

Dell EMC RecoverPoint for Virtual Machines

Version 5.1

Administrator's Guide

P/N 302-003-966

REV 01

Copyright © -2017 Dell Inc. or its subsidiaries. All rights reserved.

Published July 2017

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Figures		5
Tables		7
Preface		9
Chapter 1	Getting started	11
	RecoverPoint concepts.....	12
	Activating entitlements.....	12
	Licensing, support, and registration.....	14
	Exporting the post-deployment form to a CSV file.....	16
	Registering RecoverPoint by email or phone.....	16
	Registering datastores.....	17
	Registering ESX clusters.....	17
	Registering vCenters.....	18
Chapter 2	Protecting your data	21
	Overview.....	22
	Protecting virtual machines.....	22
	Stopping protection.....	25
Chapter 3	Managing and monitoring the system	27
	Monitoring the system.....	28
	Reviewing recovery activities.....	28
	Monitoring consistency group replication.....	29
	Consistency group transfer states.....	29
	Managing the system.....	29
	Licenses.....	29
	Support settings.....	30
	RecoverPoint for VMs registration.....	30
	Registering an external host.....	30
	Policy templates.....	31
	Managing group sets.....	31
	Creating a group set.....	31
	Editing an existing group set.....	32
	Enabling a group set.....	32
	Removing a group set.....	32
	Managing consistency groups.....	33
	Enabling or disabling a consistency group.....	33
	Editing group or link policies.....	33
	Creating bookmarks.....	35
	Adding a copy.....	38
	Managing copies.....	39
	Editing copy policies.....	39
	Configuring copy VM failover networks.....	40

	Managing virtual machines and VMDKs.....	41
	Orchestration.....	41
	Automation.....	47
Chapter 4	Testing and recovery	51
	Testing and recovery.....	52
	Testing a copy.....	52
	Failing over.....	54
	Recovering production.....	56
Chapter 5	Troubleshooting	59
	Finding the vRPA cluster management IP.....	60
	Identifying a RecoverPoint for VMs system.....	60
	Recovering from a cluster disaster.....	60
	Detecting bottlenecks.....	60
	Load balancing.....	62
	System alerts.....	63
	Collecting system information.....	63
	Collecting RecoverPoint for VMs splitter logs.....	64
	Manual copy VM network configuration (RE-IP).....	64
	Copy VM network configuration guidelines.....	67
	Changing the network adapter configuration of a protected virtual machine	72

FIGURES

1	RecoverPoint for VMs.....	12
---	---------------------------	----

FIGURES

TABLES

1	Add vCenter server.....	19
2	KVSS syntax.....	37
3	RPA and cluster bottlenecks.....	60
4	Consistency group and link bottlenecks.....	61
5	Glue script samples.....	65
6	Virtual machine network settings available through the GUI.....	67
7	Network settings only available through the CSV file.....	71

TABLES

Preface

As part of an effort to improve product lines, we periodically release revisions of software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your technical support professional if a product does not function properly or does not function as described in this document.

Note

This document was accurate at publication time. Go to Online Support (<https://support.emc.com>) to ensure that you are using the latest version of this document.

Purpose

This document includes conceptual information on managing a RecoverPoint for Virtual Machines system.

Audience

This document is intended for use by storage administrators who are responsible for managing the RecoverPoint for Virtual Machines system.

Related documentation

The following publications provide additional information:

- *RecoverPoint for Virtual Machines Release Notes*
- *RecoverPoint for Virtual Machines Quick Start Installation Poster*
- *RecoverPoint for Virtual Machines Basic Configuration Installation Guide*
- *RecoverPoint for Virtual Machines Installation and Deployment Guide*
- *RecoverPoint for Virtual Machines Product Guide*
- *RecoverPoint for Virtual Machines Administrator's Guide*
- *RecoverPoint for Virtual Machines CLI Reference Guide*
- *RecoverPoint for Virtual Machines Deployment REST API Programming Guide*
- *RecoverPoint for Virtual Machines REST API Programmer's Guide*
- *RecoverPoint for Virtual Machines Security Configuration Guide*
- *RecoverPoint for Virtual Machines Scale and Performance Guide*
- *RecoverPoint for Virtual Machines FAQ*
- *Recoverpoint for Virtual Machines Simple Support Matrix*

In addition to the core documents, we also provide White papers and Technical Notes on applications, arrays, and splitters.

Typographical conventions

This document uses the following style conventions:

Bold

Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)

<i>Italic</i>	Used for full titles of publications referenced in text
Monospace	Used for: <ul style="list-style-type: none">• System code• System output, such as an error message or script• Pathnames, filenames, prompts, and syntax• Commands and options
<i>Monospace italic</i>	Used for variables
Monospace bold	Used for user input
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

Technical support, product, and licensing information can be obtained as follows:

Product information

For documentation, release notes, software updates, or information about products, go to Online Support at <https://support.emc.com>.

Technical support

Go to Online Support and click Service Center. You will see several options for contacting Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your sales representative for details about obtaining a valid support agreement or with questions about your account.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to techpubcomments@emc.com.

