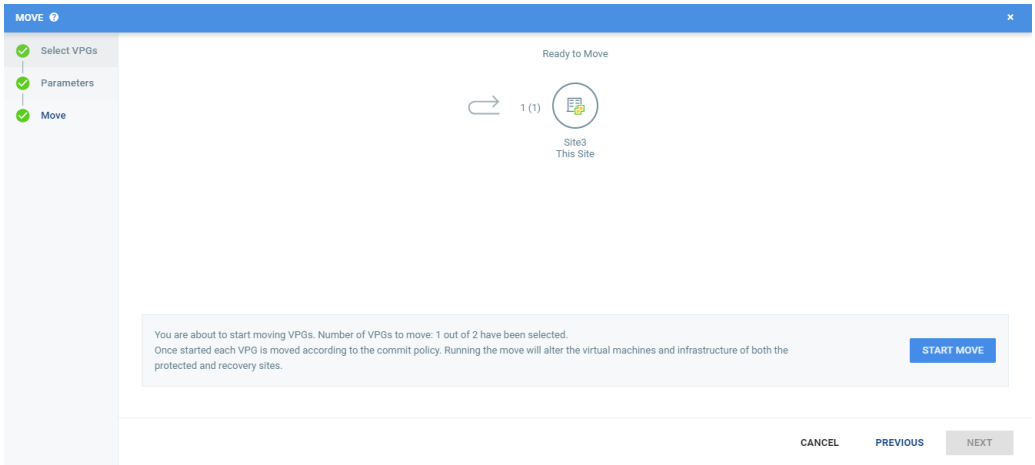

8. Click **START MOVE** to start the migration.
   A warning message appears, presenting a summary of your Commit Policy.
9. Review the Commit Policy summary, and either click **Change Settings**, or click **START MOVE** to start the migration.

   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.

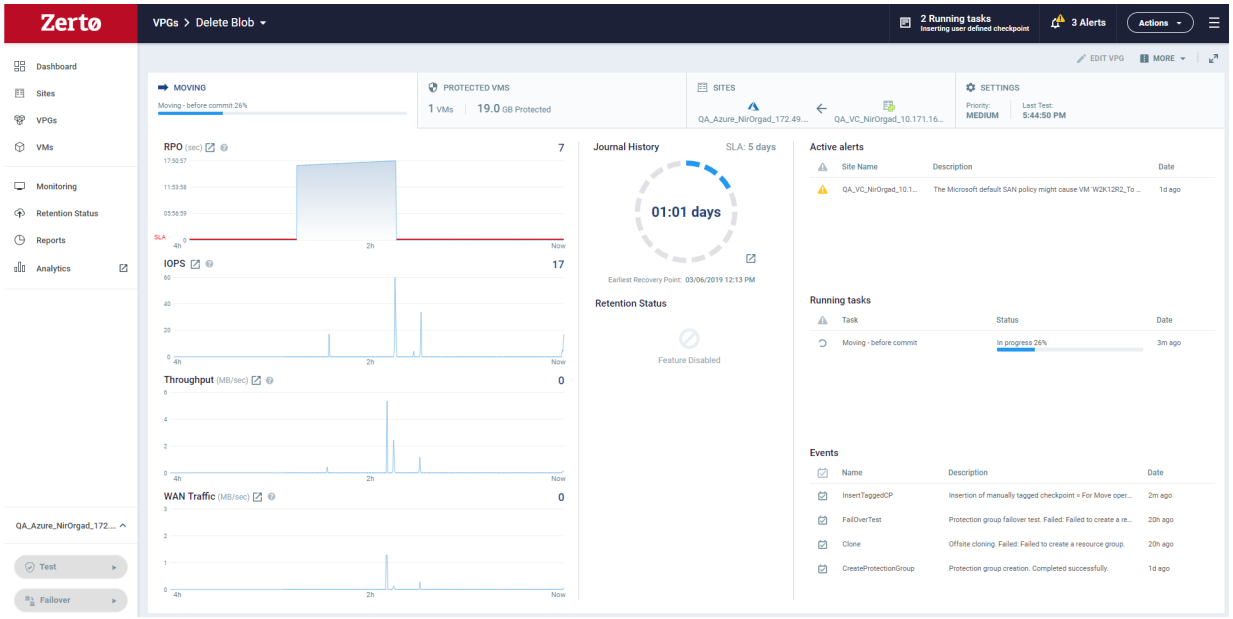

You can also commit or roll back the operation via the TASKS popup dialog in the status bar, or by selecting **MONITORING > TASKS**.

- 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 is m5.xlarge. 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 143. 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 13: MIGRATING A VPG FROM AWS

This section describes a planned migration of a VPG comprised of VMs that are either created in or moved from AWS, to a recovery site. You can migrate the VPG to the following platforms:

- VMware vSphere
- Microsoft Hyper-V
- Microsoft Azure

The disks all have to exist in the Zerto Storage Account.

The following topics are described in this section:

- "The Move Process", below
- "Moving Protected Virtual Machines From AWS to a Remote Site", on page 164
- "Reverse Protection For a Moved VPG", on page 167

The Move Process

The Migration process uses the MOVE wizard to move groups of protected virtual machines from AWS to a recovery site in a planned migration.You can initiate the MOVE operation from either the AWS site or the recovery site.

The MOVE process differs from a FAILOVER in the following ways:

- You cannot select a checkpoint to restore the virtual machine to.
- To ensure data integrity during a MOVE, the protected virtual machines are powered off completely and a final checkpoint is created so that there is no data loss before the move is implemented.

When you perform a planned migration of virtual machines to a recovery site, Zerto assumes that both sites are healthy and that you plan to relocate the virtual machines in an orderly fashion without loss of data. By setting a commit policy that enables checking the recovered machines before committing the migration, you can check the integrity of the recovered machines. If the machines are in the expected state and functioning correctly, you can commit to the migration.

The migration process has the following basic steps:

- Shutting down the protected virtual machines gracefully. This ensures data integrity.
- 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 the disk.
- Transferring all the latest changes that are still in the queue to the recovery site, including the new checkpoint.
- Creating the virtual machines at the remote site in the production network and attaching each virtual machine to its relevant virtual disks, configured to the checkpoint specified for the migration. 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 if re-IP cannot be performed.
- Powering on the virtual machines making them available to the user. If applicable, the boot order defined in the VPG settings is used to power on the machines.

  Note: If the virtual machines do not power on, the process continues and the virtual machines must be manually powered on. The virtual machines cannot be powered on automatically in a number of situations, such as when there are not enough resources in the resource pool or the required MAC address is part of a reserved range or there is a MAC address conflict or IP conflict, for example, if a clone was previously created with the MAC or IP address.
- Committing the Move operation. The default is to automatically commit the migration 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 Move, or rolled back, aborting the operation.
- The data from the journal is promoted to the machines. The machines can be used during the promotion and Zerto ensures that the user sees the latest image, even if this image, in part, includes data from the journal.
- If **Keep Source VMs** is not selected, the protected virtual machines are removed from AWS.

Moving Protected Virtual Machines From AWS to a Remote Site

You can move the virtual machines in a virtual protection group from AWS, to a remote site, where the virtual machines are replicated. As part of the process, you can also set up reverse protection, where you create a VPG on the remote site for the virtual machines being moved, pointing back to the original site. This is commonly use, for example, when the protected site has planned downtime.

To initiate a migration:

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

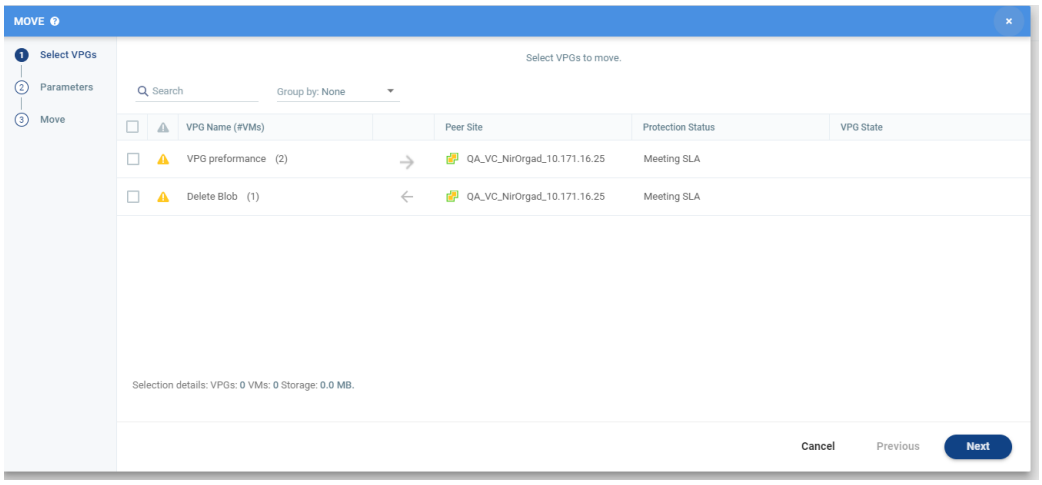

2. Select the VPGs to move. By default, all VPGs are listed.
   At the bottom, the selection details show the amount of data and the total number of virtual machines selected.
   The **Direction** arrow shows the direction of the process: **From** the protected site **To** the peer, recovery, site.
3.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

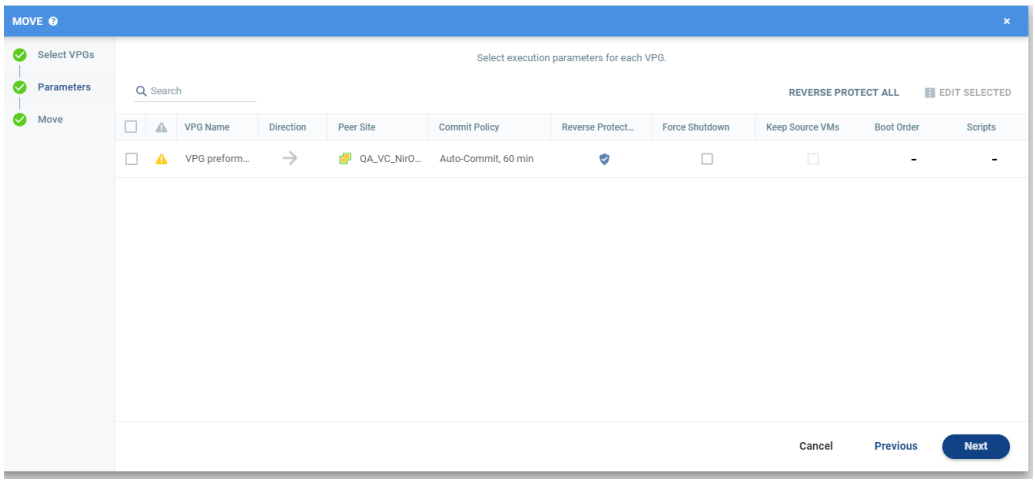

   You can change the following values to use for the recovery:
   - **Commit Policy**
   - **VM Shutdown:** Select **Yes** / **No**. Selecting YES places the VMs in a **Stopped (Deallocated)** state.
   - **Reverse Protection** settings. For more information see "Reverse Protection For a Moved VPG", on page 167.
   - **Keep Source VMs:** Prevents removal of the protected virtual machines at the AWS site.
     Note: If reverse protection is specified, the **Keep Source VMs** option is grayed out

4. To change the **Commit Policy**, click the field or select the VPG and click **EDIT SELECTED**.
   a) To commit the recovery operation **automatically**, with **no** testing, select **Auto-Commit** and **0** minutes.
   b) Select **None** if you do **not** want an automatic commit or rollback. You must manually commit or roll back.
   c) To test **before** committing or rolling back, specify an **amount of time** to test the recovered machines, in minutes. This is the amount of time that the commit or rollback operation is **delayed**, 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**.

   Testing that involves I/O is done on **scratch volumes**.
   - The more I/Os generated, the more scratch volumes are used, until the maximum size is reached, at which point no more testing can be done.
   - The maximum size of all the scratch volumes is determined by the journal size hard limit and cannot be changed.
   - The scratch volumes reside on the storage defined for the journal.
5. To specify reverse protection, whereby the virtual machines in the VPG are failed over to the recovery site and then protected in the recovery site, back to the original site, either:
   - Click **REVERSE PROTECT ALL**. This activates reverse protection on all the VPGs that you plan to Failback. The system default values for this procedure will be assigned to all the VPGs.
     -or-
   - Click the **Reverse Protection** field.

   If you want to configure the VPG for reverse protection, click the **REVERSE** link.

   The **Edit Reverse VPG** window is displayed and you can edit the reverse protection configuration. The parameters are the same as described when you create a VPG, with the following differences:
   - You cannot add or remove virtual machines in the reverse protected VPG.
   - If VMware Tools or Microsoft Integration Services are available for vSphere or Hyper-V respectively, for each virtual machine in the VPG, the IP address of the originally protected virtual machine is used. Thus, during failback the original IP address of the virtual machine on the site where the machine was originally protected is reused. However, if the machine does not contain the utility, DHCP is used. The host version must be 4.1 or higher for re-IP to be enabled.When committing the migration, you can reconfigure reverse protection, regardless of the reverse protection settings specified here. For more information see "Reverse Protection For a Moved VPG", on page 167.
6. Click **NEXT**.

   The **MOVE** step is displayed. The topology shows the VPGs and the virtual machines that are about to be moved from AWS to the recovery site.

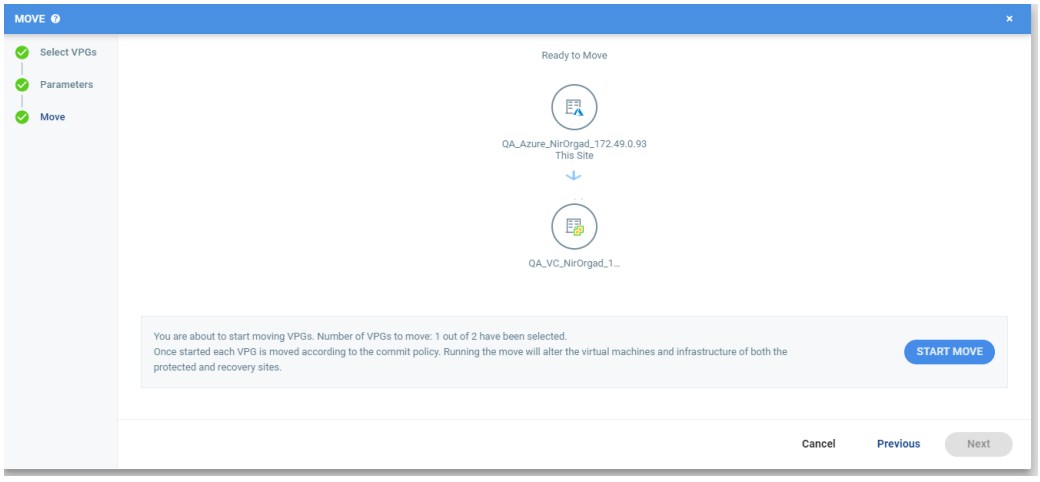

7. Click **START MOVE** to start the Move.
8. If a commit policy was set with a timeout greater than zero, as described in step 4, you can check the moved virtual machines on the recovery site before they are removed from the protected site.

Note: If a virtual machine exists on the recovery site with the same name as a virtual machine being migrated, the machine is moved 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 Move operation, is written to thin-provisioned virtual disks, one per virtual machine in the VPG. These virtual disks are automatically defined when the machines are created on the recovery site for testing. The longer the test period the more scratch volumes are used, until the maximum size is reached, at which point no more testing can be done. The maximum size of all the scratch volumes is determined by the journal size hard limit and cannot be changed. The scratch volumes reside on the storage defined for the journal. Using these scratch volumes makes committing or rolling back the Move operation more efficient.

9. 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 and, if necessary set, or reset, the reverse protection configuration. If the protected site is still up and you can set up reverse protection, you can reconfigure reverse protection by checking the **Reverse Protection** checkbox and then click the **Reverse** link. Configuring reverse protection here overwrites any of settings defined when initially configuring the move.
     - Click **Rollback** to roll back the operation, removing the virtual machines that were created on the recovery site and rebooting the machines on the protected site. The **Rollback** dialog is displayed to confirm the rollback.
   - You can also commit or roll back the operation in the **TASKS** popup dialog in the status bar or under **MONITORING** > **TASKS**.

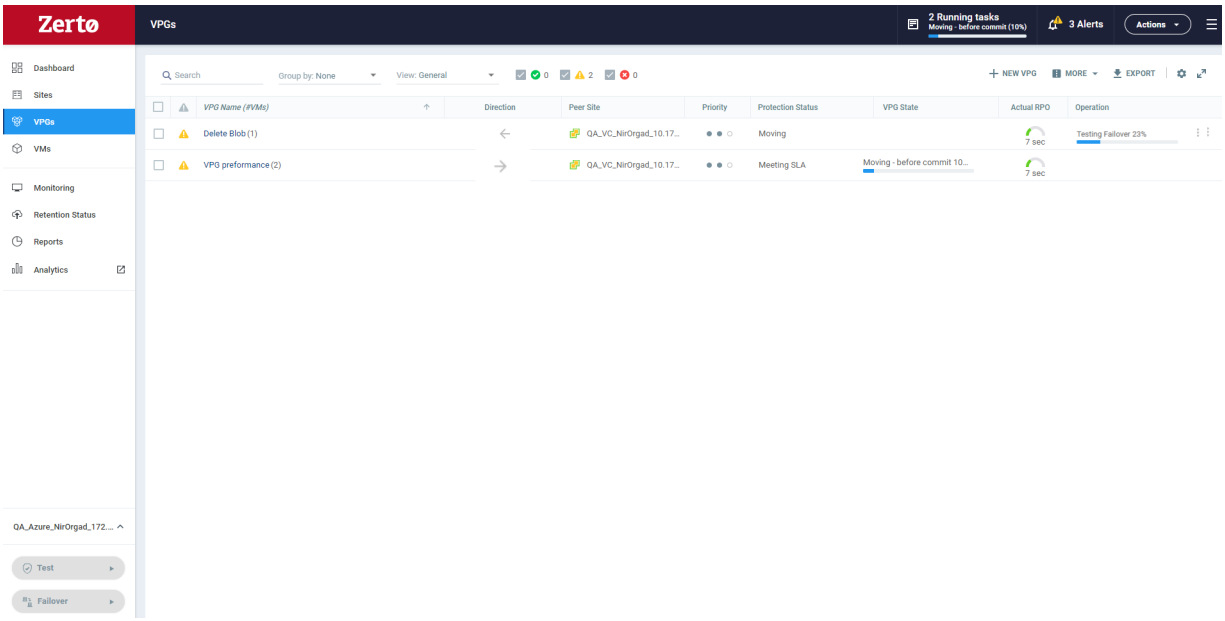

After the virtual machines are up and running and committed in the recovery site, the powered off virtual machines in the protected site are removed from AWS. Finally, data is promoted from the journal to the moved virtual machines.

During promotion of data, you cannot move a host on the moved virtual machines. If the host is rebooted during promotion, make sure that the VRA on the host is running and communicating with the Zerto Virtual Manager before starting up the recovered virtual machines.

Note: If the virtual machines do not power on, the process continues and the virtual machines must be manually powered on. The virtual machines cannot be powered on automatically in a number of situations, such as when there are not enough resources in the resource pool or the required MAC address is part of a reserved range or there is a MAC address conflict or IP conflict, for example, if a clone was previously created with the MAC or IP address.

Reverse Protection For a Moved VPG

When you specify reverse protection, the virtual machines are failed over to AWS and then protected using the values specified during the move. Data is promoted from the journal to the failed over virtual machines and then synchronization with the original site is performed so that the VPG is fully protected. A sync is required since the recovered machines can be updated while data is being promoted.