CHAPTER 1

Getting started

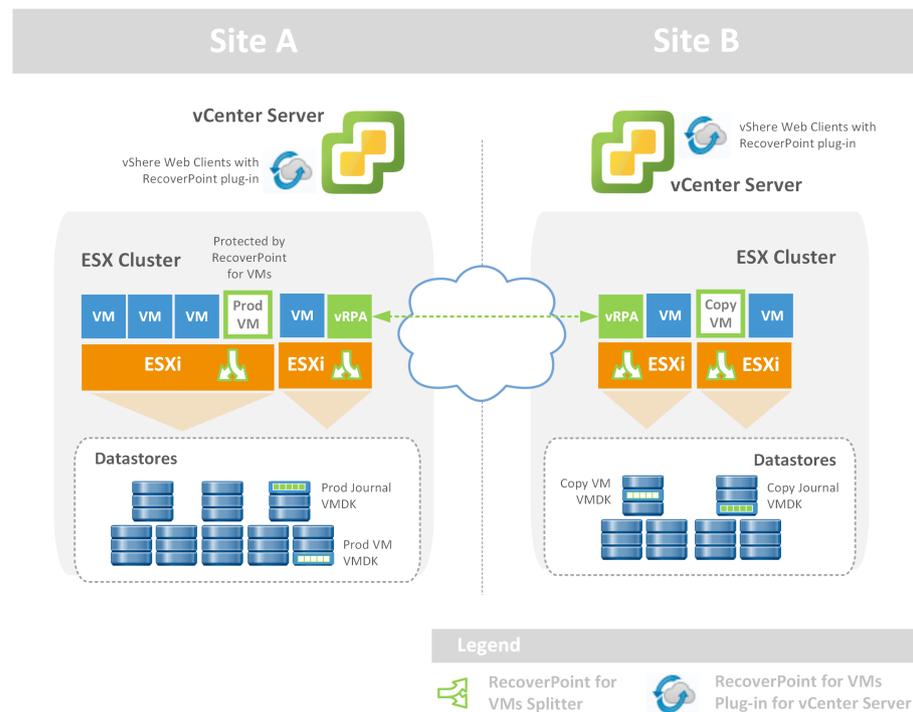
This chapter contains instructions for getting started using RecoverPoint for Virtual Machines.

- [RecoverPoint concepts](#)..... 12
- [Activating entitlements](#)..... 12
- [Licensing, support, and registration](#)..... 14
- [Exporting the post-deployment form to a CSV file](#)..... 16
- [Registering RecoverPoint by email or phone](#)..... 16
- [Registering datastores](#)..... 17
- [Registering ESX clusters](#)..... 17
- [Registering vCenters](#)..... 18

RecoverPoint concepts

All of the operational concepts are described in detail in the *RecoverPoint for VMs Product Guide*. Ensure that you are familiar with the concepts that are described in the product guide before initiating any of the procedures that are outlined in this document.

Figure 1 RecoverPoint for VMs



Activating entitlements

Once the RecoverPoint for Virtual Machines sales order is approved, a License Authorization Code is automatically sent to the email addresses provided during order entry, containing all of the customer's entitlements. Each entitlement must be activated and then saved as a license file before it can be added to the RecoverPoint system.

- Licenses can be partially or fully activated.
- Licenses are based on the number of supported virtual machines per vCenter Server. Only production VMs are counted in the number of supported virtual machines per vCenter Server.
- Licensing is enforced using the vCenter Server ID.
- All vCenter Servers must be registered in RecoverPoint for VMs before their licenses can be added. vCenter Server registration is performed using the Deployment Manager. For the procedures, refer to "Connecting vRPA clusters" and "Registering additional vCenter Servers" in the *RecoverPoint for VMs Installation and Deployment Guide*.
- When reaching the maximum number of virtual machines that are supported per vCenter Server, you cannot protect new virtual machines or enable disabled

consistency groups. Replication of existing virtual machines and consistency groups continues.

Procedure

1. To activate RecoverPoint for VMs entitlements, access the entitlements on emc.support.com:
 - If you have the License Authorization Code email, open it and click the **Click here** link. Clicking the link automatically accesses Powerlink Licensing on the Online Support site, and searches for the entitlements associated with the License Authorization Code.
 - If you do not have the License Authorization Code email but you do have the License Authorization Codes or sales order numbers, log in to Online support at <http://support.emc.com>, and
 - a. Select **Support > Service Center** from the main menu.
 - b. Select **Get and Manage Licenses**.
 - c. Select **RecoverPoint for Virtual Machines**.
 - d. Type the customer's License Authorization Code and click **Activate** to search for all inactive entitlements that are associated with a customer's profile, or access all of the features of the Licensing site by clicking **Manage Entitlements**. Whichever option you chose, the **Search Entitlements to Activate** screen is displayed.
2. Activate the entitlements and download the license files.
 - a. In the **Search Entitlements to Activate** screen, select an entitlement to activate. Each entitlement must be selected and activated separately.
 - b. Click **Start Activation Process**.
 - c. In the **Search Machines** dialog box, click **Add a Machine**.
 - d. In the **Add Machine** dialog box, type a new machine name, and click **Save**. A unique machine name must be specified for each entitlement.

A machine name is like a folder. It is used to group items together logically.
 - e. In the **Register** screen, verify the machine name, and click **Next**.
 - f. In the **Activate** screen, type the **Locking ID**, and click **Next**.

The Locking ID is the field that is displayed in the Machine Information column. Its value is the entity that the license is enforced for, namely, the vCenter Server ID. To find the vCenter Server ID, type `https://<vCenterServerIP>/mob` into the browser address bar or SSH client, and type the credentials to log in to the vCenter Server. Select **Content > About**. The instanceUuid is the vCenter Server (Locking) ID that the license is enforced for.
 - g. In the **Confirm** screen, type the email addresses of the recipients of the license file in the **Email to** field of the **Additional Email Options** section, and click **Finish**. Separate multiple email addresses with commas.
 - h. In the **Complete** screen, click **Save to File** to download the license file and save the file locally. The resulting license file has a `*.lic` extension and is in plain text format (can be opened in any text editor).
 - i. Repeat this procedure for all inactive entitlement in each License Authorization Code email.

3. After you have turned all of the entitlements into license files, physically transfer the license files to the computer from which you will be running RecoverPoint for VMs.

Licensing, support, and registration

The first time that you run RecoverPoint for VMs, the **Getting Started Wizard** guides you through configuring the basic RecoverPoint for VMs settings and ensuring the system is up and running.

Before you begin

- To enable support, a permanent RecoverPoint for VMs license must exist in the system. Ssystem reports and alerts do not work with a temporary license.
- Best practice is to keep both system reports and alerts, and compression and encryption enabled.
- System reports and alerts require a valid method of transfer (SMTP, Secure Remote Services, or FTPS) is configured. Secure Remote Services is the recommended method of transfer.
- To transfer system reports and alerts using SMTP or Secure Remote Services, ensure that port 25 is open and available for SMTP traffic.
- To transfer system reports and alerts using FTPS, ensure that ports 990 and 989 are open and available for FTPS traffic.

Procedure

1. In the Getting Started Wizard **Welcome** screen, click **Next Add Licenses**.
2. In the **Add Licenses** screen, click **Add**. In the **Add license** dialog box, type the location of the license file or click **Browse** to locate the file. Click **OK**. Click **Next Enable Support**.
3. To provide communication between the RecoverPoint for VMs system and the System Reports database, in the **Enable Support** screen, select **Enable pre-emptive support for RecoverPoint for VMs**.
4. Define the transfer method:
 - To transfer system notifications through an SMTP server, in the **Transfer Method** section, select **SMTP**. In the **SMTP server address** field, specify the IP address or DNS name of the dedicated SMTP server, in IPv4 format. In the **Sender address** field, specify the email address to send the system notifications from. Click **OK**.
 - To transfer system notifications through RecoverPoint's built-in FTPS server, in the **Transfer Method** section, select the **FTPS radio button**. Click **OK**.
 - To transfer system notifications through the Secure Remote Services gateway, in the **Transfer Method** section, select the **EMC Secure Remote Services radio button**. In the **ESRS gateway IP address** field, specify the IP address of the Secure Remote Services gateway in IPv4 format. Click **OK**.
5. Click **Test Connectivity**. Wait 10 minutes. Then, create an `ssh` connection to the cluster management IP address, run the `get_events_log` command, and look for event 1020 "Failed to send system report".
 - If this event does not appear in the event logs, the system notifications mechanism has been correctly configured.

- If you do receive an event 1020Failed to send system report, check whether there is an issue with the selected method of transfer. If a problem exists, fix it, configure support, and click **Test Connectivity** again. If the problem persists, contact Customer Support.
6. Click **Next Register RecoverPoint**.
 7. In the **Register RecoverPoint** screen, register or re-register each cluster in the RecoverPoint system after every RecoverPoint system installation, after connecting vRPA clusters in a RecoverPoint system, and after upgrading a RecoverPoint system.
 - a. To display the **Update Post-Deployment Form Details** dialog box, click **Edit Settings...**
 - b. In the **Update Post-Deployment Form Details** dialog box, update the form information.

Option	Description
Company name	The name of the company
Connect in method	<p>The method that is used to allow remote connectivity to the RecoverPoint environment. Enabling this feature is recommended as it allows secure access to the RecoverPoint environment to gather logs and resolve issues as expeditiously as possible.</p> <p>If you already have a Secure Remote Services Gateway servicing other products, use the Secure Remote Services Config Tool to add the RecoverPoint devices to the list of Secure Remote Services monitored environments. Once the device is added, click the request update button to send the new device information to EMC and contact the local Customer Engineer to approve the update. Refer to the <i>EMC Secure Remote Services Gateway Operation Guide</i> for further instructions on Config Tool usage.</p> <p>If you do not have a Gateway at the site, contact the Account Manager to find out more about the benefits of Secure Remote Services.</p>
Location	The city, state, and country where the customer is located.
Sales order numbers	The customer or Customer Engineer should provide this information.
Site (party) ID	The unique ID of the customer site. This value is automatically inserted and taken from the license file and can only be modified by contacting Customer Support.
VCE	Indicate whether this RecoverPoint implementation is operating within a VCE™ (Vblock) environment. VCE = VMware+Cisco+EMC.
Activity type	Type the kind of activity you are performing (upgrade, installation)
Resource performing this upgrade/ installation	Type the role of the person performing this upgrade or installation activity

Option	Description
Connect home	The method that is used to send configuration reports and alerts to EMC. Enabling this feature is recommended as it allows EMC to pro-actively address issues within the RecoverPoint environment, should they arise. "Configure System Reports and Alerts" to establish connectivity for the RecoverPoint environment.

If the company does not have outside connectivity, and therefore, you cannot configure system reports and alerts, you can skip the rest of the steps in this procedure.

8. Verify that the details in the **Registration** screen are correct.
9. Click **Send Form** to automatically send the post-deployment form to the Install Base group. In the **Send Form** dialog box, type the email address of the person at the company that is in charge of RecoverPoint maintenance and operation. A service request is opened and sends an email to the specified verification email address from Customer Support to verify that the registration details were updated successfully in the Install Base.

Exporting the post-deployment form to a CSV file

Save the RecoverPoint registration information or register RecoverPoint by email or phone by exporting the RecoverPoint post-deployment form and all of its contents as a comma-delimited *.csv file.

Procedure

1. Select **Administration > vRPA Clusters > Registration**.
2. Select the vRPA cluster for which you want to export a post-deployment form.
3. Click the **Export to CSV** button and save the file to the computer.
4. Open the exported file in Microsoft Excel. The Excel **Text Import Wizard** is displayed to help you set the import options. In the **Excel** dialog box, select **Delimited**, and click **Next**. In the **Delimiters** field, select **Comma**, and click **Next**. Click **Finish**.