When you specify **Reverse Protection**, the following occurs:

- Reverse protection from AWS results in an Initial sync.
- The virtual machines are recovered on the recovery site and then protected using the values specified during the Move.
- The original virtual machines are removed from the original protected site.
- On the target site the data is promoted from the journal to the recovered virtual machines
- The recovery site, which is now the protected site, is synchronized with the original protected site so that the VPG is fully protected.
- A sync is required since the recovered machines can be updated while data is being promoted.

Notes:

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

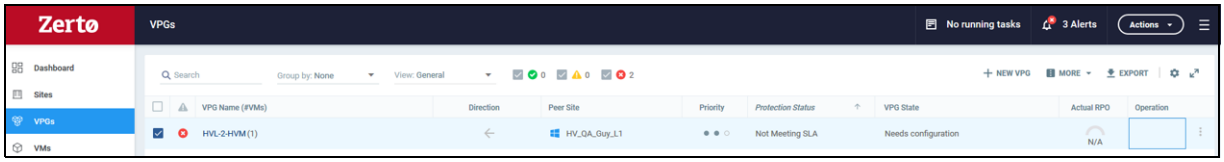

- Clicking **EDIT VPG** displays the **Edit VPG** wizard with the settings filled in, using the original settings for the virtual machines in the VPG from the original protected site, except for the volumes. To start replicating the virtual machines in the VPG, specify the disks to use for replication and optionally, make any other changes to the original settings and click **DONE**.

CHAPTER 14: MANAGING FAILOVER LIVE TO AWS

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

- "The Failover Live Process", below
- "Initiating Failover Live", on page 169
- "What Happens When the Protected Site is Down", on page 174
- "Initiating Failover Live During a Test", on page 175

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 retention process is running, first abort the retention process to enable the failover to run.

The Failover Live 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 157.

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 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.
- The failover operation creates new instances on AWS in the production network and attaches 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 is m5.xlarge. 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 143. 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 Failover Live

You can initiate a failover live, whereby the virtual machines in the virtual protection group (VPG) or specific virtual machines in a virtual protection group (VPG) 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 Failover Live During a Test", on page 175.

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 live:

1. In the Zerto User Interface click **FAILOVER**.

   The **Failover** wizard is displayed.

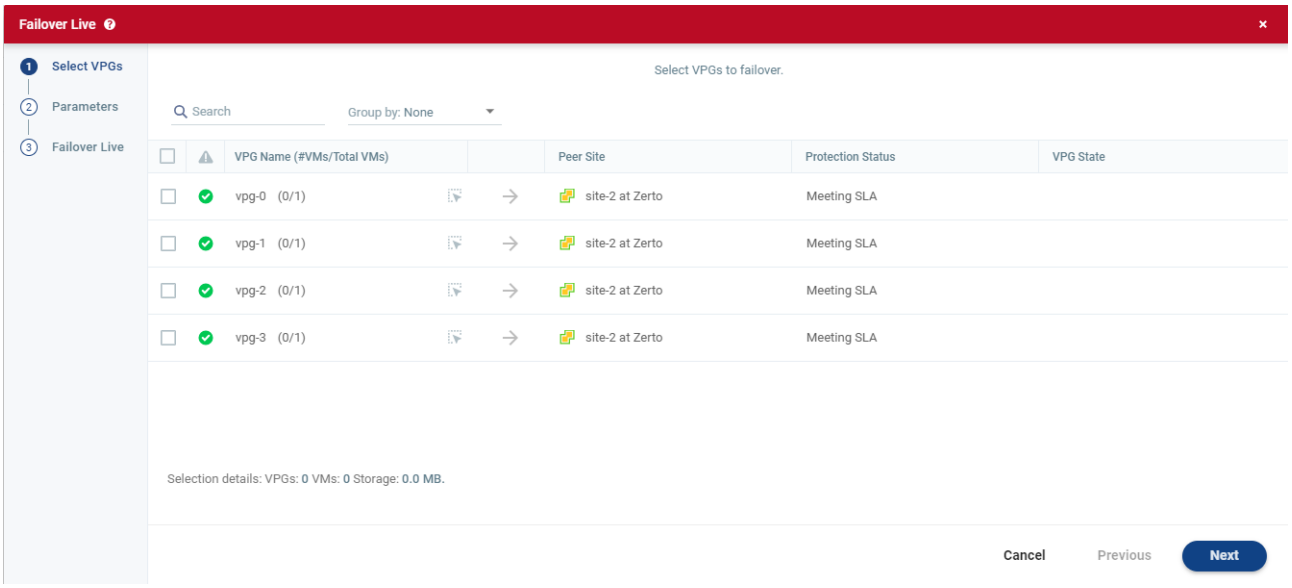

2. Select the VPGs to failover. By default, all VPGs are listed.
   a) To select specific VMs in a VPG, click the icon next to each VPG to get a list of VMs. The Select VMs to Failover dialog is displayed. By default, all VMs are selected.

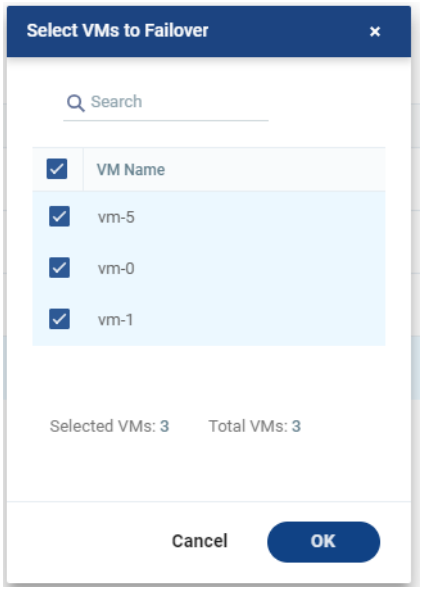

b) Select the VMs to failover.

Note: Selecting specific VMs in a VPG to failover is not supported when replicating from a **vCD site**.

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.

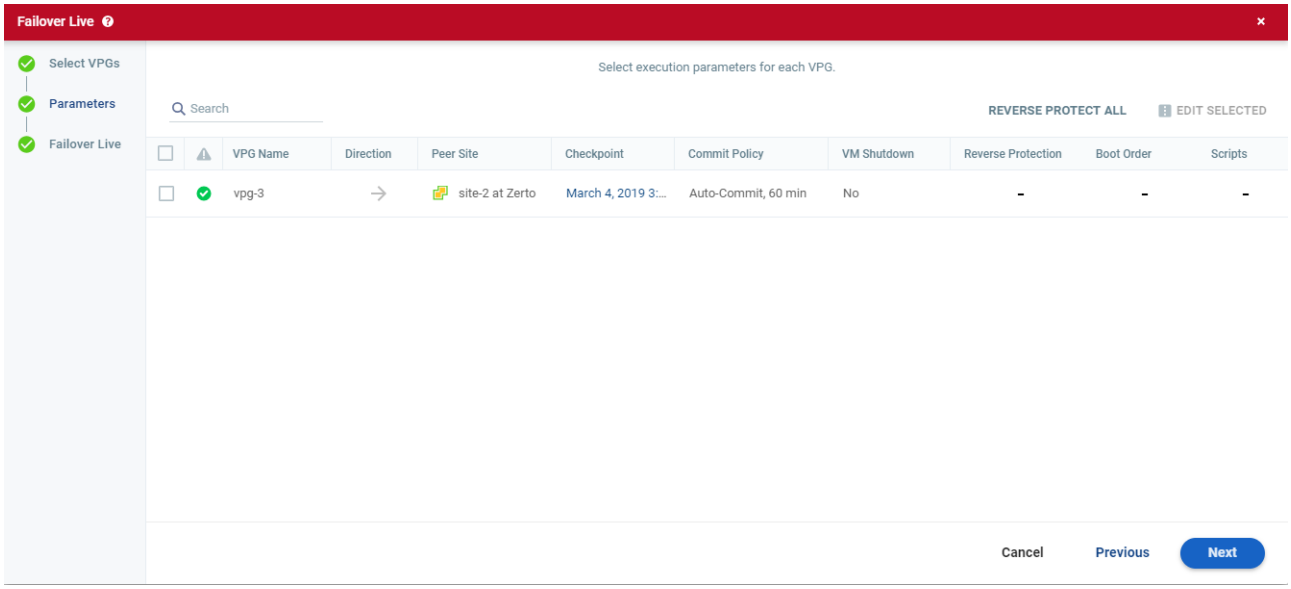

You can change the following values to use for the recovery:

- The commit policy
- The checkpoint to use
- The shutdown policy
- Reverse protection settings

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

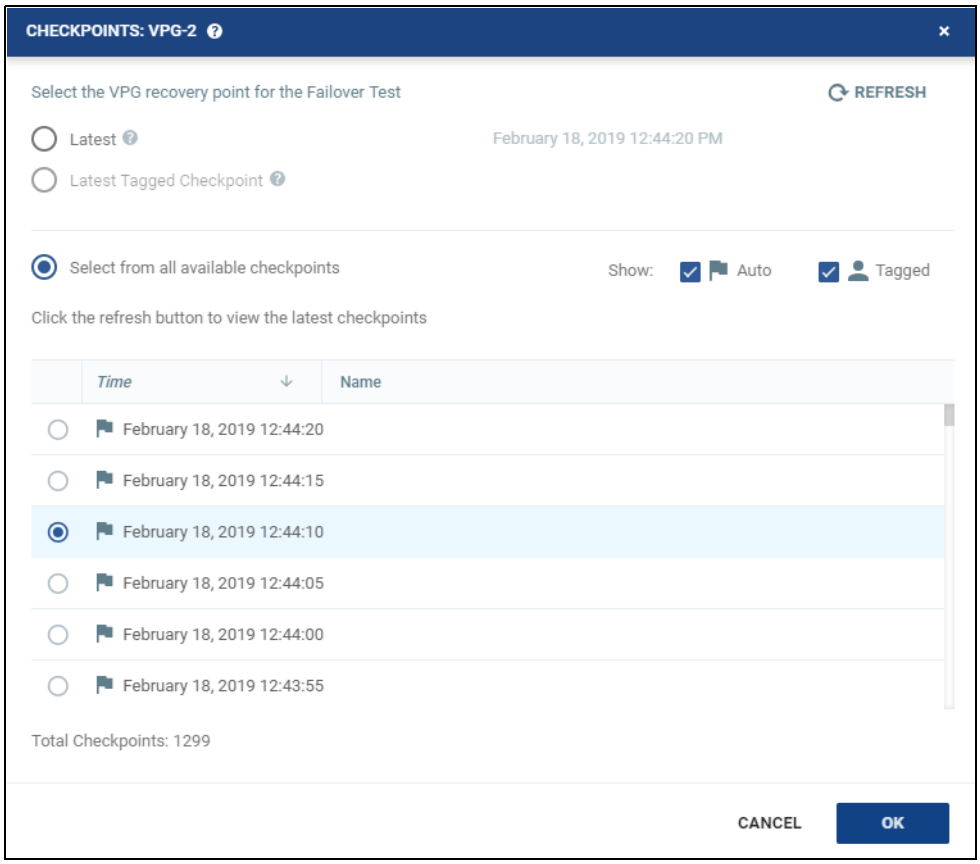

When selecting the point to recover to, remember the following:

- 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 reminder, **Click the refresh button to view the latest checkpoints** 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.

Use the following filters:

- **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.
- **Select from all available checkpoints:** If you do not want to use the latest checkpoint, 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.

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

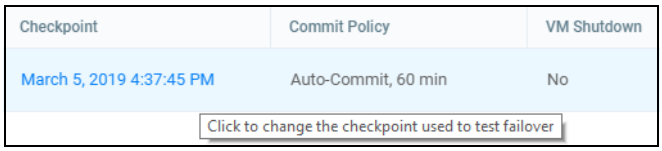

A window appears, displaying a list of the VPGs' checkpoints.

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.

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**. A warning appears informing the user if the selected VMs were not protected with the selected checkpoint.
8. To change the commit policy, click on the field or select the VPG and click **EDIT SELECTED**.
   a) To commit the recovery operation **automatically**, with **no** testing, select **Auto-Commit** and **0** minutes.
   b) Select **None** if you do **not** want an automatic commit or rollback. You must manually commit or roll back.
   c) To test **before** committing or rolling back, specify an **amount of time** to test the recovered machines, in minutes. This is the amount of time that the commit or rollback operation is **delayed**, 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**.
9. 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.
10. To specify reverse protection, whereby the virtual machines in the VPG are failed over to the recovery site and then protected in the recovery site, back to the original site, either:
    - Click **REVERSE PROTECT ALL**. This activates reverse protection on all the VPGs and/or VMs that you plan to Failback. The system default values for this procedure will be assigned to all the VPGs.
      -or-
    - Click the **Reverse Protection** field.

    If you want to configure the VPG for reverse protection, click the **REVERSE** link.
11. The **Edit Reverse VPG** window is displayed and you can edit the reverse protection configuration. The parameters are the same as described when you create a VPG, with the following differences:
    - You cannot add or remove virtual machines in the reverse protected VPG.
    -
    - If VMware Tools or Microsoft Integration Services are available for vSphere or Hyper-V respectively, for each virtual machine in the VPG, the IP address of the originally protected virtual machine is used. Thus, during failback the original IP address of the virtual machine on the site where the machine was originally protected is reused. However, if the machine does not contain the utility, DHCP is used. The host version must be 4.1 or higher for re-IP to be enabled.
    - The **Failover Test** field is empty and does not display a default. Note that although this field is empty, the Failover Test will succeed. The settings can be changed on the recovered VMs.

12. When committing the Failback, you can reconfigure reverse protection, regardless of the reverse protection settings specified here. For more information see "Reverse Protection For a Moved VPG", on page 167.
13. Click **NEXT**.
14. 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.

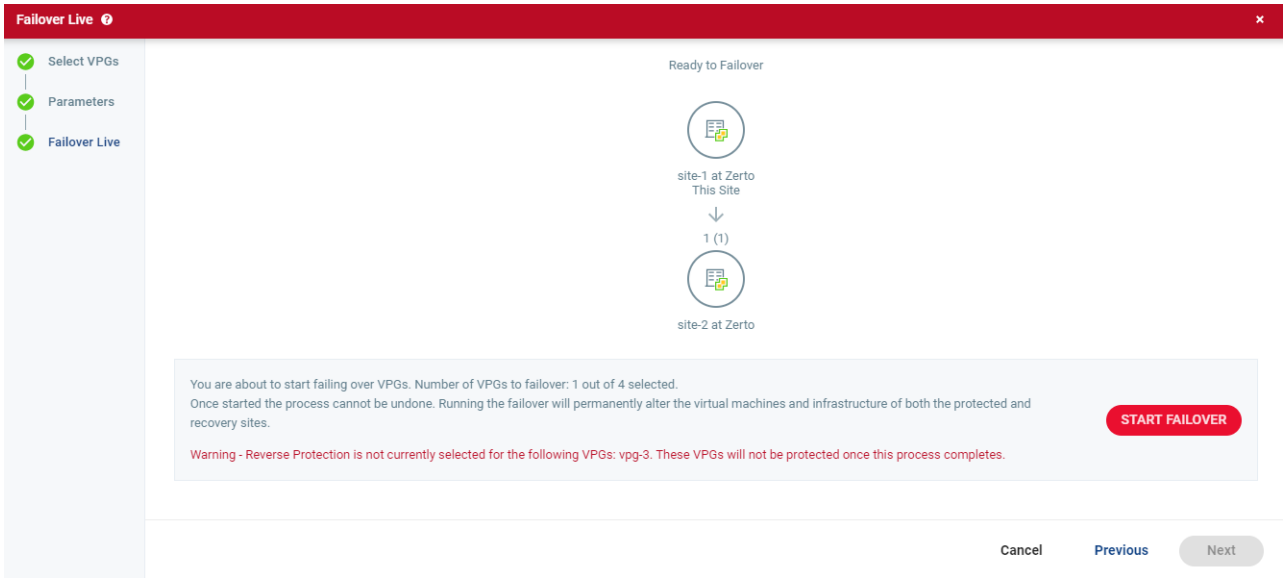

15. Click **START FAILOVER**.

    A warning message appears, presenting a summary of your Commit Policy.
16. Review the Commit Policy summary, and either click **Change Settings**, or 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.

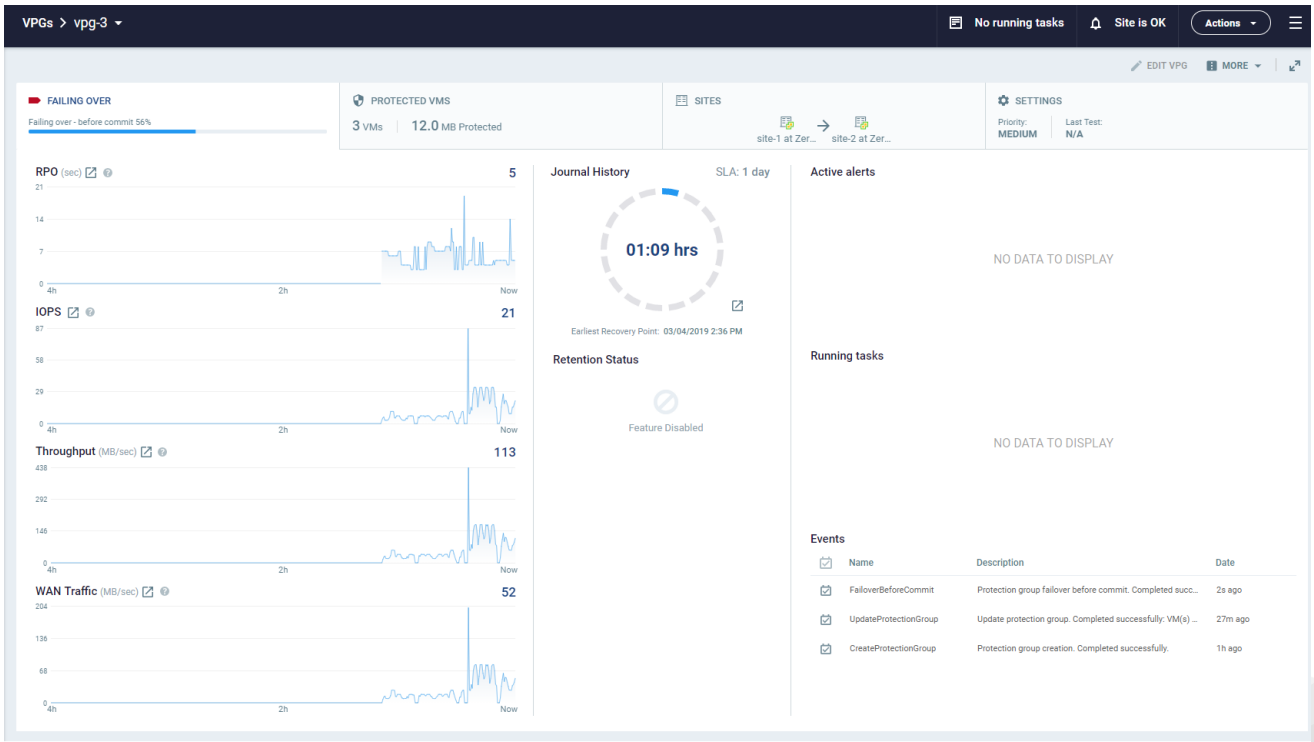

The failover starts by creating the new instances on AWS to the point-in-time specified: either the last data transferred from the protected site or to one of the checkpoints written in the journal.

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.