Registering RecoverPoint by email or phone

Registers RecoverPoint if the company does not have external connectivity, and therefore you cannot register the RecoverPoint system online.

Before you begin

- Register the RecoverPoint system after:
 - Installing a RecoverPoint system
 - Connecting RPA clusters in a RecoverPoint system
 - Upgrading a RecoverPoint system
- The registration process is incomplete if valid values are not provided for every field in the post-deployment form.

Procedure

1. Access <https://support.emc.com>

2. Search for the term *Post-Deployment Form*
3. Download and fill out the RecoverPoint and RecoverPoint for VMs Post-Deployment Form
4. Send the information to the Install Base group:
 - Customers and partners: Email the post-deployment form to the Install Base group at rp.registration@emc.com.
 - Employees:
 - (Preferred) Use the IB Portal at <http://emc.force.com/BusinessServices>.
 - Call in the information to the Install Base group at 1-866-436-2411 – Monday to Friday (normal Eastern Time Zone working hours).

Registering datastores

Registers a datastore at a vRPA cluster.

Procedure

1. Browse to the datastore management tab.

Select **Administration > vRPA Clusters > Related Objects**. Select the vRPA cluster at which you want to register datastores, and click **Add...** under the **Datastores** widget.

The **Register Datastores** dialog box is displayed.

2. In the **Register Datastores** dialog box:
 - a. Select the vCenter server that manages the datastores.
 - b. Select one or more datastores to register.
 - c. Click **OK**.

Results

The specified datastore is registered at the selected vRPA cluster.

Registering ESX clusters

Registers ESX clusters at a vRPA cluster.

Before you begin

ESX clusters that host protected or copy VMs must be registered before you can protect virtual machines. See [protect VMs](#).

Procedure

1. In the vSphere Web Client home page, select **RecoverPoint for VMs Management > Administration > vRPA Clusters**.
2. Select the vRPA cluster at which you want to register ESX clusters.
3. Select the **ESX Clusters** tab.
4. Click **Add**.
5. In the **Register ESX Clusters** dialog box:
 - a. Select the ESX cluster that you want to register.

- b. Click **OK**.

Results

The specified ESX cluster is registered at the selected vRPA cluster.

Registering vCenters

Registers or edits the registration details of a vCenter server at a specific vRPA cluster, or all vRPA clusters in the RecoverPoint for VMs system.

Before you begin

- Best practice is to configure the vCenter Server to require a certificate, because once RecoverPoint has read the certificate, it does not need further access to the location.
- The default certificate locations are:
 - **Windows 2003 Server:**
C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\SSL\rui.crt.
 - **Windows 2008 Server:**
C:\Users\All Users\Application Data\VMware\VMware VirtualCenter\SSL\rui.crt.

For more information about the location of the security certificate, refer to "Replacing vCenter Server Certificates in VMware vSphere 5.0, 5.5 and 6.0," available at www.vmware.com.

Procedure

1. In the vSphere Web Client home page, select **RecoverPoint for VMs Management > Administration**.

2. Access the **vCenter Registration** information:

There are two ways to access the registration details of the vCenters in RecoverPoint.

- To manage the registration of all vCenter servers in a RecoverPoint for VMs system select **vCenter Servers > Registration** and use the **Edit** icon to edit the vCenter settings. Use this option to:
 - Edit the vCenter server information, upload a new vCenter certificate, or delete an existing certificate.
 - Propagate the changes to the specified vCenter server at the specified vRPA cluster using the **Apply** button.
 - Propagate the changes to all vRPA clusters in the system using the **Apply changes to all clusters** button.
 - To manage the registration of a vCenter server at a specific RPA cluster select **vRPA Clusters > vCenter Servers**, select a vRPA cluster, and:
 - To edit the registration details of an existing vCenter server at the selected vRPA cluster, click the **Edit** icon.
 - to register a new vCenter server at the selected vRPA cluster, click the **Add** button.
3. In the **Register vCenter Server** dialog box, type the following information:

Table 1 Add vCenter server

Setting	Description
vCenter Server IP	IP address of the vCenter Server. This IP address is also the display name of the vCenter Server in RecoverPoint.
Port	Port number of the vCenter Server. Default = 443 (HTTPS).
Username	vCenter Server username.
Password	vCenter Server password.
Certificate	To specify a certificate, browse to and select the certificate file.

4. Click **OK**.

Results

The specified vCenter Server is registered at the specified vRPA clusters with the specified details.

CHAPTER 2

Protecting your data

This chapter contains instructions for protecting your data using RecoverPoint for Virtual Machines.

- [Overview](#)22
- [Protecting virtual machines](#)22
- [Stopping protection](#) 25

Overview

- Host and storage VMware vSphere vMotion are both supported for vRPAs and for protected virtual machines.
- When production VMDKs are expanded, RecoverPoint for VMs automatically expands all corresponding copy VMs in the group. Refer to "Expanding copy VMDKs automatically."
- VMware snapshots:
 - Taking VMware snapshots is supported only on production VMs and not on copy VMs. Using the `VM Restore` operation to restore a production VM from a VMware snapshot or clone causes a full sweep.
 - Before creating a VMware snapshot of a vRPA, it must be detached from the vRPA cluster. See the *RecoverPoint for VMs Installation Guide* for the procedure for detaching vRPAs from a cluster.
 - Promoting VMware snapshots on a copy VM is not supported.
- Replicating a virtual machine with fault tolerance enabled is not supported.
- To clone a copy VM, you must first use the **Test a Copy** wizard to enable image access, and then power off the virtual machine. If the virtual machine is not powered off, a VMware snapshot is created, which can cause data corruption.
- The resources (memory and CPU) of a running vRPA can be increased but not decreased. The changes take effect only when the vRPA is rebooted.
- Any `suspend` or `resume` tasks (or hibernation-like functionality) of a running vRPA virtual machine is not supported and may lead to data corruption.
- Collecting vSphere logs from a running vRPA (for instance, using vSphere **File > Export > Export System Logs** command) may cause inconsistencies in the datastore or other unexpected behavior.
- Upgrading VMware Tools on the vRPAs is unsupported and may lead to undesired results. There is a default version of VMware tools that is installed and shipped with vRPAs, and this special version of VMware tools should never be upgraded.

Protecting virtual machines

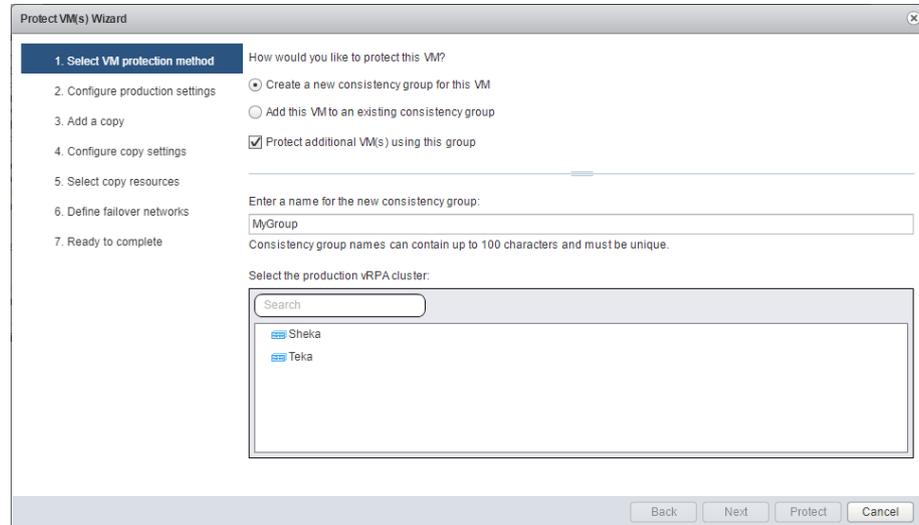
Before you begin

Perform [Licensing, support, and registration](#) on page 14, [Registering ESX clusters](#) on page 17, and [Registering datastores](#) on page 17.

To protect a virtual machine:

Procedure

1. In vSphere, right-click the virtual machine to protect and select **All RecoverPoint for Virtual Machines Actions > Protect**.
2. In the **Select VM protection method** screen:



- **Create a new consistency group for this VM.** Type a descriptive name for the new consistency group. Select the production vRPA cluster. If you want to add additional virtual machines to protect, mark the **Protect additional VM(s) using this group** checkbox, select the additional virtual machines to protect in the consistency group, and click **Add**. If you do not want to add additional virtual machines, click **Next**.
- **Add this VM to an existing consistency group.** Select an existing consistency group. If you want to add additional virtual machines to protect, mark the **Protect additional VM(s) using this group** checkbox select the additional virtual machines to protect in the consistency group, and click **Add**. If you do not want to add additional virtual machines, click **Next**.

Note

When a virtual machine is added to an existing consistency group, if the virtual machine image is larger than the allotted journal size, the system automatically enters *one-phase distribution mode*.

3. In the **Configure production settings** screen:
 - a. Expand and configure the **Advanced options** per virtual machine:
 - **VMDK(s):** Displays the number of included VMDKs at the relevant production copy, and their total size. Uncheck a VMDK to exclude it from replication.
 - **Protection policy:** Default = **Enabled**. Selecting **Automatically protect new VMDKs** ensures all new VMDKs are automatically protected.
 - **Disk provisioning:** Default = **Same as source**. Defines the way in which the copy VMDKs are to be provisioned; **Same as source**, **Thick provisioning**, or **Thin provisioning**. When **Thick provisioning** or **Same as source** are selected, if the production VMDKs are thick (either eager or lazy), the copy VMDKs are thick provisioned lazy zero.
 - **Hardware changes:** Default = **Enabled**. Automatically replicates the hardware settings of all production virtual machines to their copy VMs whenever an image is accessed on the copy VMs. When enabled, RecoverPoint for VMs replicates the virtual machine version, CPU, memory, resource reservations, and network adapter status and type. Replication of SR-IOV Passthrough Adapter is not supported. If the ESX

at a copy does not support the production VM version, no hardware resources are replicated.

- **MAC address replication to local copy VMs on the same vCenter:**
Default = `Disabled`. If two *remote copies* of the same production VM are on the same vCenter and in the same network, you cannot power on both copy VMs simultaneously, as they have the same MAC address. Therefore, by default, the MAC address of remote copy VMs on a different vCenter than their production VMs is replicated to the copy. However:
 - When **Replicate hardware changes** is disabled, **MAC address replication** is also disabled.
 - To avoid IP conflicts, by default, the MAC address is not replicated for local copy VMs on the same vCenter as their production VMs. If a copy VM is not on the same network and ESX as its production VM, select **Enable for local copy VMs managed by this vCenter** to replicate the MAC address.

b. If you chose to create a consistency group in the previous step, also:

- Define the minimum **Journal Size** for the production copy.
- Either allow RecoverPoint for VMs to automatically select a registered datastore for the specified journal size or manually select one from the table.
- If the table does not contain the required datastore, click **Register Datastore** and select the datastore to register.

c. Click **Next**.