17. 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. When committing failover, if specific VMs in a VPG are selected, a new VPG will be created in addition to the original VPG. The additional VPG includes only the VMs selected for recovery. The new VPG name is displayed as {Original-VPG-Name**-Partial**}. The original VPG will remain intact with its history.
      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 is m5.xlarge. 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 143. 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 live:", on page 169.

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 Failover Live 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 Failover Live", on page 169. 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 15: 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 section:

- "The Clone Process", below
- "Cloning Protected Virtual Machines to the Remote Site", on page 176

Note: You cannot clone virtual machines in a VPG test while a retention process 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.

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.

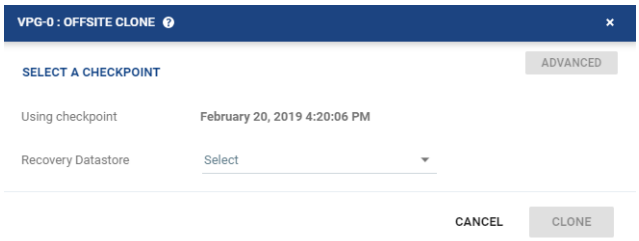

   a) To clone specific VMs, click **ADVANCED.**

The {VPG-Name}: Select VMs to Clone dialog is displayed. By default, all VMs are selected.

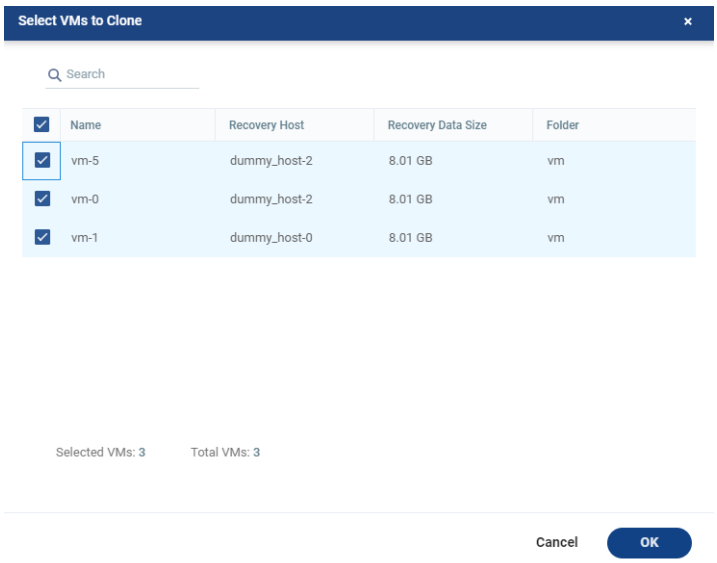

b) Select the VMs to clone.

Note: Zerto 6.0 supports cloning specific VMs if installed on the protected site only.

3. If you intend to use the last checkpoint, which is displayed in the dialog, go to step 7.

   To select the checkpoint, click **SELECT A CHECKPOINT**.

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

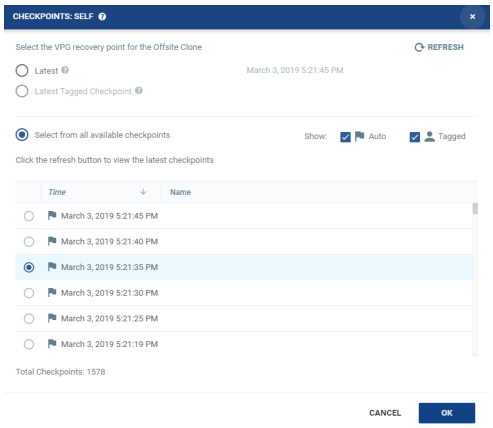

When selecting the point to recover to:

- 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.

5. 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.
6. Click **OK**.
7. 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 16: RECOVERING FILES AND FOLDERS

You can recover specific files and folders from the recovery site for virtual machines that are being protected by Zerto. 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", below
- "Recovering Files and Folders", on page 180

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 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:

1. Select the virtual machine that is protected, on which the files or folders to recover are located.
2. Select the checkpoint, at which the files and folders will be recovered.
3. Select the disk, which contains the files and folders to recover.

   Note: You can only recover files and folders from one disk at a time.
4. Mount the selected disk.
5. Select the files and folders on the disk to recover.
6. Download 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.
7. Unmount 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. Zerto does not support mounting the same volume twice, for example if you want a file from two different checkpoints.

Considerations:

- You cannot recover files or folders from a virtual machine when a test failover, live failover, move, clone, or retention process is being performed on a VPG that contains the virtual machine.
- You cannot recover files or folders from the Zerto plugin.

For additional file and folder recovery considerations or known issues, see the Release Notes > Known Issues > File Level Restore.

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 180

   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.
   - You can perform a test failover or clone.
2. "Downloading the Files and Folders from the Disk", on page 183

   Note: When downloading, you may see grayed out files or folders. These grayed out files or folders are not supported, and therefore cannot be selected. For a list of unsupported files, see Zerto Virtual Replication - Interoperability Matrix for All Versions.

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.

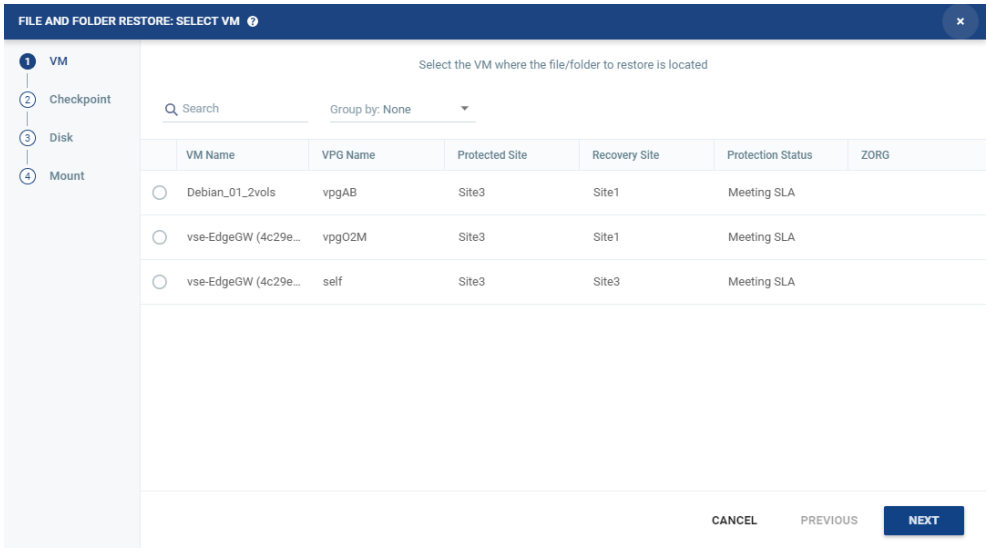

   The list of all protected virtual machines is displayed. You can only recover files or folders from one virtual machine at a time.
2. Select the virtual machine on which the file or folder is located and click **NEXT**.

The CHECKPOINT step is displayed. By default, all checkpoints are displayed.

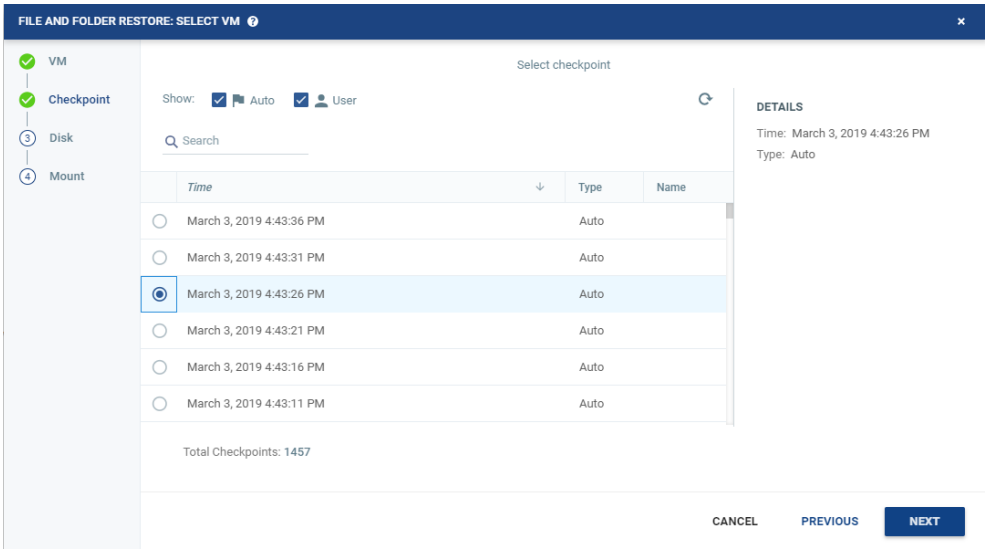

3. Select the checkpoint from which to recover the file or folder.
   - **Auto:** Checkpoints generated by the Zerto Virtual Manager are displayed.
   - **User:** Checkpoints that were added by a user, or were added by the Zerto Virtual Manager when a failover test was performed on the VPG that included 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.

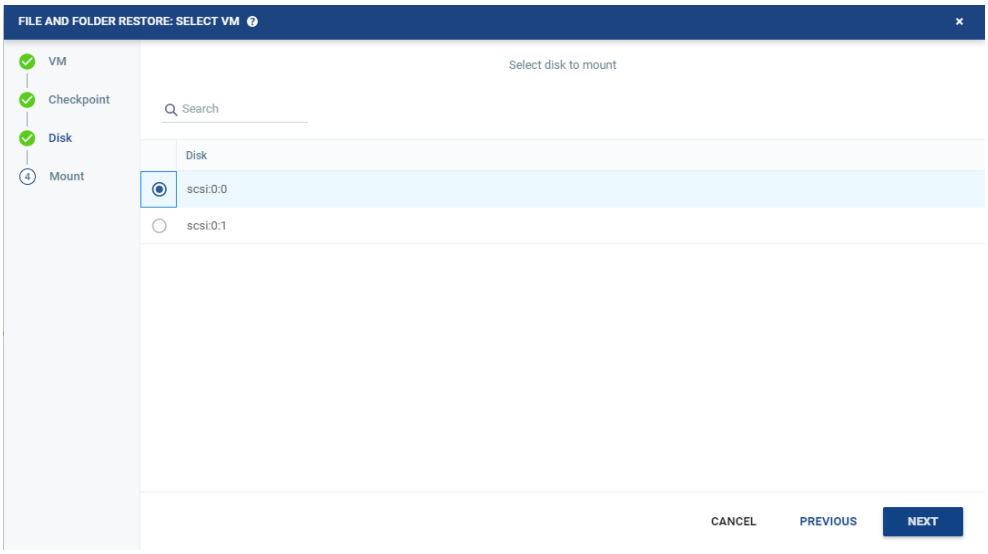

5. Select a disk to mount and click **NEXT**.

The MOUNT summary step is displayed. The settings you selected are displayed.

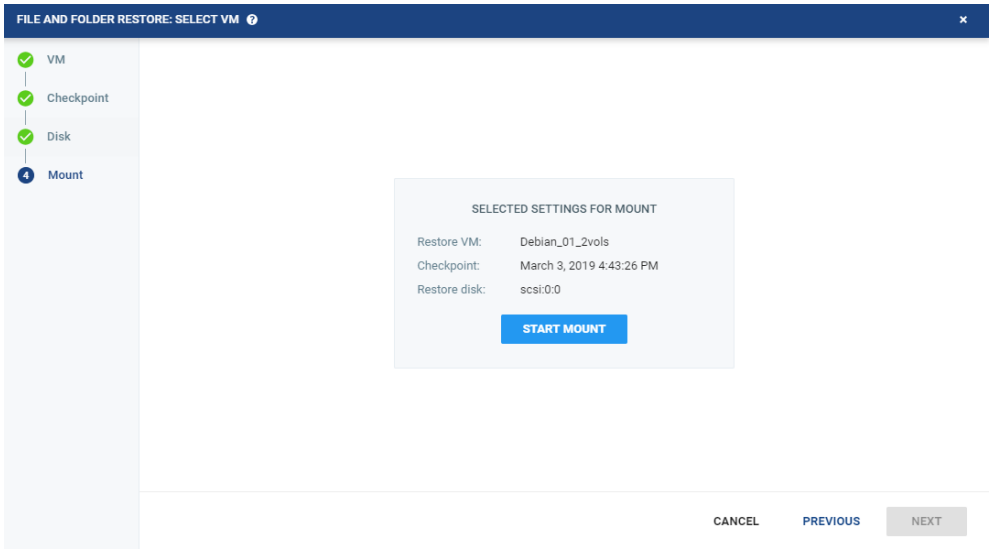

6. Click **START MOUNT** to mount the disk.
   - When the mount is completed, icons appear next to the task.
   - By clicking the folder icon ( ) you can browse the folders and files on the disk.
   - By clicking the unmount icon ( ) you can unmount the disk without restoring any files or folders.

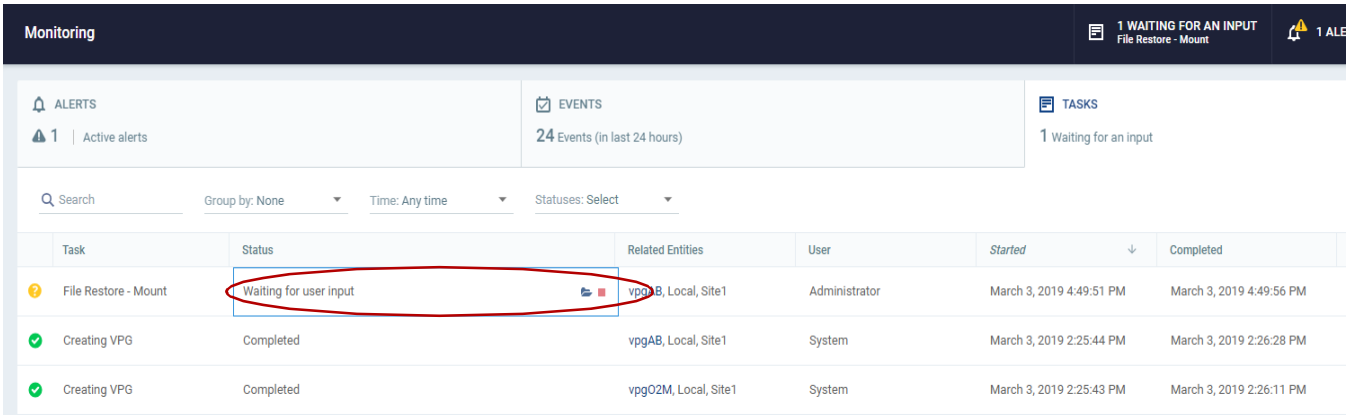

7. Continue with "Downloading the Files and Folders from the Disk" on page 183.

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 ( ).

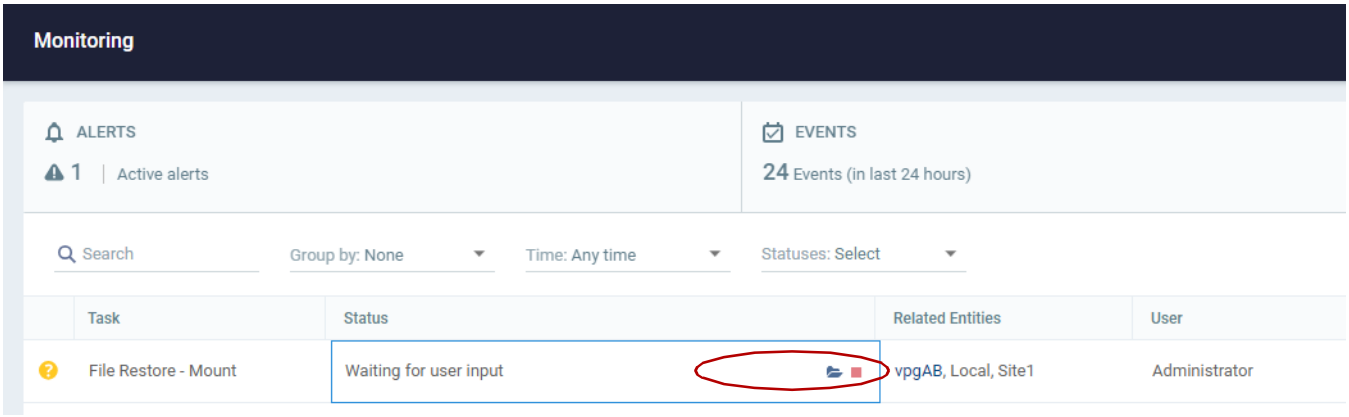

The File and Folder Restore dialog is displayed.

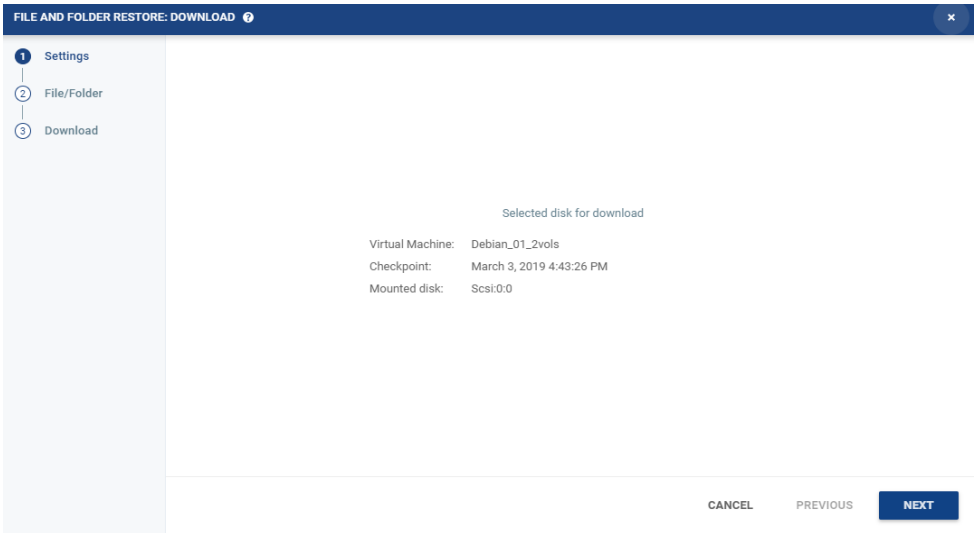

2. Click **NEXT**.

The FILE/FOLDER step is displayed.

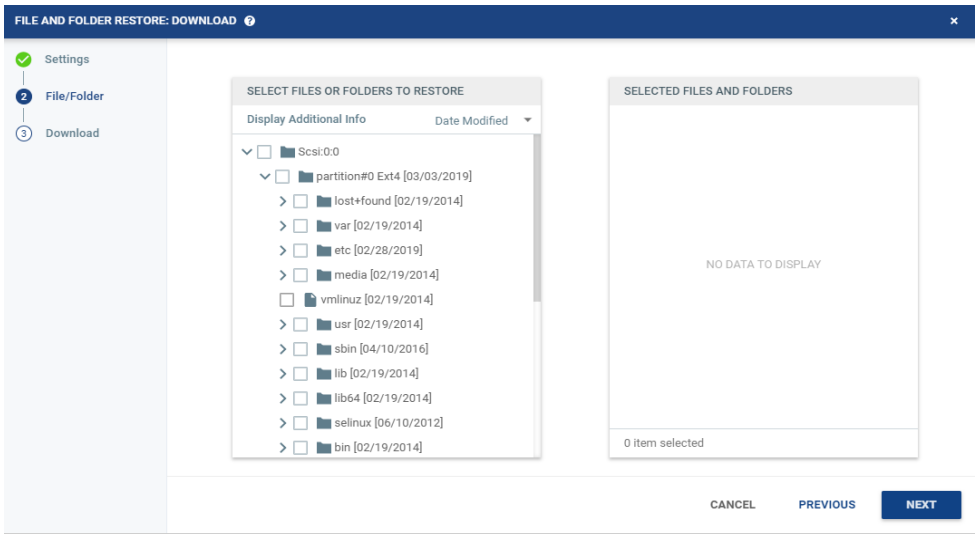

- Select the files and folders you want to download.
  - Grayed out file systems are not supported therefore cannot be selected. For a list of unsupported file system, see the Interoperability Matrix.
  - When a selected folder contains unsupported files, only the supported files will successfully download. Meaning, only the supported files will appear in the downloaded zip file

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.

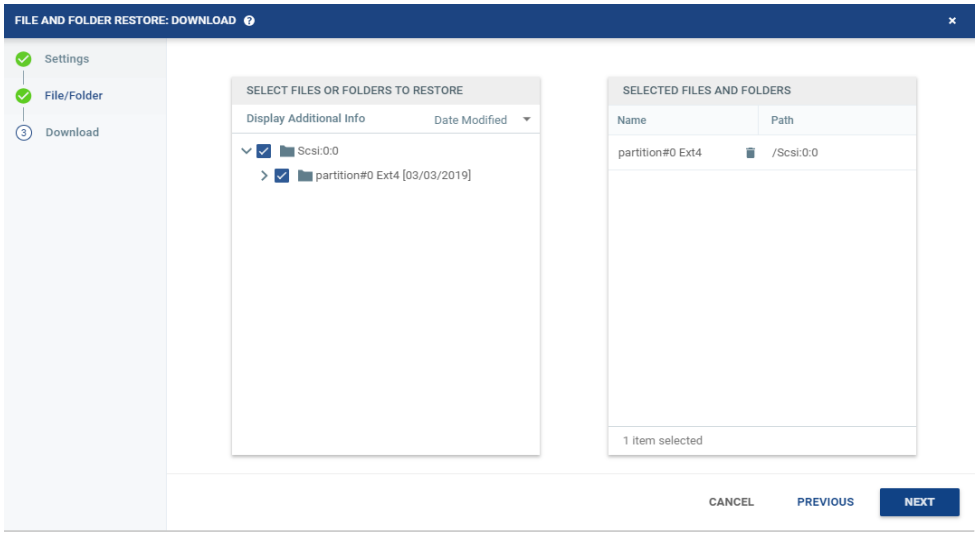

3. Click **NEXT**.

   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.

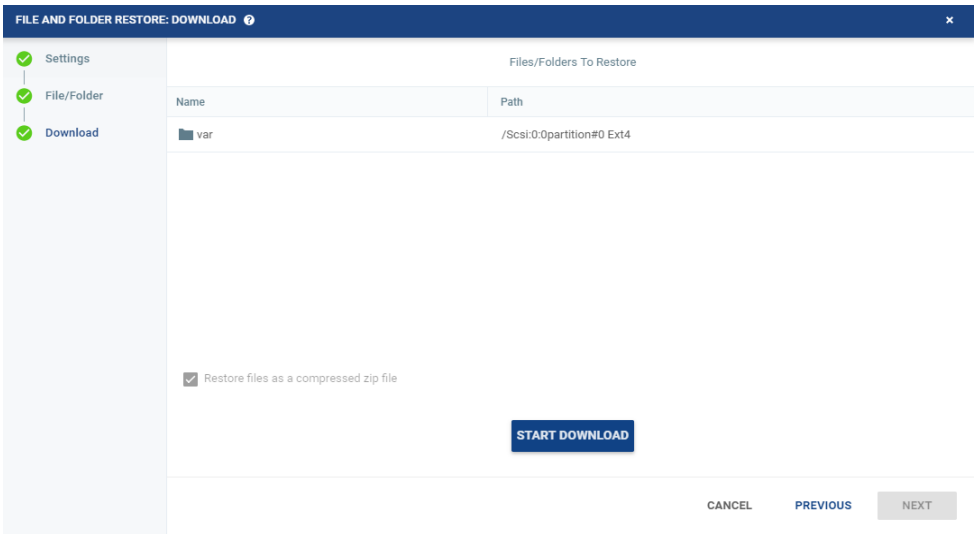

4. Click **START DOWNLOAD**.

   The files and folders are downloaded by default to the downloads folder on the computer where you run the Zerto User Interface.

   When a selected folder contains unsupported files, only the supported files will successfully download. Meaning, only the supported files will appear in the download folder.

   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`.

5. **Unmount the disk** after the files or folders are downloaded. To unmount the disk, click the unmount icon (■).

CHAPTER 17: FAILING BACK FROM AWS

If you have moved, or failed over a VPG to AWS, you can failback these VPGs to the following platforms:

- VMware vSphere
- Microsoft Hyper-V
- Microsoft Azure

You can failback both Windows and Linux machines.

The following topics are described in this section:

- "The Failback (Move) Process", below
- "Initiating a Failback (Move)", on page 187
- "Reverse Protection For a Moved VPG", on page 190

The Failback (Move) Process

The failback process uses the MOVE wizard, following a disaster, to recover failed over virtual machines from AWS, back to the original site. You can initiate the MOVE operation from either the AWS site or the recovery site.

The MOVE process differs from a FAILOVER in the following ways:

- You cannot select a checkpoint to restore the virtual machine to.
- To ensure data integrity during a MOVE, the protected virtual machines are powered off completely and a final checkpoint is created so that there is no data loss before the move is implemented.

When you perform a planned migration of virtual machines to a recovery site, Zerto assumes that both sites are healthy and that you plan to relocate the virtual machines in an orderly fashion without loss of data. By setting a commit policy that enables checking the recovered machines before committing the Failback (Move), you can check the integrity of the recovered machines. If the machines are in the expected state and functioning correctly, you can commit the MOVE.

The Failback (MOVE) process has the following basic steps:

- Creating the virtual machines at the remote site in the production network and attaching each virtual machine to its relevant virtual disks, configured to the checkpoint specified for the failback. 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 if re-IP cannot be performed.
- Powering on the virtual machines making them available to the user. If applicable, the boot order defined in the VPG settings is used to power on the machines.

  Note: If the virtual machines do not power on, the process continues and the virtual machines must be manually powered on. The virtual machines cannot be powered on automatically in a number of situations, such as when there are not enough resources in the resource pool or the required MAC address is part of a reserved range or there is a MAC address conflict or IP conflict, for example, if a clone was previously created with the MAC or IP address.
- The default is to automatically commit the Failback (Move) 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 Failback (Move), or rolled back, aborting the operation.
- The data from the journal is promoted to the machines. The machines can be used during the promotion and Zerto ensures that the user sees the latest image, even if this image, in part, includes data from the journal.
- If **Keep Source VMs** is not selected, the protected virtual machines are removed from AWS.

- Based on the Instance Size and Type selected when creating a VPG from on-premise to AWS,upon failback from AWSto on-premise, the virtual machine will inherit the CPU and Memory from the selected AWSinstance.

Failback (Move) After the Original Site is Operational

To Failback (Move) to the original protected site, the VPG that is now protecting the virtual machines at the recovery site have to be configured, if reverse protection was not selected, and then a **Delta Sync** is performed with the disks at the original protected site. Once the VPG is in a protecting state the virtual machines can be moved back to the original protected site, as described in "Migrating a VPG to a Recovery Site", on page 335.

Initiating a Failback (Move)

You can initiate a Failback (Move), whereby the virtual machines in the virtual protection group (VPG) are replicated to a set checkpoint in the recovery site. As part of the process you can also set up reverse protection, whereby you create a VPG on the recovery machine for the virtual machines being replicated.

You can initiate a Failback (Move) to the last checkpoint recorded in the journal, even if the protected site is no longer up. You can initiate a Failback (Move) during a test.

If the protected site is operational, the Failback (Move) can be initiated from the protected site, however if the protected site is down, you can initiate the Failback (Move) from the recovery site.

To initiate a Failback:

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

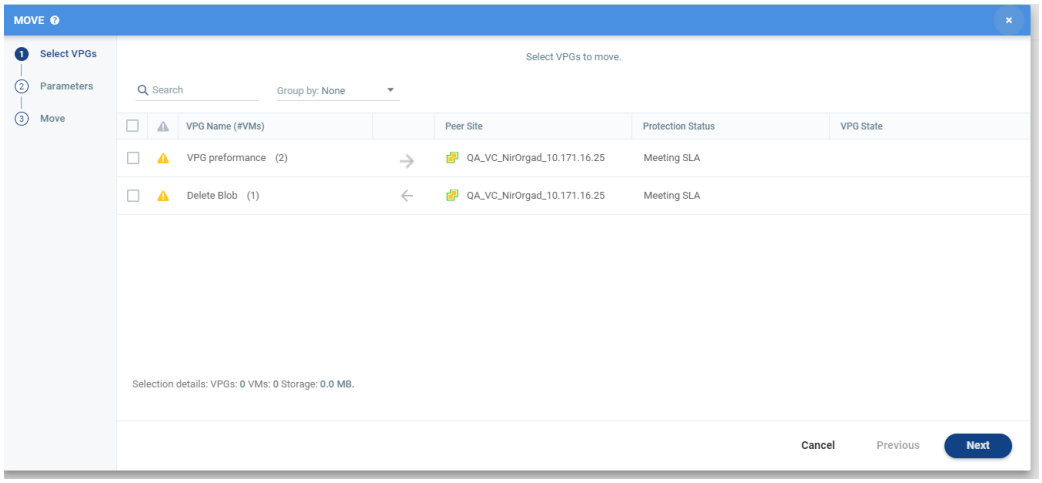

2. Select the VPGs to failback. By default, all VPGs are listed.
   At the bottom, the selection details show the amount of data and the total number of virtual machines selected.
   The Direction arrow shows the direction of the process: **From** the protected site **To** the peer, recovery, site.
3. Click **NEXT**.

The **EXECUTION PARAMETERS** step is displayed.

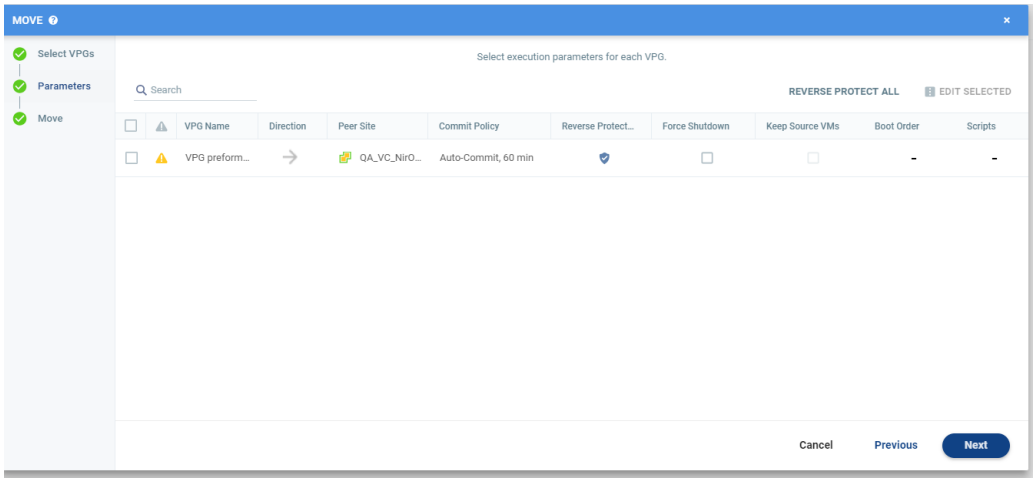

You can change the following values to use for the recovery:

- **Commit Policy**
- **Force Shutdown**: Select **Yes** / **No**. Selecting YES places the VMs in a **Stopped (Deallocated)** state.
- **Reverse Protection** settings. For more information see "Reverse Protection For a Moved VPG", on page 190.
- **Keep Source VMs**: Prevents removal of the protected virtual machines at the AWS site.
  Note: If reverse protection is specified, the **Keep Source VMs** option is grayed out.

4. To change the **commit policy**, click the field or select the VPG and click **EDIT SELECTED**.
   a) To commit the recovery operation **automatically**, with **no** testing, select **Auto-Commit** and **0** minutes.
   b) Select **None** if you do **not** want an automatic commit or rollback. You must manually commit or roll back.
   c) To test **before** committing or rolling back, specify an **amount of time** to test the recovered machines, in minutes.
   This is the amount of time that the commit or rollback operation is **delayed**, 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**.

   Testing that involves I/O is done on **scratch volumes**.
   - The more I/Os generated, the more scratch volumes are used, until the maximum size is reached, at which point no more testing can be done.
   - The maximum size of all the scratch volumes is determined by the journal size hard limit and cannot be changed.
   - The scratch volumes reside on the storage defined for the journal.
5. To specify reverse protection, whereby the virtual machines in the VPG are failed over to the recovery site and then protected in the recovery site, back to the original site, either:
   - Click **REVERSE PROTECT ALL**. This activates reverse protection on all the VPGs that you plan to Failback. The system default values for this procedure will be assigned to all the VPGs.
     -or-
   - Click the **Reverse Protection** field.

   If you want to configure the VPG for reverse protection, click the **REVERSE** link.

   The **Edit Reverse VPG** window is displayed and you can edit the reverse protection configuration. The parameters are the same as described when you create a VPG, with the following differences:
   - You cannot add or remove virtual machines in the reverse protected VPG.
   - By default, reverse protection is to the original protected disks. You can specify a different storage to be used for the reverse protection.
   - If VMware Tools or Microsoft Integration Services are available for vSphere or Hyper-V respectively, for each virtual machine in the VPG, the IP address of the originally protected virtual machine is used. Thus, during failback the original IP address of the virtual machine on the site where the machine was originally protected is reused. However, if the machine does not contain the utility, DHCP is used. The host version must be 4.1 or higher for re-IP to be enabled.

   When committing the Failback, you can reconfigure reverse protection, regardless of the reverse protection settings specified here. For more information see "Reverse Protection For a Moved VPG", on page 190.

6. Click **NEXT**.

   The **MOVE** step is displayed. The topology will show the VPGs and the virtual machines that are about to be moved from AWS to the original site.

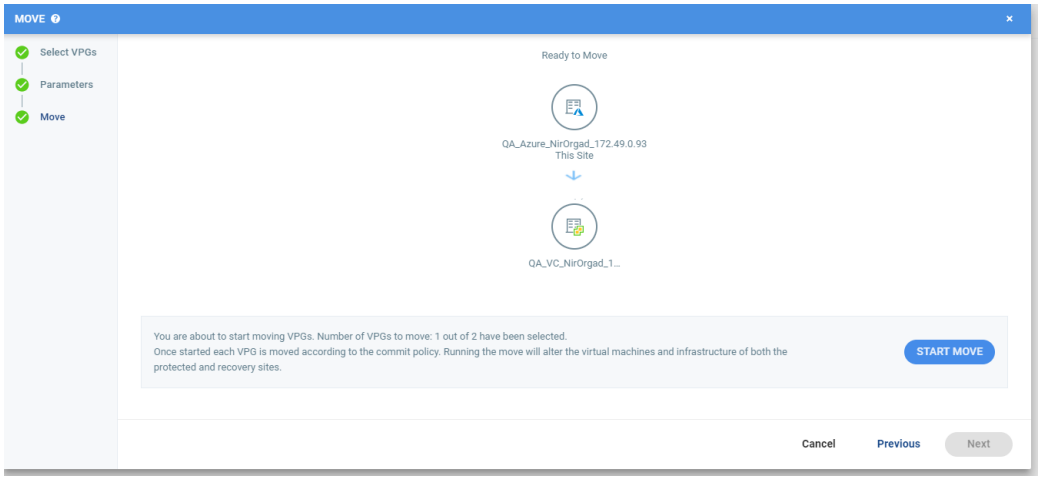

7. Click **START MOVE** to start the Failback.
8. If a commit policy was set with a timeout greater than zero, as described in step 4, you can check the moved virtual machines on the recovery site before they are removed from the protected site.

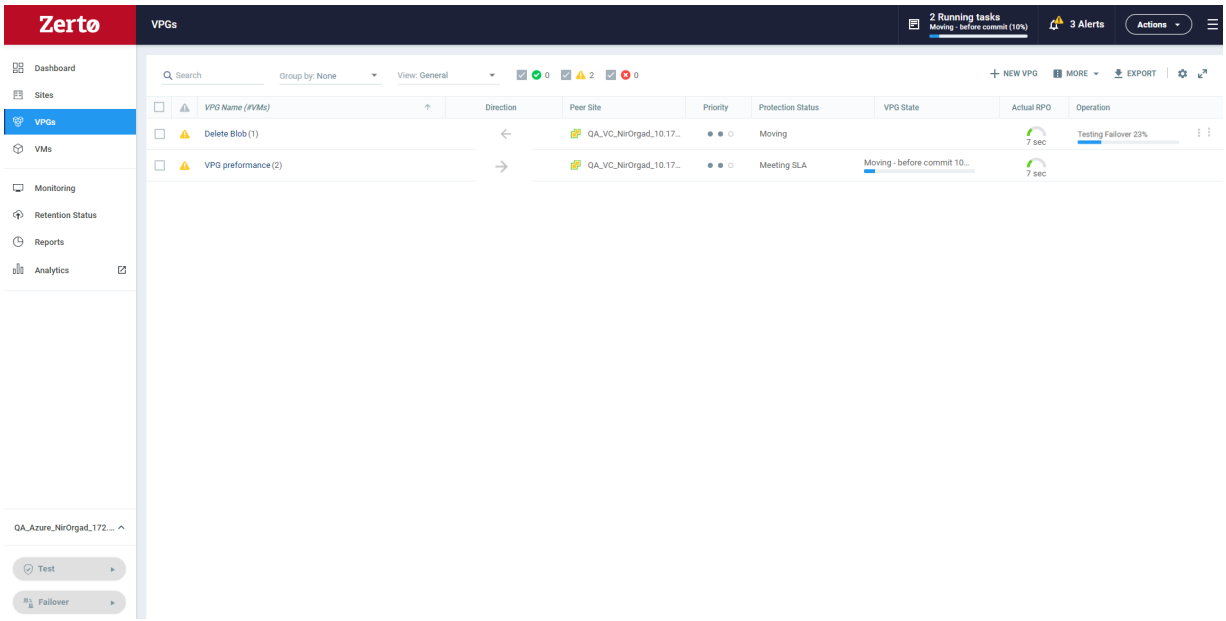

Note: If a virtual machine exists on the recovery site with the same name as a virtual machine being migrated, the machine is moved and named in the peer site with a number added as a suffix to the name, starting with the number 1

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 Move operation, is written to thin-provisioned virtual disks, one per virtual machine in the VPG. These virtual disks are automatically defined when the machines are created on the recovery site for testing. The longer the test period the more scratch volumes are used, until the maximum size is reached, at which point no more testing can be done. The maximum size of all the scratch volumes is determined by the journal size hard limit and cannot be changed. The scratch volumes reside on the storage defined for the journal. Using these scratch volumes makes committing or rolling back the Move operation more efficient.

Note: You cannot take a snapshot of a virtual machine before the Move 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. Taking a snapshot of a machine that is in the process of being moved will corrupt that machine.

9. 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 and, if necessary set, or reset, the reverse protection configuration. If the protected site is still up and you can set up reverse protection, you can reconfigure reverse protection by checking the Reverse Protection checkbox and then click the Reverse link. Configuring reverse protection here overwrites any of settings defined when initially configuring the failover.
     - 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 in the TASKS popup dialog in the status bar or under MONITORING > TASKS.

After the virtual machines are up and running and committed in the recovery site, the powered off virtual machines in the protected site are removed from AWS. Finally, data is promoted from the journal to the moved virtual machines.