4. In the **Add a copy** screen, type a name for the copy and select the vRPA cluster to use to manage the group data.
5. In the **Configure copy settings** screen, configure the copy journal:
 - Define the minimum **Journal Size** for the production copy.
 - Either allow RecoverPoint for VMs to automatically select a registered datastore for the specified journal size or manually select one from the table.
 - If the table does not contain the required datastore, click **Register Datastore** and select the datastore to register.
 - Load the replication policy for this copy from a template or manually define it. By default, the replication policy is set to `Synchronous` mode and the `RPO` (Recovery Point Objective) is set to `25 Seconds`. The RPO is the point in time to which you are required to recover data, for a specific application, as defined by the organization. RPO defines the maximum lag that is allowed on a link, and is set manually in `Bytes`, `KB`, `MB`, `GB`, `TB`, `Writes`, `Seconds`, `Minutes`, `Hours`.
6. In the **Select copy resources** screen, select how to protect the virtual machine at the target vRPA cluster:
 - To **Automatically create new copy VM(s)**, expand the tree, select an ESX host, cluster, or resource pool. Note that ESX clusters are only displayed after [Registering ESX clusters](#) on page 17. Click **Next**. Repeat this process for all production VMs. In the **Select copy storage** screen, select the datastore where you want to place the data of the copy VM. All VMDKs are mapped to a single datastore. Click **Next**.

- To **Manually select an existing VM to use as the VM copy**, expand the tree, select an ESX host, cluster, or resource pool, and select an existing virtual machine. Repeat this process for all production VMs, click **Next**, and skip the next step.
7. In the **Define failover networks** screen, select the networks to be used after failover, and click **Next**.

NOTICE

Skip this step or select *system-defined* to allow RecoverPoint for VMs to automatically select the networks to be used after failover.

8. In the **Ready to complete** screen, if you do not want to start protection immediately, uncheck **Start replicating this group when I click Protect**.
9. In the **Ready to complete** screen, review the group details.
Click **Add a Copy** to add more copies to the group, **Edit...** to modify an existing copy, or **Protect** to start replication from the production to all copies.

Stopping protection

To stop replication of a production VM:

Procedure

1. In the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon > Protection** tab. Click **Virtual Machines**.
2. Select the virtual machine to stop replicating. Click the **Unprotect** icon:



Results

Replication stops and the virtual machine is removed from its consistency group. The copy VM is not automatically deleted. If there are no other virtual machines in the consistency group, the consistency group is removed. If other virtual machines remain in the consistency group, the journal is lost.

CHAPTER 3

Managing and monitoring the system

This chapter contains the instructions for managing and monitoring RecoverPoint for Virtual Machines.

- [Monitoring the system](#)..... 28
- [Managing the system](#)..... 29
- [Managing group sets](#)..... 31
- [Managing consistency groups](#)..... 33
- [Managing copies](#)..... 39
- [Managing virtual machines and VMDKs](#)..... 41

Monitoring the system

The RecoverPoint for VMs Dashboard provides a high-level overview of the RecoverPoint system. It presents important system information to help you analyze and monitor the RecoverPoint environment.

Procedure

1. To access the **Dashboard**, in the **vSphere Web Client** home page, click the **RecoverPoint for VMs** icon.
2. To monitor the environment at the system level, click **Dashboard** to browse through the **Overall Health**, **Recovery**, **Components**, **System Limits**, and **Events** sub-tabs.
3. To monitor consistency groups, select **Protection > Consistency Groups**.
4. To search the consistency groups list:
 - a. Select **Consistency Groups** in the left navigation pane.
 - b. Click the **Select filter columns** icon if you want to limit the search to specific columns.
 - c. To type the filter criteria, click inside the **Filter** text box.
5. To monitor a consistency group:
 - a. Expand the **Consistency Groups** tree and select a specific consistency group in the left pane.
 - b. To monitor the write-rate and throughput of the selected group, select the **Details**, **Topology**, or **Statistics** sub-tabs.

Reviewing recovery activities

Use recovery activity reports to display each of the steps in the recovery activities (testing a copy, recovering production, and failing over), the time each step took, and the completion status of the step.

Before you begin

Before you begin, note the following about activity reports:

- Users have access to up to 10 reports per consistency group.
- Reports can only be produced through the vCenter GUI.
- Users must export or view the reports in the vCenter GUI.
- Users must manually export reports from vCenter.
- Reports can only be exported to CSV format.
- Every vRPA clock must be synchronized within their time zone to prevent inconsistencies in the report timestamps.
- In case of disaster, reports may be missing or inaccurate.

Procedure

1. In the RecoverPoint for VMs plug-in, browse to the **Reports** tab.
2. Expand the consistency group tree.
3. Select the consistency group that you want the report for.

The Recovery Activities are displayed in the right-most pane.

4. Select the desired activity.
 - To export the selected activity report, click the **Export to CSV** button.
 - To remove an activity report from the list, click the **Remove** button.
 - To change the time zone, click **Change to GMT** or **Change to local time**.

The Activity Report for that activity is displayed. Within each activity report, you can expand the report to view each step.

Monitoring consistency group replication

Procedure

1. In the vSphere Web Client, select the virtual machine that you wish to monitor.
2. Select the **Manage** tab and the **RecoverPoint for VMs** subtab.

Results

A graphical representation of the virtual machine's consistency group is displayed.

Consistency group transfer states

This topic lists the possible transfer states of a consistency group or consistency group copy. Hover your mouse over a transfer state on screen to see the names of the consistency groups in that state.