During promotion of data, you cannot move a host on the moved virtual machines. If the host is rebooted during promotion, make sure that the VRA on the host is running and communicating with the Zerto Virtual Manager before starting up the recovered virtual machines.

Note: If the virtual machines do not power on, the process continues and the virtual machines must be manually powered on. The virtual machines cannot be powered on automatically in a number of situations, such as when there are not enough resources in the resource pool or the required MAC address is part of a reserved range or there is a MAC address conflict or IP conflict, for example, if a clone was previously created with the MAC or IP address.

Reverse Protection For a Moved VPG

When you specify reverse protection, the virtual machines are failed over to AWS and then protected using the values specified during the failback. Data is promoted from the journal to the failed back virtual machines and then synchronization with the original site is performed so that the VPG is fully protected A sync is required since the recovered machines can be updated while data is being promoted.

When you specify Reverse Protection, the following occurs:

- Reverse protection from AWS results in an Initial sync.
- The virtual machines are recovered on the recovery site and then protected using the values specified during the failover.
- The original virtual machines are removed from the original protected site.
- On the target site the data is promoted from the journal to the recovered virtual machines.
- The recovery site, which is now the protected site, is synchronized with the original protected site so that the VPG is fully protected.

Notes:

- For the Move operation to complete successfully, when Reverse Protection is specified, the original protected site must have enough storage available to fail back the moved virtual machines.
- 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.
- If you do not specify Reverse Protection, the VPG definition is kept with the status Needs Configuration and the reverse settings in the VPG definition are not set.
- Clicking **EDIT VPG** displays the Edit VPG wizard with the settings filled in, using the original settings for the virtual machines in the VPG from the original protected site, except for the volumes. To start replicating the virtual machines in

the VPG, specify the disks to use for replication and optionally, make any other changes to the original settings and click **DONE**.

Zerto includes reporting for the following:

- "Recovery Reports", below
- "Resources Report", on page 193
- "VPG Performance", on page 197

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**.

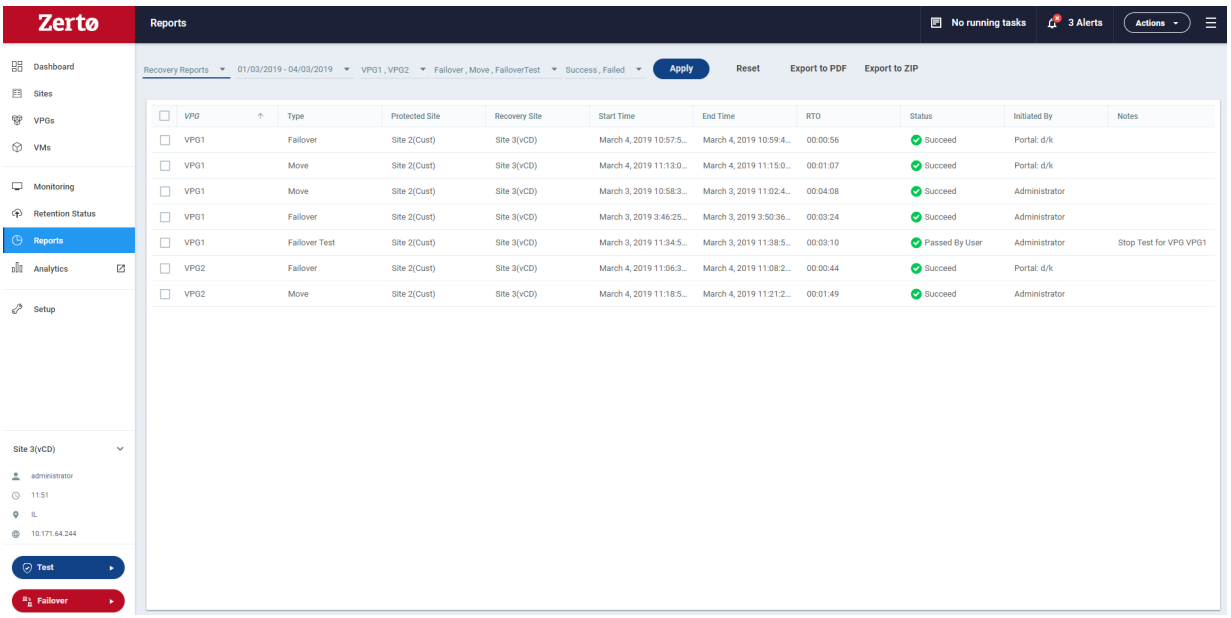

You can filter the tests by the following:

- **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 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 the 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 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 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 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.
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.

PARAMETER	DESCRIPTION
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 Policy	The target vCD storage policy 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.
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.

PARAMETER	DESCRIPTION
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 193.
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.
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.

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 section:

- "Ensuring the Zerto Virtual Manager is Running", below
- "Troubleshooting: "Needs Configuration" Problems", on page 198
- "Troubleshooting VRA Problems", on page 199
- "Zerto Diagnostics Utility", on page 199
- "Collecting Zerto Logs", on page 199
- "Understanding the Logs", on page 205

For details about Zerto Virtual Manager alarms, alerts, and events, refer to *Zerto 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 is installed, that the Zerto Virtual Manager Windows service is started.

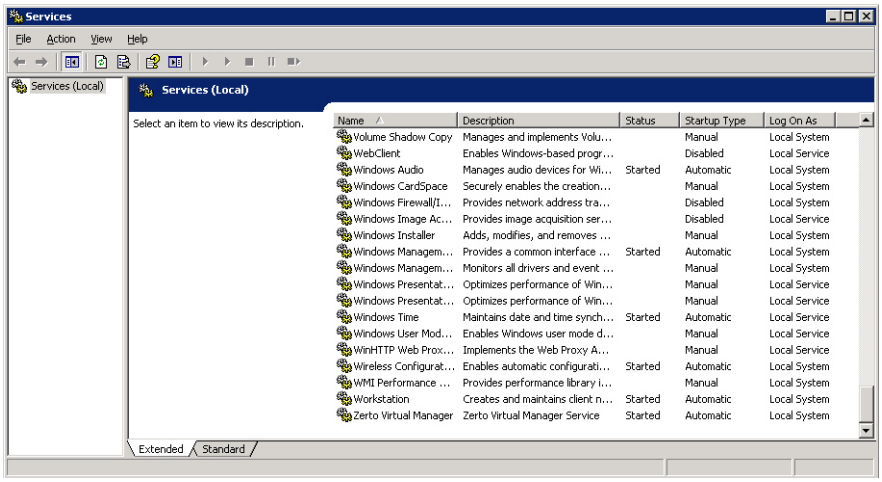

Troubleshooting: "Needs Configuration" Problems

When the VPG status changes to `Needs Configuration`, the settings in the VPG need to be checked and, when necessary, updated.

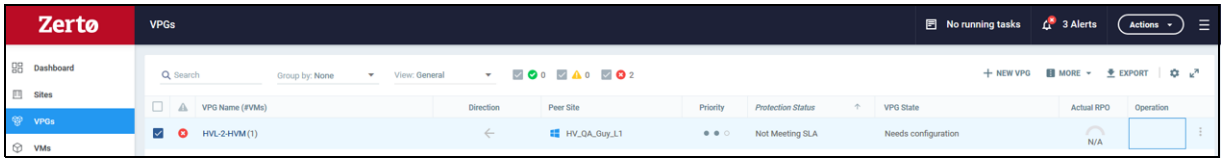

The following scenarios result in the VPG status changing to Needs Configuration:

- 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 Diagnostics Utility

The Zerto solution 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 203.
- Check the connectivity between Zerto 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. See "Reconfiguring the Microsoft SQL Server Database Used by the Zerto Virtual Manager", on page 134.

Note: A separate installation kit is available for download from the Zerto Support Portal downloads page that installs the Zerto Diagnostics utility as a standalone utility on any Windows machine that has Microsoft .NET Framework 4.7.2 installed. The installation executable is included as part of the standalone utility installation kit and it requires an additional 1.8GB of free disk space.

Collecting Zerto Logs

Virtual replication logs can be collected to help Zerto support resolve problems related to Zerto. Virtual replication logs can be collected in the following ways:

"Using Remote Log Collection", below

"Using the Zerto Diagnostics Tool", on page 200

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.

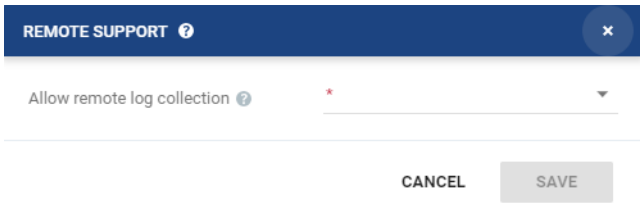

2. Click the drop down menu to display the remote log collection options.
3. Select the remote log collection option you wish to allow:
   - **Never:** Remote log collection is not allowed (default). If remote log collection is currently allowed, the remote connection will be terminated if you select 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 the duration of a specific case**: You will be prompted to select the Case number from the drop-down list. The list contains all the active cases opened under the account that the Zerto Virtual Manager is registered to.

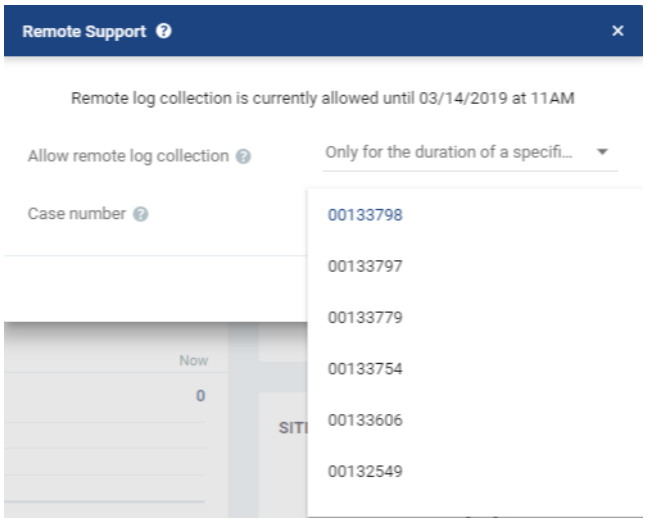

   Remote log collection will be allowed for as long as the case is active or until remote log collection is terminated by selecting the **Never** option.
4. Click **Save**.

Using the Zerto Diagnostics Tool

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, for example VRA logs, as well as logs generated by VMware, for example, vCenter Server logs or host logs. The Zerto 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 203.

To collect logs for Zerto support to use when troubleshooting:

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

   The Zerto Diagnostics menu dialog is displayed.

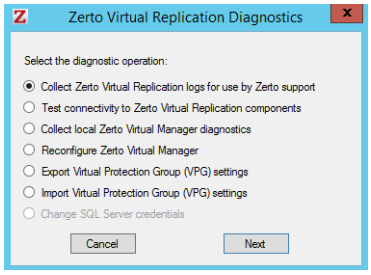

2. Select the Collect the Zerto logs for use by Zerto support option.
3. Click **Next**.

   The Initialize dialog is displayed.

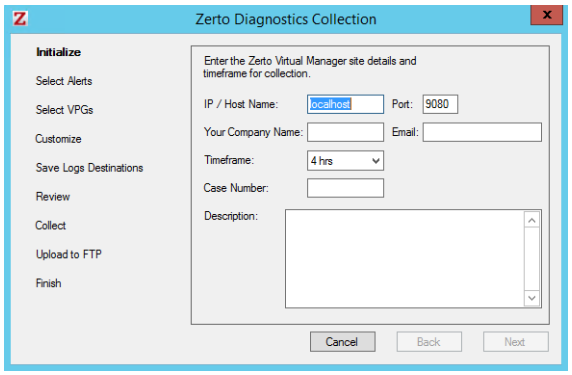

4. Specify the following and click **Next**.

   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 Zerto and if any alerts have been issued, they are displayed in the Select Alerts dialog.

   If there are no alerts, this dialog is skipped.

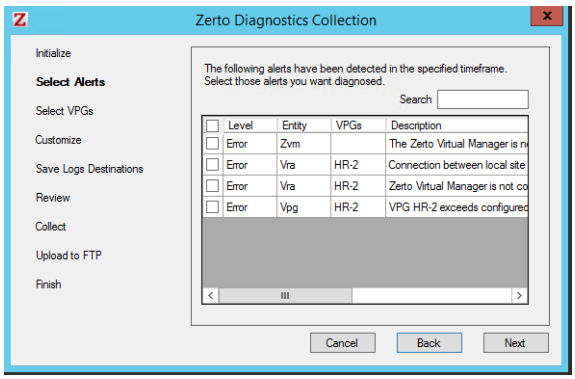

5. Select any alerts that need analyzing from the list and click **Next**.
   The Select VPGs dialog is displayed.

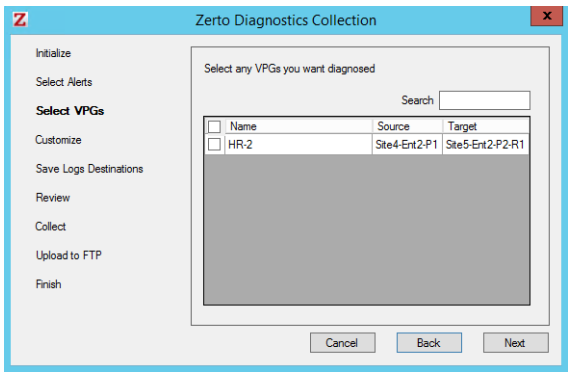

6. Select the VPGs that you want analyzed and click **Next**.
   The Customize dialogs are displayed. With AWS these values are not relevant.
7. Click **Next** until the Save Log Destinations dialog is displayed.

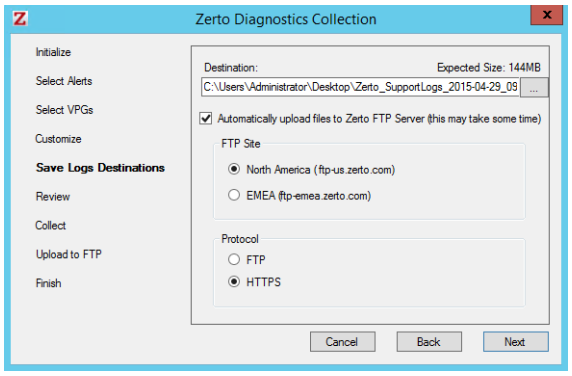

8. Specify destination for the files that you want collected.
   Destination: The name and location where the log collection will be saved.
   Automatically upload files to Zerto FTP Server: When this option is checked, the log collection is automatically uploaded to a specified FTP site.
   If you choose 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.

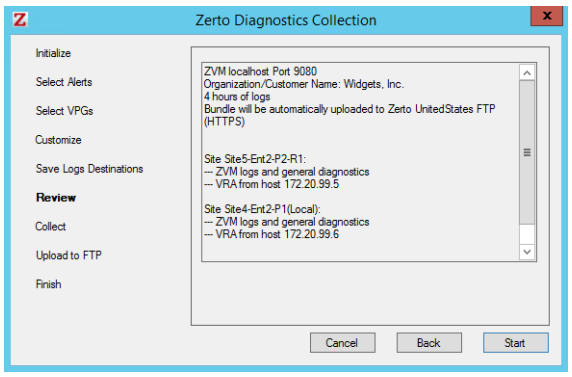

Check that you have specified everything you want to collect and if you want to make changes, click **Back** to change the selection.

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:

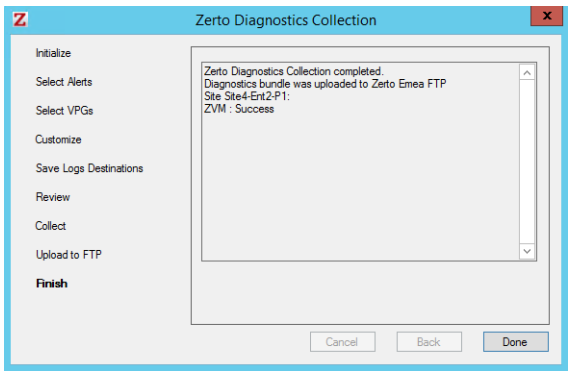

12. Click **Done** to return to the Zerto Virtual Replication Diagnostics menu dialog.
13. Send the log to Zerto support, unless the Automatically upload files to Zerto FTP Server option was specified, in which case it is automatically sent to Zerto.

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 > Zerto Diagnostics.

   The Zerto 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.

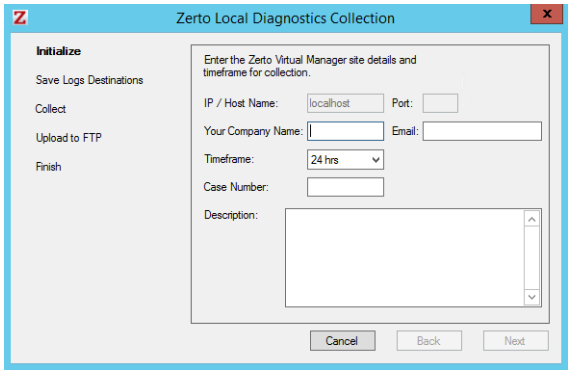

3. Specify the details that you want collected.

   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.

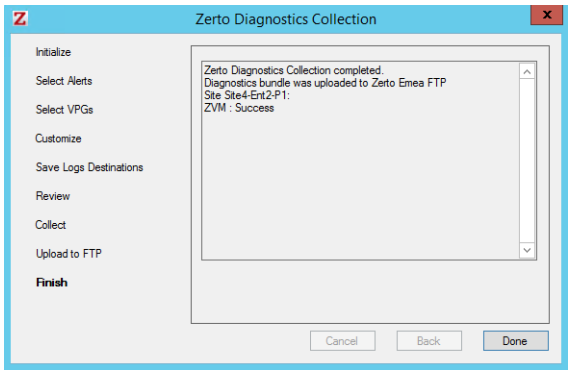

7. Click **Done** to return to the Zerto Diagnostics menu dialog.
8. Send the log to Zerto support, unless the Automatically upload files to Zerto FTP Server option was specified, in which case it is automatically sent to Zerto.

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\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 20: 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 section:

- "About Dialog", on page 206
- "Add Checkpoint Dialog", below
- "Add Site Dialog", on page 208
- "Advanced Journal Settings Dialog", on page 208
- "Advanced VM Replication Settings Dialog", on page 210
- "Advanced VM Settings for Cloud Dialog", on page 210
- "ALERTS", on page 211
- "Boot Order Dialog", on page 212
- "Checkpoints Dialog", on page 213
- "Edit NIC Dialog", on page 214
- "Edit vNIC Dialog", on page 215
- "Edit VM Settings Dialog", on page 216
- "License Dialog", on page 216
- "New Repository Dialog", on page 217
- "Offsite Clone Dialog", on page 219
- "Open Support Case Dialog", on page 220
- "Remote Support Dialog", on page 221
- "Restore VPG - NICs Dialog", on page 222
- "Restore VPG - Volumes Dialog", on page 222
- "Site Settings Dialog", on page 222
- "Stop Failover Test Dialog", on page 226
- "TASKS", on page 227

About Dialog

In the About window, you can do the following:

- View the version of Zerto being run.
- Enable or disable the **Zerto CALLHOME feature**. The Zerto CALLHOME feature enables support notification and analytics for the following purposes:
  - To improve Zerto.
  - To send notifications to the user when a new Zerto version is available, or when new hypervisor versions are supported by Zerto.
- Enable or disable Zerto Virtual Manager to send data to the SaaS platform for monitoring purposes, using the Zerto Mobile App. This action is done by licensed Zerto Virtual Manager users.

When clicking About, the following options appear:

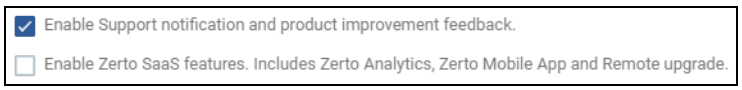

To perform these actions, do the following:

1. In the Zerto User Interface, in the top right of the header, click ☰, and then click **About**.
   The version and build of Zerto installed in the site are displayed.
2. The Zerto CALLHOME feature is selected by default; **Enable Support notification and product improvement feedback**.
   Note: This option is grayed out for Microsoft Azure and AWS.

If the user deselects Enable Support notification and analytics, a warning appears notifying the user that deselecting this option will stop Zerto from sending notifications when new Zerto updates are available, or when new hosts are supported.

3. If you want Zerto to send information to our Online Services and Zerto Mobile App, and enable remote upgrade, select **Enable Zerto SaaS features. Include Zerto Analytics, Zerto Mobile App and Remote upgrade.** This is selected by default.

   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.

   - If the user deselects Enable Online Services and Zerto Mobile, a warning appears notifying the user that deselecting this option will stop Zerto from sending information to Online Services and to the Zerto Mobile Application, rendering these services inoperable for the entire installation.

Add Checkpoint Dialog

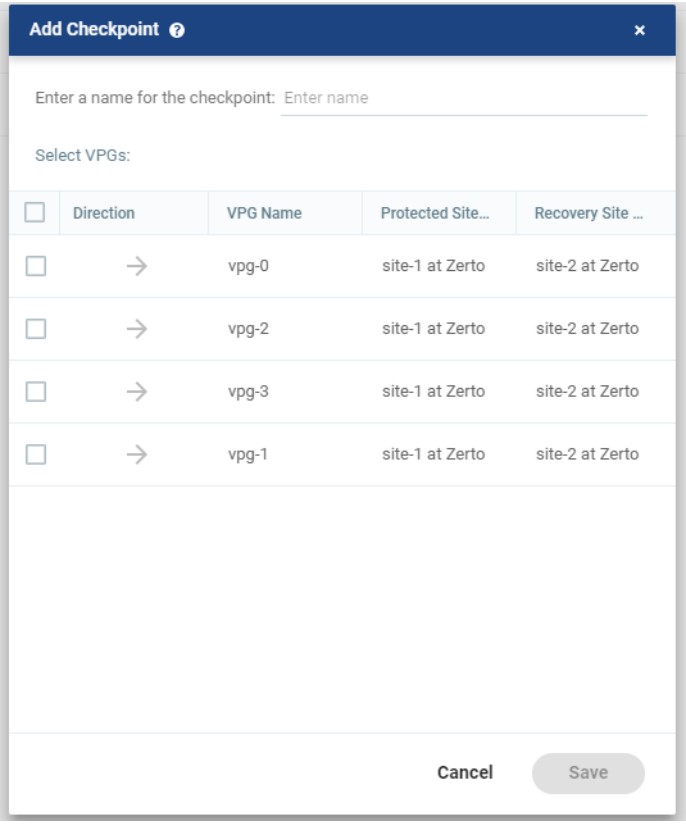

Checkpoints are recorded automatically every few seconds in the journal. These checkpoints ensure crash-consistency and are written to the virtual machine journals by the Zerto Virtual Manager. Each checkpoint has a timestamp set by the Zerto Virtual Manager. In addition to the automatically generated 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.
- **Direction:** 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

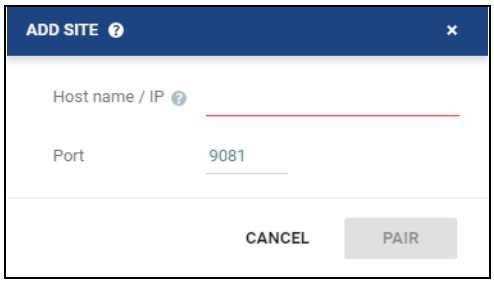

Pair sites

- **Host name/IP:** 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 Journal Settings Dialog

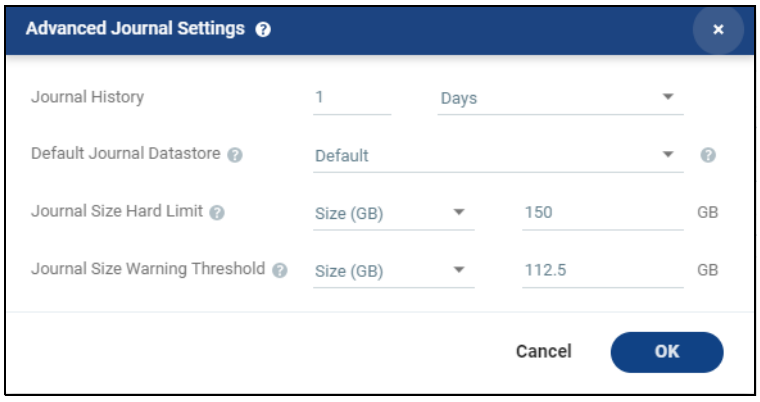

4. Select the journal settings.

SETTING & DESCRIPTION	SELECT...
Journal History	
The time that all write commands are saved in the journal. The longer the information is saved in the journal, the more space is required for each journal in the VPG.	■ Number of **hours** from **1 to 23** ■ Number of **days** from **1 to 30**
Default Journal Storage (Hyper-V), or **Default Journal Datastore** (vSphere)	
The storage/datastore used for the journal data for each virtual machine in the VPG. **Note:** This field is **not** relevant when replicating to a **vCD** recovery site.	■ Select the storage/datastore accessible to the host. When you select a specific journal storage/datastore, the journals for each virtual machine in the VPG are stored in this storage/datastore, regardless of where the recovery storage/datastore is for each virtual machine. All protected virtual machines are recovered to the hosts that can access the specified journal storage/datastore.
Journal Size Hard Limit	

The maximum size that the journal can grow, either as a percentage or a fixed amount.
 	
The journal is always **thin-provisioned**.
 	
Note: The Journal Size Hard Limit applies independently **both** to the Journal History **and** also to the Scratch Journal Volume.
 	
For Example: If the Journal Size Hard Limit is configured to a maximum size of 160 GB limit, then during Failover Test, both the Journal History and the Scratch Journal Volume together can take up to 320 GB. Each one with a maximum size of 160 GB limit.
	■ **Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/ datastore.
	
	If Unlimited is selected, Size and Percentage options are **not** displayed.
	
	■ **Size (GB):** The maximum journal size in GB.
	  ■ The **minimum** journal size, set by Zerto, is **8GB** for Hyper-V and vSphere environments, and **10GB** for Microsoft Azure environments.
	
	■ **Percentage:** The percentage of the virtual machine volume size to which the journal can grow.
	  ■ This value can be configured to more than 100% of the protected VM's volume size.

Journal Size Warning Threshold	
The size of the journal that triggers a warning that the journal is nearing its hard limit.	■ **Unlimited:** The size of the journal is unlimited and it can grow to the size of the recovery storage/ datastore. If Unlimited is selected, Size and Percentage options are **not** displayed. ■ **Size* (GB):** The size in GB that will generate a warning. ■ **Percentage*:** The percentage of the virtual machine volume size that will generate a warning. *****The values of **Size** and **Percentage** must be **less** than the configured **Journal Size Hard Limit** so that the warning will be generated when needed. In addition to the warning threshold, Zerto will issue a message when the free space available for the journal is almost full.

5. Click **OK**.

Advanced VM Replication Settings Dialog

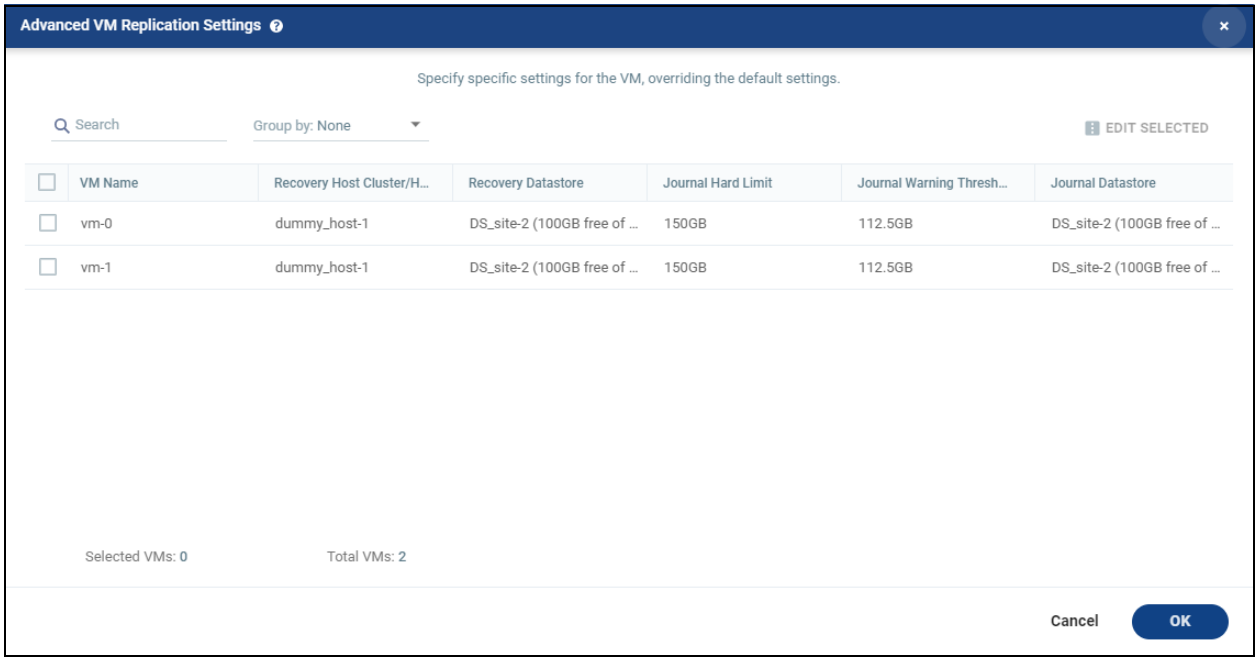

Displays the replication settings for the virtual machines in the VPG. You can choose to edit information in one field by clicking the field and updating the information. You can choose to edit information for several virtual machines at the same time by selecting the virtual machines and clicking **EDIT SELECTED**.

For more details, see "Edit VM Settings Dialog", on page 216.

Advanced VM Settings for Cloud Dialog

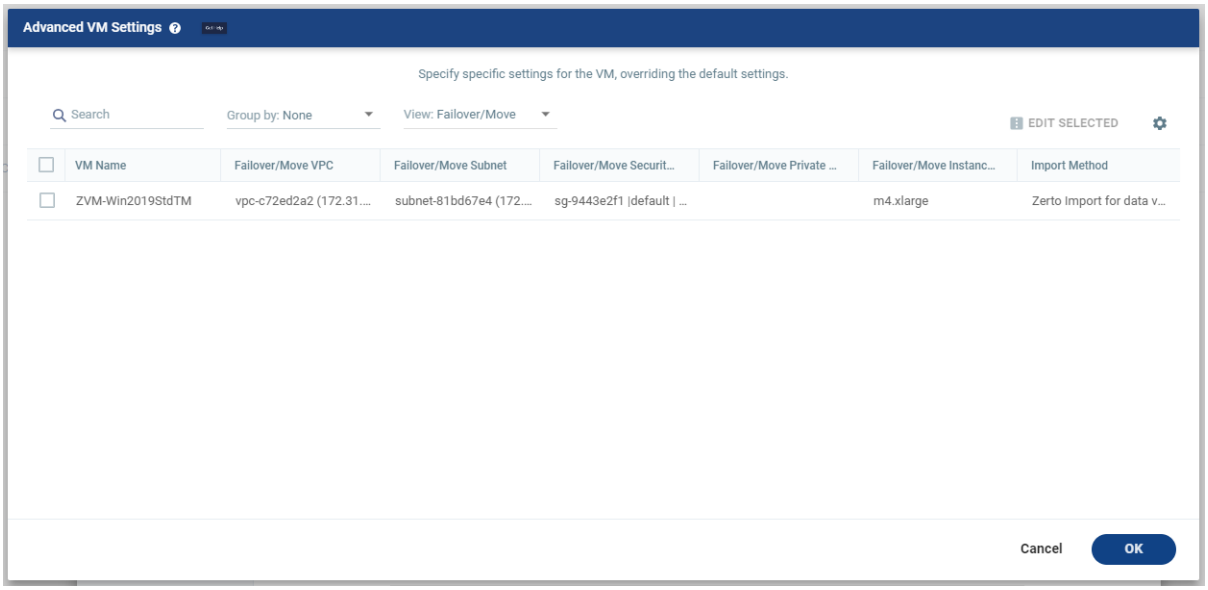

Displays the recovery settings for the virtual machines in the VPG. By default, the recovery settings for failover and move are displayed. You can display the recovery settings for failover tests by selecting **TEST**. You can choose to edit information in one field by clicking the field and updating 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

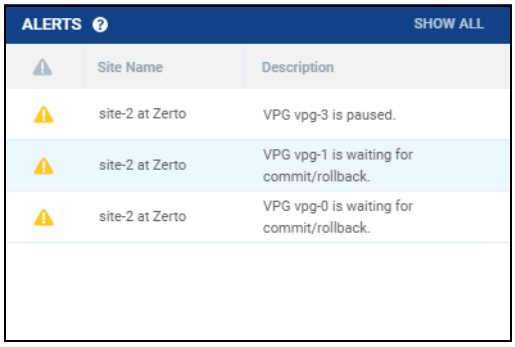

Monitor the recent alerts by clicking the ALERTS area in the status bar at the bottom of the Zerto User Interface. The following information is displayed for the most recent alerts:

- The alert status.
- The site where the alert is issued.
- A description of the alert.

Click **All Alerts** to access MONITORING > ALERTS.

Boot Order Dialog

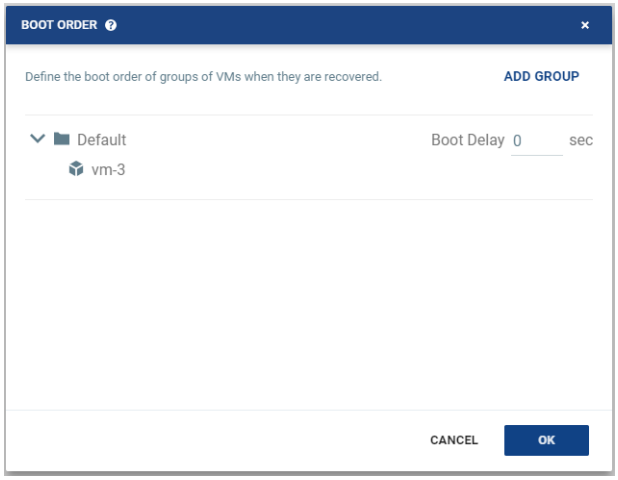

To specify the boot order of virtual machines in a VPG. When machines are started up on recovery, for example, after a move operation, the virtual machines in the VPG are not started up in a particular order. If you want specific virtual machines to start before other machines, you can specify a boot 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:** 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

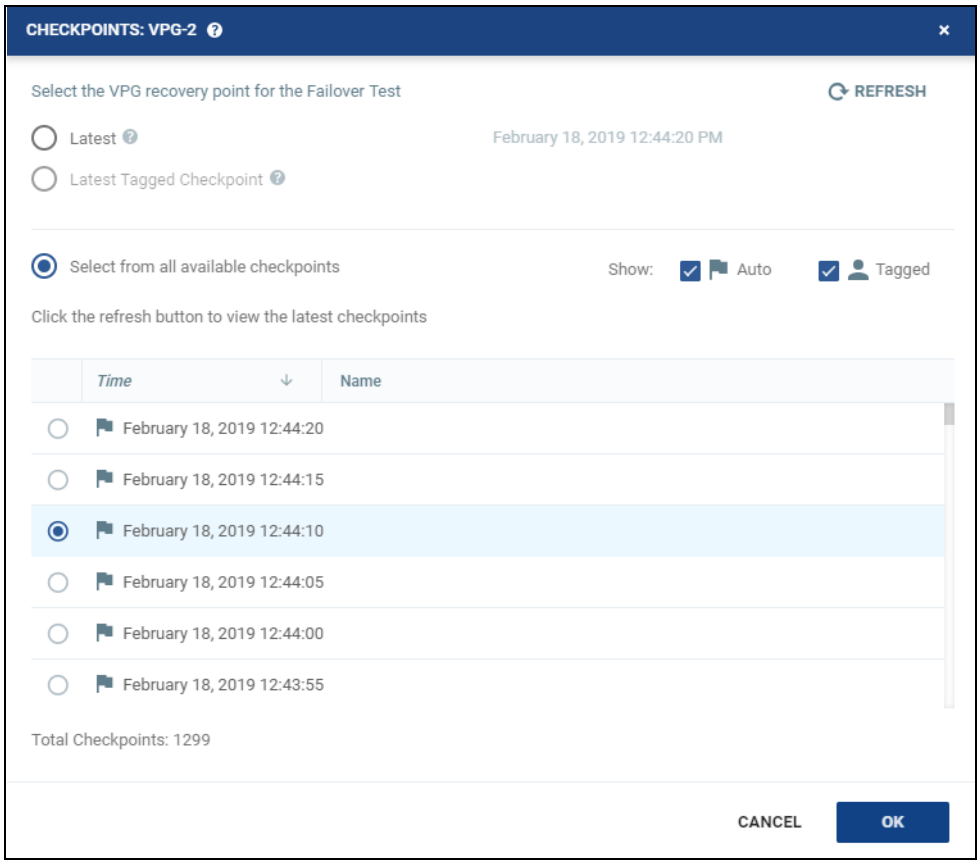

When selecting the point to recover to, remember the following:

- 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 reminder, **Click the refresh button to view the latest checkpoints** 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.

Use the following filters:

- **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.
- **Select from all available checkpoints:** If you do not want to use the latest checkpoint, 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.

Edit NIC Dialog

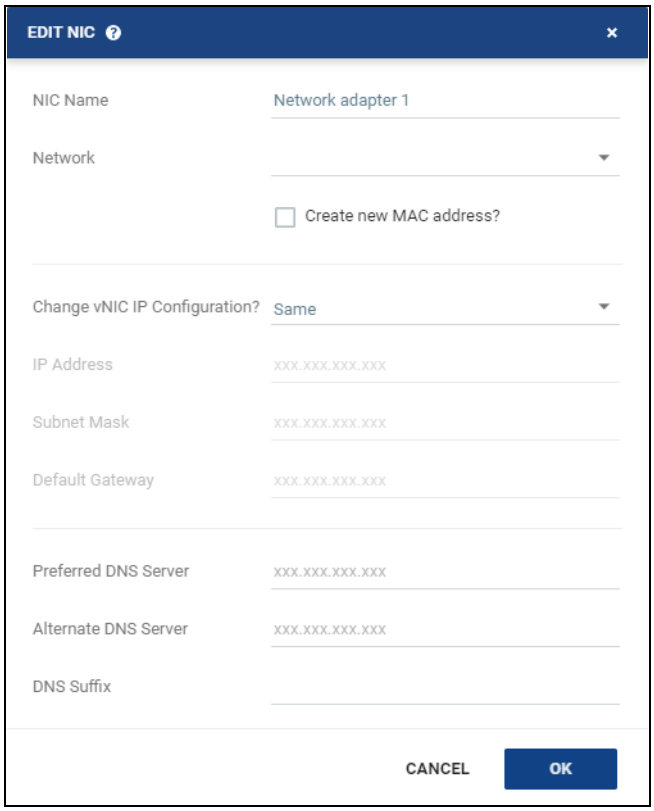

Specify the NIC settings when restoring a retention set to the recovery site.