Transfer State	Description
Active	Data is being transferred asynchronously to a copy.
Active (Synchronized)	Data is being transferred synchronously to a copy.
Paused by system	Data is not being transferred to a copy, because transfer has been paused by the system.
Init (n%)	A copy is being initialized or undergoing a full sweep.
High-load (n%)	The system enters a temporary high-load state while data is being transferred to a copy. High-load occurs when the journal is full and cannot accept new writes. The system will attempt to resolve the high-load state without user action.
High-load	The system enters a permanent high-load state while data is being transferred to a copy. A permanent high-load can occur after a temporary high-load. The system pauses replication and waits for user action.
N/A	Data is not being transferred to a copy, because the copy has been disabled by the user.

Managing the system

Licenses

Before you begin

When there is no license that is installed in the system, clicking **Add** in the Licensing screen opens the **Getting Started Wizard** to guide you through the process of enabling support and registering the RecoverPoint for VMs system.

Procedure

1. Browse to the license management tab.
Select **Administration > vRPA Clusters > Licensing**.
2. Do one of the following:
 - To add a new license to the system, click **Add..**
 - To delete an existing license form the system, click **Remove**.

Support settings

Procedure

1. Select **Administration > vRPA Clusters > Support** tab.
2. Select a vRPA cluster.
3. To edit the cluster support settings, click **Edit Settings...** .

After you finish

For details on how to configure the support settings, see [licensing, support, and registration](#).

RecoverPoint for VMs registration

Procedure

1. Select **Administration > vRPA Clusters > Registration** tab.
2. Select a vRPA cluster.
3. To manage the cluster registration settings, click **Edit Settings...** .

After you finish

For details on how to register RecoverPoint for VMs, See [licensing, support, and registration](#).

Registering an external host

Defines the external host on which user scripts are run during virtual machine start-up sequences.

Before you begin

- SSH must be installed on the external host.
- Only one external host can be configured per vRPA cluster.
- Define the external host before defining virtual machine start-up scripts in a virtual machine startup Sequence. For information on how to define start-up scripts, see [VM start-up sequence](#).

Procedure

1. Browse to the external host management tab.
Select **Administration > vRPA Clusters > Related Objects**. Select the vRPA cluster for which you want to define an external host, and click **Edit...** under the **External Host** widget.
2. In the **Edit External Host Configuration** dialog box, type the **Name**, **IP**, **User**, and **Password** of the external host for the selected vRPA cluster.

3. Optionally:
 - To verify connectivity with the external host, click **Check Connectivity** .
 - To unregister the external host from the specified vRPA cluster, click **Remove**.

Policy templates

Adds a policy template, edits, imports, or removes an existing policy template.

Procedure

1. Browse to the policy template management tab:
Select **Administration > vRPA Clusters > vRPA System** .
2. Manage the **Policy Templates**:
 - To add a policy template, click **Add**.
 - To modify the settings of an existing policy template, select a policy template and click **Edit....**
3. In the **Add/Edit Policy Template** dialog box:
 - To configure a link policy template, type a name for the policy template and [define the group or link policy settings](#).
 - To configure a copy policy template, type a name for the policy template and [define the copy policy settings](#).
4. Optionally:
 - To delete a policy template, select a policy template and click **Remove**.
 - To import to all RPA clusters in the system, click **Import**, and select a policy template.

Managing group sets

This section describes how to manage group sets in the RecoverPoint for VMs system.

Creating a group set

A group set is a collection of consistency groups to which the system applies parallel bookmarks at a user-defined frequency. Group sets are useful for consistency groups that are dependent on one another or that must work together as a single unit.

Procedure

1. In the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon > Protection** tab.
2. Click **Consistency Groups**
3. Click the **Add Group Set** icon:



4. In the **Add Group Set** dialog box, type a name for the group set.

5. Choose the vRPA cluster from which to select consistency groups.
6. To add to the group set, select one or more consistency groups.
7. To enable parallel bookmarking, select **Enable Parallel Bookmarking** and set the bookmarking frequency value. Click **OK**.

If any of the groups in the group set are part of another group set that has parallel bookmarking enabled, you cannot enable parallel bookmarking for that group set.

Editing an existing group set

Before you begin

If any of the groups in the group set are part of another group set that has parallel bookmarking enabled, you cannot enable parallel bookmarking for that group set.

Procedure

1. In the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon** > **Protection** tab. Click **Group Sets**.
2. Select the group set to edit. Click the **Edit Group Set** icon:

3. In the **Edit Group Set** dialog box, if desired, modify the group set name.
4. Select or clear consistency groups to include or exclude them from the group set.
5. Enable or disable parallel bookmarking by selecting or clearing **Enable Parallel Bookmark**. Click **OK**.

Enabling a group set

Enabling a disabled group set causes a full sweep and then starts replication in all of the consistency groups in the group set.

Procedure

1. In the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon** > **Protection** tab. Click **Group Sets**.
2. Select the group set to enable. Click the **Enable Group Set** icon:



Removing a group set

Procedure

1. In the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon** > **Protection** tab. Click **Group Sets**.
2. Select the group set to remove. Click the **Remove Group Set** icon:



Managing consistency groups

This section describes how to manage consistency groups in the RecoverPoint for VMs system.

Enabling or disabling a consistency group

Enabling a consistency group starts replication and causes a full sweep. Disabling a consistency group stops all replication, deletes journals.

Procedure

1. In the **vSphere Web Client** home page, click **RecoverPoint for VMs Management icon** > **Protection** tab. Click **Consistency Groups**.
2. Select the consistency group that you want to enable or disable. Click the **Enable Group** icon or the **Disable Group** icon:



Editing group or link policies

To load an existing link policy template, click on the **Load link policy from template** link and select the template.

Procedure

1. In the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon** > **Protection** tab. Click **Consistency Groups**.
2. Expand the list of consistency groups and select the consistency group whose policies you wish to edit.
3. Click on the **Edit group policy** link to change the consistency group name, primary RPA, or group priority.