- **NIC Name:** The name of the selected NIC.
- **Network:** The network to use for the restored virtual machine.
- **Create new MAC address:** The Media Access Control address (MAC address) to use. The default is to use the same MAC address for the restored virtual machine that was used in the protected site. Check the box to create a new MAC address on the restore site.
- **Change vNIC IP Configuration:** Whether or not to keep the default virtual NIC (vNIC) IP configuration. The vNIC IP is changed after the restore has completed when VMware Tools are installed.
  - If **Static** is selected, the IP address, subnet mask, and default gateway must be set.
  - If **DHCP** is selected, the IP configuration and DNS server configurations are assigned automatically, to match the protected virtual machine.
- **IP Address:** The IP for the restored virtual machine. This can be the same IP as the original protected virtual machine.
- **Subnet Mask:** The subnet mask for the network. The default value is **255.255.255.0**.
- **Default Gateway:** The default mask for the network.
- **Preferred DNS Server:** The IP address of the primary DNS server to handle Internet protocol mapping.
- **Alternate DNS Server:** The IP address of the alternate DNS server.
- **DNS Suffix:** The DNS name excluding the host.

Edit vNIC Dialog

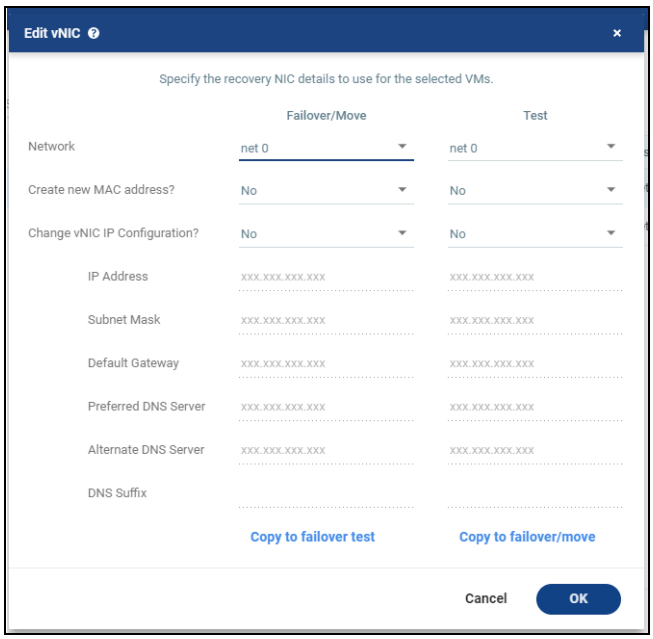

Specify the network details to use for the recovered virtual machines after a failover or move operation, in the Failover/Move column, and for the recovered virtual machines when testing replication, in the Test column.

6. In each column, **specify** the following:
   - **Network:** The network to use for this virtual machine.
   - **Create new MAC address?:** Whether the Media Access Control address (MAC address) used on the protected site should be replicated on the recovery site. The default is to use the same MAC address on both sites. Note that if you check this option, to create a new MAC address, and the current IP address is not specified, the protected virtual machine static IP address might not be used for the recovered virtual machine.
   - **Change vNIC IP Configuration?**: Whether or not to keep the default virtual NIC (vNIC) IP configuration. The vNIC IP is only changed after recovery for virtual machines with VMware Tools running.
     See the Zerto Interoperability Matrix for the list of operating systems for which Zerto supports Re-IPing.
     - To change the vNIC IP, in the Failover/Move or Test column, select **Yes**. If you select to use a static IP connection, set the **IP address**, **subnet mask**, and **default gateway**.
     - Optionally, change the preferred and **alternate DNS server IPs** and the DNS suffix.
     - If you leave the DNS server and suffix entries empty, or select to use DHCP, the IP configuration and DNS server configurations are assigned automatically, to match the protected virtual machine. You can change the DNS suffix.
     - If the virtual machine has multiple NICs but is configured to only have a single default gateway, fill in a **0** for each octet in the Default gateway field for the NICs with no default gateway.

   During a failover, move, or test failover, if the recovered virtual machine is assigned a different IP than the original IP, then after the virtual machine has started, it is automatically rebooted so that it starts up with the correct IP. If the same network is used for both production and test failovers, Zerto recommends changing the IP address for the virtual machines started for the test, so that there is no IP clash between the test machines and the production machines.
   - **Copy to failover test:** Select this to copy the settings in the Failover/Move column to the **Test** column.
   - **Copy to failover/move:** Select this to copy the settings in the Test column to the **Failover/Move** column.

Edit VM Settings Dialog

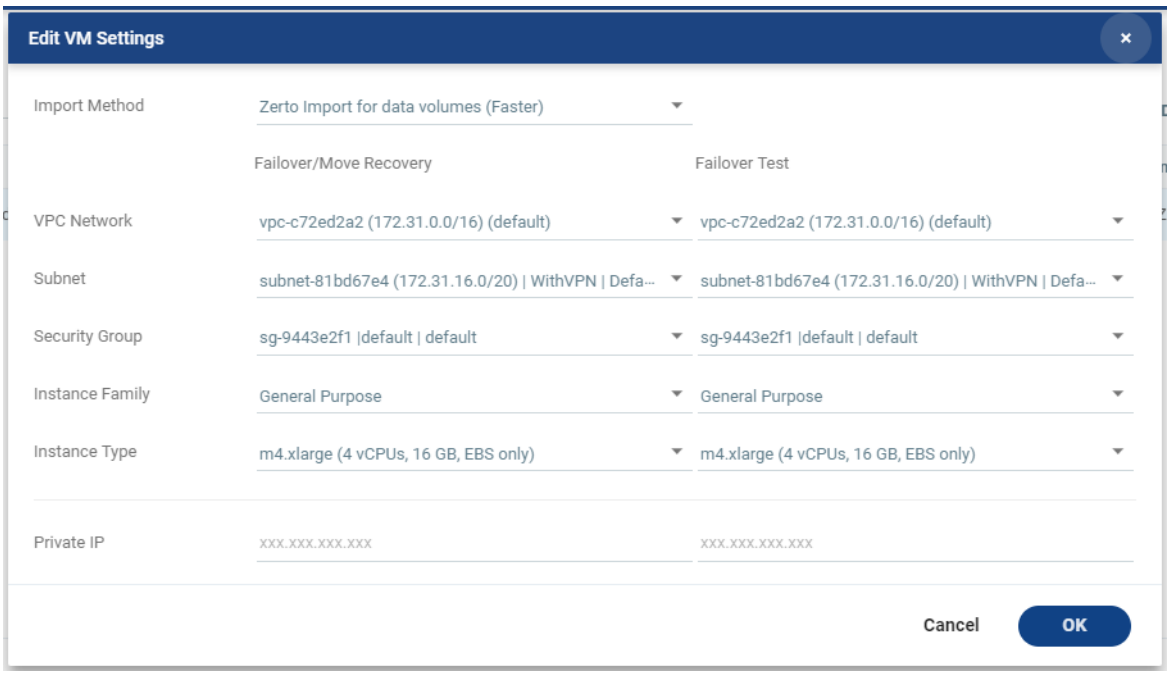

Edit the network settings for one or more virtual machines in a VPG that will be recovered to AWS. There are recovery settings for failovers and moves, and for failover tests.

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.

License Dialog

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 by clicking ☰ in the top right of the header and selecting **License**.

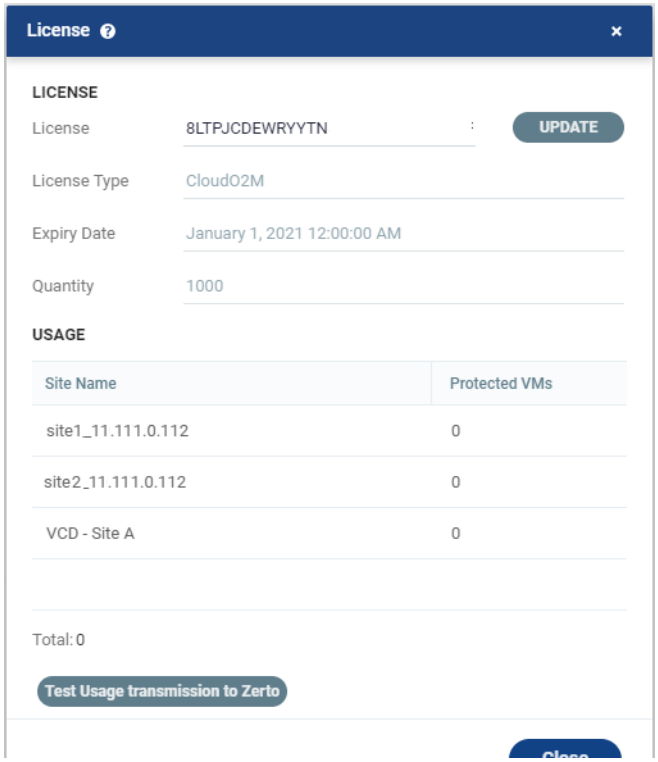

The Zerto license includes the following details:

- **License:** The license key itself.
- **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.
- **Usage:** The sites using the license and the 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 with this key.

New Repository Dialog

Define Repositories on the recovery site where retention sets can be stored. The Repositories are connected over the network, leveraging SMB(CIFS) or NFS protocols, on top of a simple network share or PBBA.

The Repository is defined as part of the retention policy configuration.

To create a Repository for Long Term Retention:

1. From the recovery site, in the Zerto User Interface, click **Setup > REPOSITORIES > NEW REPOSITORY**.

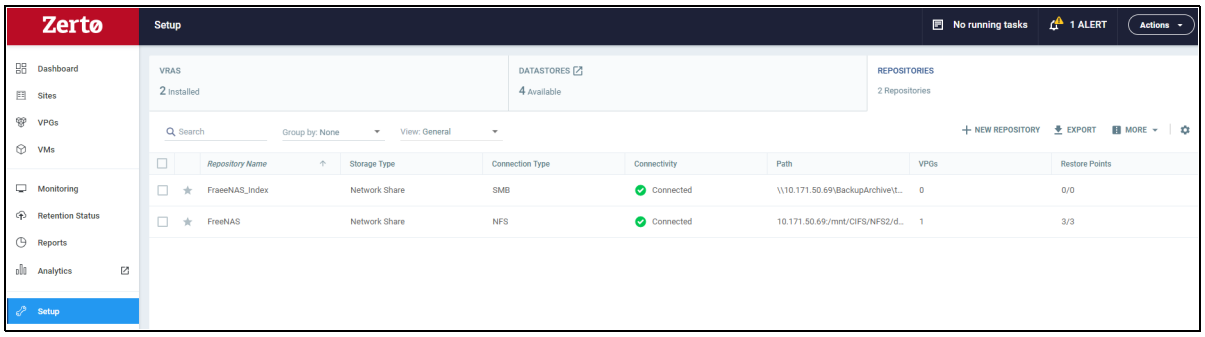

   The New Repository window is displayed.

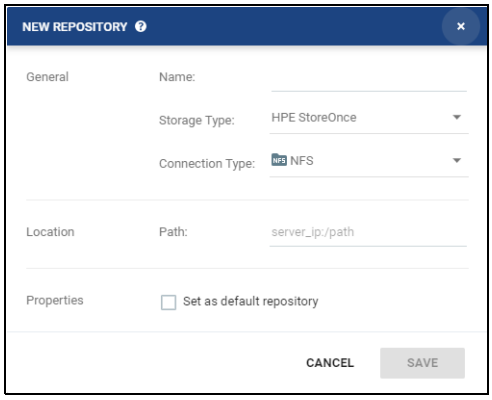

2. In the General area, specify the following:
   - **Name:** A unique name for the Repository.
   - **Storage Type:** The type of network share, selected PBBAs or other de-duplicated appliances. See the Interoperability Matrix for supported versions.
   - **Connection Type:** The type of Repository:

■ **Network Share - NFS**	The network share is NFS based
■ **Network Share - SMB**	The network share is SMB based. Username and password must be provided to access the share.
	Username: The username can be entered using either of the following formats: ■ **username** ■ **domain\username**
	Password: Password to access the Network Share.

3. In the Location area, define the **Path**. This is the path where the Repository will reside.

   The path must be accessible from the DSS, so if the Repository is on a different domain, the domain must be included in the path.

4. In the Properties area, you can check the **Set as default repository** checkbox. This will use the Repository as the default when defining the retention policy in a VPG.
5. Click **SAVE**. The Repository is created.

Offsite Clone Dialog

On-premise environments:

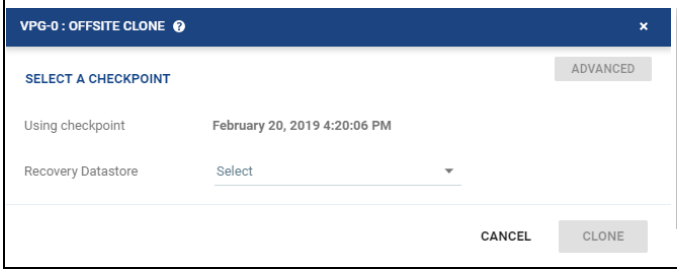

Cloud environments:

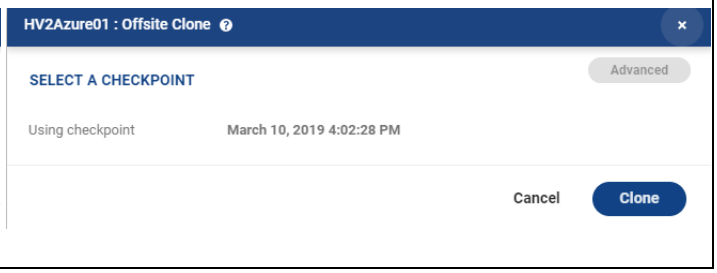

To create a clone of each virtual machine in a VPG on the recovery site in the production network. The clone is a copy of the protected virtual machines on the recovery site, while the virtual machines on the protected site remain protected and live.

- **SELECT A CHECKPOINT:** Opens the Checkpoints Dialog to select the checkpoint to use to make the clone.

  The Using checkpoint field is updated with your selection.
- (On-Premise) **Recovery Datastore:** Select the datastore to use for the recovery virtual machines.
- **Advanced:** Select specific VMs to clone.

For more details, see Checkpoints Dialog.

Open Support Case Dialog

Support cases can be opened directly in the Zerto User Interface.

Creating a support case in the Zerto User Interface simplifies the submission process since much of the information that is required when entering a case using the Zerto Support Portal, such as the version and build numbers, is automatically added to the case when it is submitted via the Zerto User Interface.

In addition, when the case is submitted, a snapshot of the current environment is also attached to the case. 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 case quickly and, whenever possible, without the need to request more information from you.

Note: The clocks on the machines where Zerto 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 case can fail.

To open a support case:

1. In the Zerto User Interface, click ☰ in the top right of the header and select Open a Case.
   The Open Support Case window for the site opens.

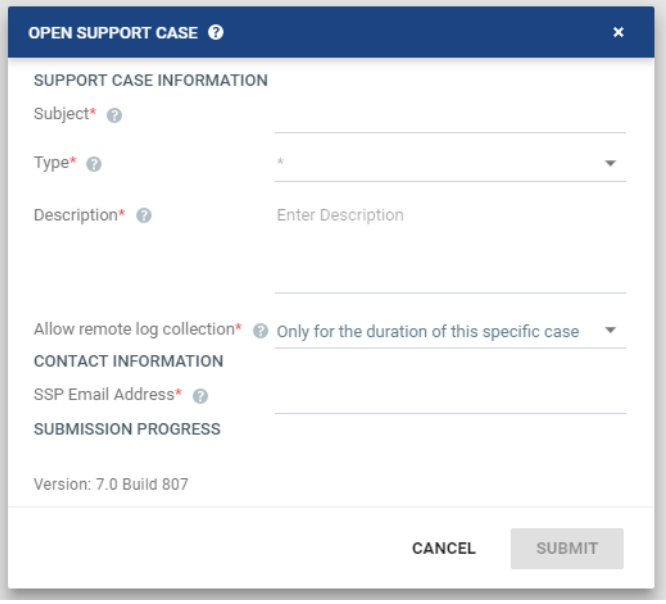

2. Specify the case details:
   - **Subject:** The subject of the support case.
   - **Type:** The type of case being opened. Available options are:
     - Problem
     - Question
   - **Description:** A description of the problem or question in addition to the information supplied in the subject.
   - **Allow remote log collection:** How many logs is Zerto allowed to collect. Available options are:
     - Only for this case
     - For the next 30 days
     - Never
   - **SSP Email Address:** A valid email address registered with Zerto, with permission to open cases.
3. Click **SUBMIT**.

The case is processed and its progress is displayed. If the email address is not valid, the case is rejected. Once the case submission starts, it cannot be canceled.

Remote Support Dialog

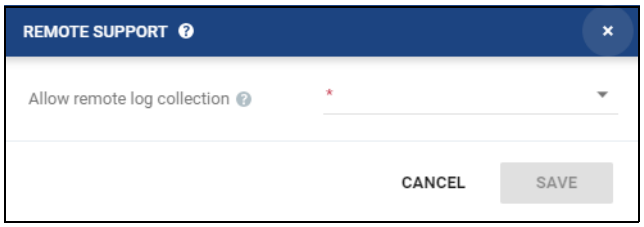

Remote Log Collection allows customers to authorize Zerto support engineers to collect logs from their environment. By using remote log collection customers can avoid having to use the Diagnostic Tool on their ZVM server in order to collect logs for analysis, a potentially complex and time-consuming procedure.

- **Never:** Remote log collection is not allowed (default). If remote log collection is currently allowed, the remote connection 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 the duration of a specific case**: You will be prompted to select the Case number from the drop-down list. The list contains all the active cases opened under the account that the Zerto Virtual Manager is registered to. 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.

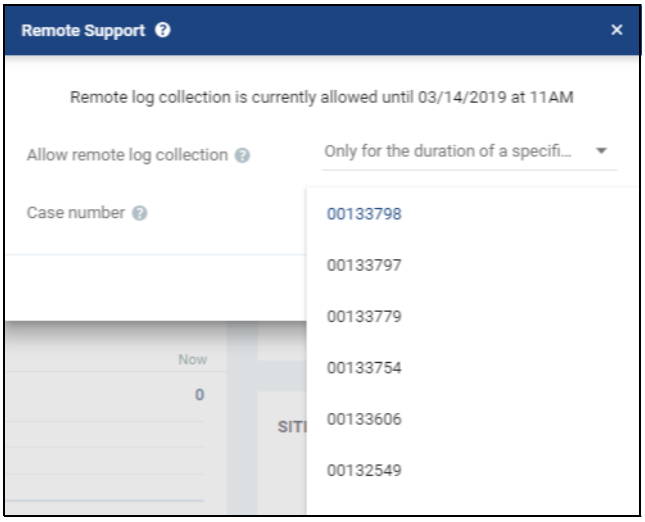

Restore VPG - NICs Dialog

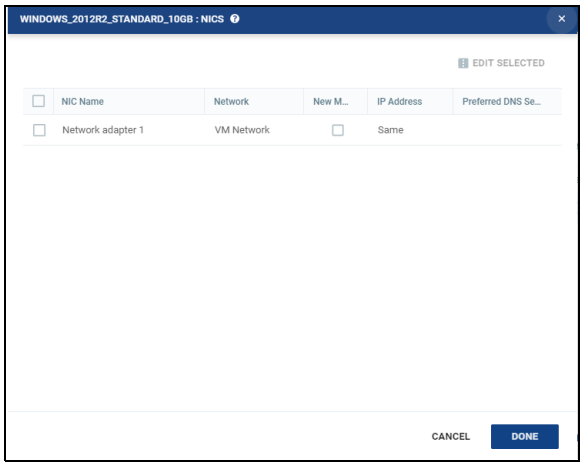

When restoring a retention set to the recovery site, this dialog shows the NIC settings for the virtual machines in the VPG. You can choose to edit information in one field by clicking the field and updating the information. You can choose to edit the NIC settings for several virtual machines at the same time by selecting the NICs and clicking **EDIT SELECTED**.

Restore VPG - Volumes Dialog

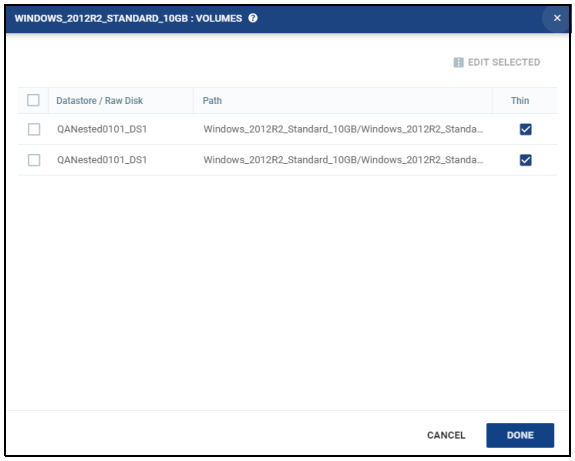

When restoring a retention set to the recovery site, this window shows the datastores for a selected virtual machine in the VPG.

- You can edit information in one field by clicking the field and updating the information.
- You can edit several datastore settings at the same time by selecting the datastores and clicking **EDIT SELECTED**.

Site Settings Dialog

Contains site-wide settings:

- "Site Information Dialog", below
- "Policies Dialog", on page 225
- "Email Settings Dialog", on page 226

■ "Reports Dialog", on page 226

Site Information Dialog

To update information about the local site:

1. In the Zerto User Interface (top right), click **SETTING** ( ) and select **Site Settings**.
   The Site Settings dialog is displayed.
   - In **vSphere** and **Hyper-V** sites, the following fields are displayed:

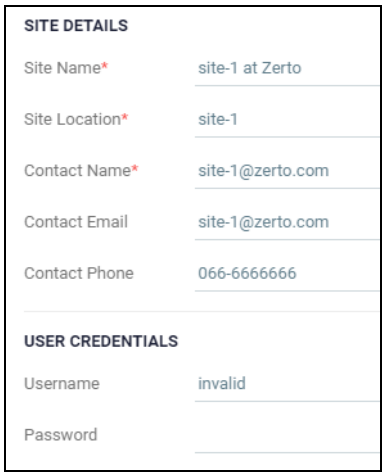

   - In **AWS** sites, the following fields are displayed:

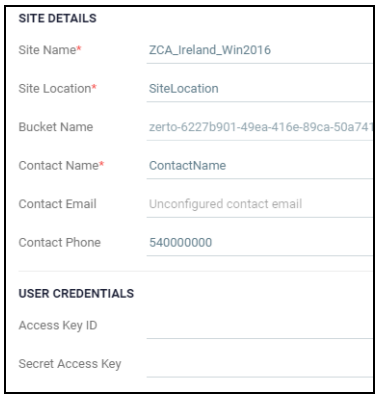

   - In **Microsoft Azure** sites, the following fields are displayed:

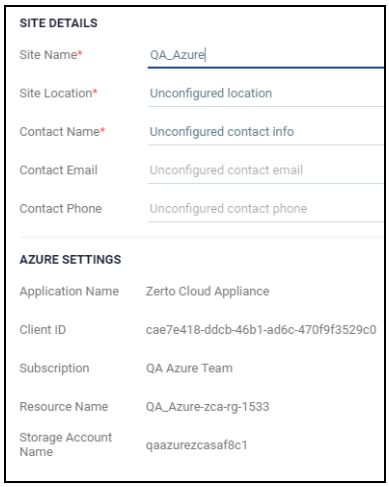

2. Define general information about the site.
   - (Mandatory) **Site Name:** The name used to identify the site.
   - (Mandatory) **Site Location:** Information such as the address of the site or a significant name with which to identify the site location.
   - (AWS environments only) **Bucket Name:** The name of the bucket that was created when Zerto was installed. This cannot be changed.
   - (Mandatory) **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. (On premise environments) To change the **User Credentials** to access the vCenter or Hyper-V SCVMM server from the Zerto Virtual Manager:
   - **User Name**: The administrator name used to access the vCenter or SCVMM server. The name can be entered using either of the following formats:
     - **username**
     - **domain\username**
   - **Password:** The password used to access the vCenter or SCVMM server for the given user name. To ensure security, after saving the settings, the password field is cleared.
4. (Azure environments only) **Azure Settings**:
   - **Application Name:** The name used to access Azure.
   - **Client ID**: A unique identifier that is associated with the access name.
   - **Subscription:** The subscription associated with the user.
   - **Resource Name:** The name of the resource the user created.
   - **Storage Account Name:** The name of the storage account created or selected for this site during installation.
5. Click **SAVE**.

Policies Dialog

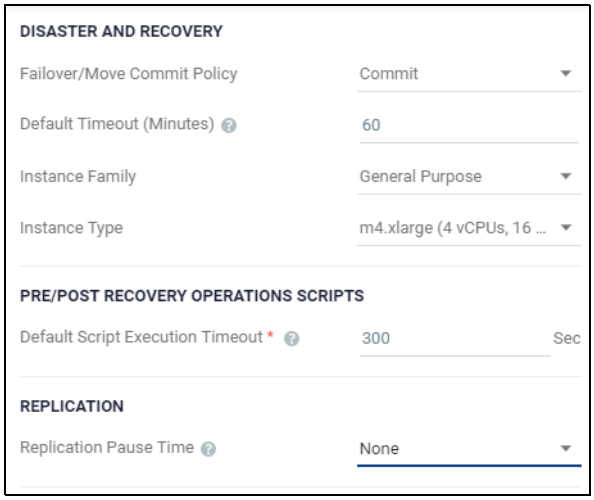

- **Failover/Move Commit Policy:** The commit policy to use during a failover or move operation. The value set here is the default for all failover or move operations from this point on but can be changed when defining a failover or move operation. The following options are available:
  - **None:** The 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.
- **Default Script Execution Timeout:** The length of time in seconds after which a script times out.
- **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.
- **For incoming replication, copy the BIOS UUID of the protected VM to the recovered VM:** Select this to preserve the **BIOS UUID** of the **protected** VM after recovery operations, in the **recovery** ZVM Site Settings window.

  Note:
  - Preserving of the BIOS UUID is **not supported** in Clone and Self-Replication recovery operations.
  - Preserving of the BIOS UUID is **not supported** in Public Cloud.
  - Cross replication is **not** supported.
-

Email Settings Dialog

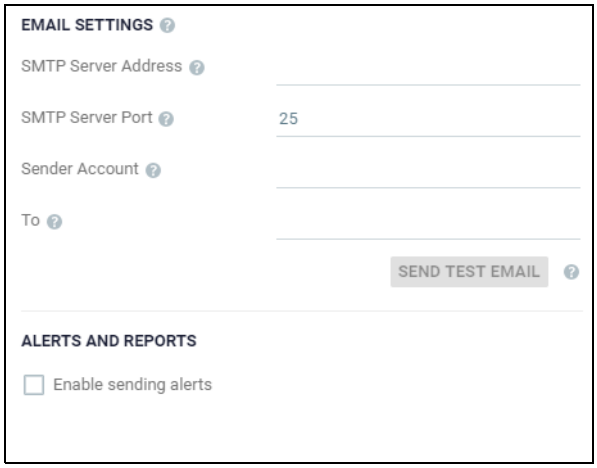

Define an email address to receive Zerto alerts and retention reports.

- **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 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.

Reports Dialog

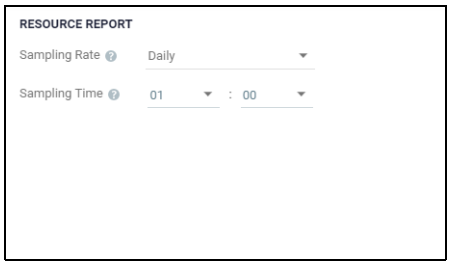

Configures the Resource Report.

- **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.

Stop Failover Test Dialog

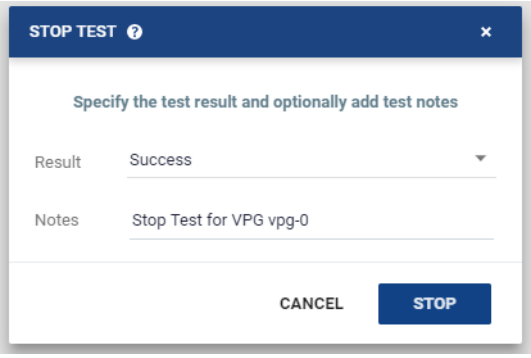

Stops the testing of the selected VPG.

- **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.
- Clicking the **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

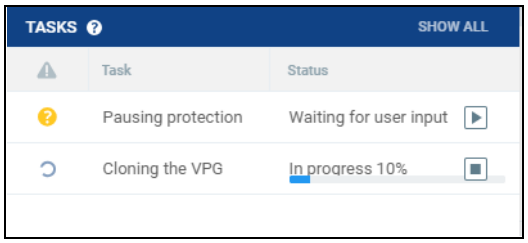

Monitor the recent tasks by clicking the TASKS area in the status bar at the bottom of the Zerto User Interface. The following information is displayed for the most recent tasks:

- 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 **SHOW ALL** to access MONITORING > TASKS.

CHAPTER 21: 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 Long Term Retention.
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 starts to maintain a smart bitmap in memory to track and record changed storage areas. Since the bitmap is kept in memory, Zerto 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 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 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).
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. It is **not possible** to perform a move during a delta 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.
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 DR includes the ability to configure both disaster recovery and Long Term Retention 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 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 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.
Long Term Retention	Providing Zerto customers a comprehensive Data Protection solution - preserve their virtualized environment production data, for a long period - from months to several years. The data is copied from the Recovery site, to a local or remotely attached storage repository, without any impact on production data or performance, via a single management platform. Data Retention is enforced by built-in scheduling and retention mechanisms. Data can later be restored easily, and with minimal RTO.
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.
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 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 takes ownership of the preseeded disk, moving it from its source folder to the folder used by the VRA.
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 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 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 retention process 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 re 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.
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 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 Windows service that manages File Level Recovery operations within Zerto. These repositories can be local or on a shared network.
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.

A

access key ID 133
alerts 211
  generated by a failover 174
  generated by a move 161
  in Dashboard 104
  in single VPG tab 108
application-consistent checkpoints 120
architecture 9
AWS
  access key ID 133
  editing VPG settings 216
  import process limit 149, 157, 168
  re-IP process 102
  secret access key 133

B

bandwidth
  freeing up 118
  in resources report 194
bitmap sync 127, 169, 228
boot order 149, 157, 169
  configuring 41, 60, 77, 212
branding the Recovery report 193

C

change rate
  estimating 51, 69, 87
checkpoint 119, 120, 171, 177, 207, 213
  add 37, 121
  adding manually 121
  application-consistent 120
  choosing 168, 177
  purpose 35
  renaming after testing failover 154
clock synchronization with NTP 147, 220
cloning ..... 176–??
  definition 176
commit policy 168
  failover and move 143
  for failover and move 225
  site setting 225
  site settings 143
compression
  to minimize bandwidth 24
connecting sites, see pairing
connectivity, checking 132

D

Dashboard 32, 103
  alerts, events, and running tasks 104, 108
  performance graphs 104
  stopping a failover test 153
default gateway, configuring 57, 93, 215
default instance type
  m3.xlarge 152, 158, 162, 168, 174
delta sync 128
  for volumes 128
diagnostics utility 132, 133
  changing IP address of ZCA 133
  checking connectivity 132
  collecting logs 201, 203
  reconfiguring access to SQL Server DB 134
  troubleshooting 199
disaster recovery
  during a test 175
  initiating 164, 169, 187
  types 37
disk
  RDM 51, 52, 69, 70, 87, 88
disk size
  estimating 51, 69, 87

E

EBS disks 11, 37, 152, 161, 178
EBS disks (AWS) 27, 28
EC2 ..... 37, 149, 152, 157, 162, 168, 174, 178
  creating new instances 152
EC2 (AWS) 27, 28
ec2-import-instance
  increasing its value 149, 157, 168
email settings 144
  configure 144
environment variable 122
  ZertoForce 122
  ZertoOperation 122
  ZertoVCenterIP 122
  ZertoVCenterPort 122
  ZertoVPGName 122
events
  in Dashboard 104
  in single VPG tab 104, 108
export to CSV
  VPG details 107

F

failback 138, 186
  after failover 187
  enable after failover 165, 188
failover 168–175, 186
  commit policy 143, 225
  description 168
  during a test 175
  failback 165, 172, 188
  generating alerts 174
  import process 149, 157, 168
  initiating 164, 169, 187
  initiating during a test 175
  overview 138
  process 186
  stopping a test 153
  testing 149
  topology 165, 173, 189
failover test
  in a sandbox 155
  overview 137
file recovery 179
folder recovery 179
force delete 126
  VPG 120
full synchronization 128
  for volumes 130

G

geometry, non-default 52, 70, 88
glossary 228–235

H

HTTPS 30, 135

I

import process to AWS 149, 157, 168
initial synchronization 34, 35, 129
instance family 225
  choosing 143
instance type 225
  choosing 143, 152, 158, 168

J

journal 119
  adding a checkpoint 121
  description 35
  history in VPG tab 108

K

Keep Source VMs 157, 161

L

license 145, 216
  updating 146
logs
  collecting 200–204
  collecting when Zerto Virtual Manager is down 203
logs, collecting ..... 199–??
logs, understanding ..... 205–??

M

m3.xlarge
  default instance type 152, 158, 162, 168, 174
MAC address 74
Microsoft SQL Server DB
  reconfiguring 134
migration, see move
monitor
  alerts, events, and running tasks 104, 108
monitoring
  one VPG 107
  virtual machines 112
  VPGs tab 105
  Zerto Virtual Replication 103
move 157–162, 166, 189
  commit policy 143, 225
  definition 157
  description 163
  generating alerts 161
  initiating 159
  overview 138
  using a scratch volume 166, 189

N

needs configuration 129
  moving without reverse protection 138
  troubleshooting 198
network
  for failing over or moving 55, 91
  for testing 55, 91
NIC
  defining network details 222
NTP clock synchronization 147, 220

O

offsite backup
  flow 38
  status in single VPG tab 108
  status in VPGs tab 106, 107
offsite backup repository
  creating 218

P

pairing 34, 134, 135
pause protection 118
performance graphs 104
  for a single VPG 107
policies
  configuring 143
PowerShell cmdlet
  Set-Checkpoint 120
preseed 52, 70, 88
preseeding 119
promotion of data 166, 190
protection
  pause 118
  resume 118
provisioned storage 106, 113

R

raw disk (RDM) 51, 69, 87
recovery 168, 186
  during a test 175
  initiating 164, 169, 187
  to Hyper-V 76
  types 37
recovery flows 37
Recovery report 192
  branding 193
recovery site, pairing 135
registration 145
re-IP 57, 75, 93, 215
  in AWS 102
  vNIC configuration 57, 93, 215
Replication Pause Time 119, 144
reports ..... 192–??
  Recovery 192
  recovery operations (test, failover, move) 192
  Resources 193
  Resources report 193
  VPG Performance 197
Resources report 193
  configuring 226
  generating with REST API 193
  output 194
retention policy 11
reverse protection, see failback
rollback
  failover 169, 174
  move 161
  setting for failover or move 143
running tasks, see tasks

S

S3 ..... 11, 35, 37, 161
sandbox 155
scratch volume 166, 189
scripts
  examples 124
  running 121
  ZertoForce environment variable 122
  ZertoOperation environment variable 122
  ZertoVCenterIP environment variable 122
  ZertoVCenterPort environment variable 122
  ZertoVPGName environment variable 122
secret access key 133
security certificate
  adding 31
Set-Checkpoint cmdlet 120
settings
  importing VPG 125
settings, for a site ..... 140–??
shadow VRA 233
signature matching, WAN optimization 24
site details
  monitoring 103
site settings
  commit policy 143, 225
  email 144
  recovery policies 143
sizing
  volumes 51, 69, 87
SLA information 35
SSL certificate, replacing 135
status
  VPG 125, 126
storage
  for replicated data 50, 68, 84, 86
  provisioned 106, 113
  sizing 50, 68, 86
storage profile 195
summary tab 103
synchronization 34
  bitmap 127
  delta 50, 68, 86
  delta sync 128
  delta sync for volumes 128
  forcing 119
  full 128, 130
  initial 34, 35, 129
  length of time 35
  status 129
  taking a long time 199
synchronization triggers
  VPG 125, 130

T

tasks 111
in Dashboard 104
in single VPG tab 108
test failover 149
description 149
overview 137
stopping 153
throttling 24
topology
for failover 165, 173, 189
for move 160
for testing failover 152
triggers
synchronization 125
VPG synchronization 130
troubleshooting ..... 198–??
collecting logs 199
collecting logs when Zerto Virtual Manager is down 203
Needs Configuration 198
using the Diagnostics utility 199
VPG syncing 199
VRA problems 199
Zerto Virtual Manager service 198

V

vCD guest customizing 73
Virtual Backup Appliance, see VBA
virtual machine
monitoring 112
virtual protection group, see VPG
Virtual Replication Appliance, see VRA
VM in Several VPGs 11, 117, 157
vNIC
configuring 74
vNIC configuration
re-IP 57, 93, 215
volume
estimating size 51, 69, 87
full synchronization 130
VPG 37
add a virtual machine 117
configuring 34
creating to vCenter Server 57
definition 34, 37
deleting 119
editing 216
force delete 126
importing settings 125
modifying 116
monitoring 105
monitoring one 107
pausing protection 118
saving details to a file 107
synchronization triggers 125, 130

synchronizing 119
testing failover 149
topology 110
waiting to be removed 120
VPG Performance report 197
VPG status
VPG waiting to be removed 120, 126
VPG statuses 125, 126
VPGs tab
monitoring 105
VRA 11, 37, 233
troubleshooting 199

W

WAN
signature matching 24
WAN bandwidth, freeing up 118
WAN optimization 24
Windows service
Zerto Virtual Manager 198

Z

ZCA, see Zerto Cloud Appliance
Zerto Cloud Appliance 30
Zerto User Interface
accessing 30
configuration 32
customizing 33
description 32
Zerto Virtual Manager 9
checking component connectivity 132
managing 132
reconfiguring 133
Windows service 198
Zerto Virtual Replication
architecture 9
benefits 22
description 7
logs 205
monitoring 103
ZertoForce script 122
ZertoOperation script 122
ZertoVCenterIP script 122
ZertoVCenterPort script 122
ZertoVPGName script 122
ZORG
preseeding 53, 71, 89
ZVM, see Zerto Virtual Manager

Zerto helps customers accelerate IT transformation by eliminating the risk and complexity of modernization and cloud adoption. By replacing multiple legacy solutions with a single IT Resilience Platform™, Zerto is changing the way disaster recovery, backup and cloud are managed. At enterprise scale, Zerto's software platform delivers continuous availability for an always-on customer experience while simplifying workload mobility to protect, recover and move applications freely across hybrid and multi-clouds. Zerto is trusted by over 6,000 customers globally and is powering resiliency offerings for Microsoft Azure, IBM Cloud, AWS, SunGard AS and more than 350 cloud services providers.

Learn more at Zerto.com

For assistance using Zerto software, contact: @Zerto Support.

Copyright © 2019, Zerto Ltd. All rights reserved.