Option	Description
Name	The name of the consistency group.
Primary RPA	The vRPA that you prefer to replicate the consistency group. When the primary vRPA is not available, the consistency group will switch to another vRPA in the vRPA cluster. When the primary vRPA becomes available, the consistency group will switch back to it.
Group priority	Only relevant for remote replication when two or more consistency groups are using the same Primary RPA. Default = Normal Select the priority assigned to this consistency group. The priority determines the amount of bandwidth allocated to this consistency group in relation to all other consistency groups using the same Primary RPA.

If the cluster is a single-vRPA cluster, the vRPA is the primary (and only) vRPA and is a single point of failure in cases of disaster. Consider adding additional vRPAs to this cluster to ensure high availability.

4. Click on the **Edit link policy** link to edit the link policy protection settings:

Option	Description
Replication Mode	
Dynamic by Latency	<p>Only relevant for synchronous replication mode. Default = Disabled</p> <p>When Enabled, RecoverPoint for VMs alternates between synchronous and asynchronous replication modes, as necessary, according to latency conditions.</p> <p>Start async replication above When the specified limit is reached, RecoverPoint for VMs automatically starts replicating asynchronously</p> <p>Resume sync replication below When the specified limit is reached, RecoverPoint goes back to replicating synchronously.</p>
Dynamic by Throughput	<p>Only relevant for synchronous replication mode. Default = Disabled</p> <p>When enabled, RecoverPoint for VMs alternates between synchronous and asynchronous replication modes, as necessary, according to throughput conditions.</p> <p>Start async replication above When the specified limit is reached, RecoverPoint for VMs automatically starts replicating asynchronously</p> <p>Resume sync replication below When the specified limit is reached, RecoverPoint goes back to replicating synchronously.</p>
RPO	RPO defines the maximum lag allowed on a link, and is set manually in MB, GB, writes, seconds, minutes, or hours.
Compression	<p>Only relevant for asynchronous remote replication. Default = None</p> <p>To compress data before transferring it to a remote vRPA cluster, select a level of compression. Compression can reduce transfer time significantly, but increases the source vRPA's CPU utilization.</p> <p>Enabling and disabling compression causes a short pause in transfer and a short initialization.</p>
Enable Deduplication	<p>Only relevant for asynchronous remote replication. Default = Disabled</p> <p>Select this to eliminate repetitive data before transferring the data to a remote vRPA cluster. Deduplication can reduce transfer time significantly, but increases the source vRPA's CPU utilization.</p>

Option	Description
	Enabling and disabling deduplication causes a short pause in transfer and a short initialization.
Snapshot Granularity	<p>Default = fixed per second.</p> <p>Fixed per write Creates a snapshot for every write operation.</p> <p>Fixed per second Creates one snapshot per second. Use this for local replication.</p> <p>Dynamic The system determines the snapshot granularity according to available resources. Use this for remote replication.</p>

Creating bookmarks

A bookmark is a name that is applied to a snapshot to identify it for future use. Bookmarks can be applied to consistency groups or group sets. Crash-consistent bookmarks are created using RecoverPoint for VMs plug-in for vCenter. Application-consistent bookmarks are created using RecoverPoint's VSS-based utility, called KVSS.

Creating crash-consistent bookmarks

Procedure

- In the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon > Protection tab**.
 - To apply a bookmark to a consistency group, Click **Consistency Groups**.
 - To apply a bookmark to a group set, click **group sets**.
- Select the consistency group or group set to which you want to apply a bookmark. Click the **Add bookmark** icon:

- In the **Apply a Bookmark** dialog box, define the bookmark settings:
 - Type a bookmark name.
 - Specify the **Mark as** choice:
 - Bookmark name** - Type a name for the bookmark.
 - Mark as** - Chose one of:
 - Crash-Consistent** - Labels bookmark as Crash-Consistent.
 - Application-Consistent** - Labels bookmark as Application-Consistent. Selecting Application-Consistent does not create an application-consistent snapshot, it only labels the snapshot as application-consistent.
 - Consolidation policy** - To specify how the consolidation policy will be managed the next time that the process runs.

- **Never consolidate this bookmark**
- **This bookmark snapshot must survive Daily/Weekly/Monthly consolidations:**
 - **Daily** - Snapshot survives daily consolidations, but is consolidated weekly and monthly.
 - **Weekly** - Snapshot survives daily and weekly consolidations, but is consolidated monthly.
 - **Monthly** - Snapshot survives daily, weekly, and monthly consolidations.

Creating application-consistent bookmarks

Before you begin

KVSS bookmarks are created using the `kvss.exe bookmark` command.

The working folder for running KVSS commands is `%SystemDriver%/EMCRecoverPointVSSProvider/`.

When using KVSS to apply bookmarks:

- Surround parameter values with quotation marks.
- You can use the `vssadmin list writers` command to obtain a list of registered writers on the host virtual machine.
- You can use the `kvss.exe list` command to display the components of each of the writers found using the `vssadmin list writers` command.
- You can run the `kvss.exe set_credentials` command once per Windows user to define the `ip`, `user`, and `password`. After doing so, you will not need to type these values again.
- If they are separated by a space, you can type multiple writers and groups simultaneously.
- Only the application on which KVSS is run is application consistent, and only when run on the same virtual machine. Best practice is to name the bookmark to contain both the name of the application and the virtual machine.
- Upgrade the vRPA clusters before upgrading KVSS. An older version of KVSS works with a vRPA cluster running a newer version of RecoverPoint for VMs. A newer version of KVSS does not work with a vRPA cluster running an older version of RecoverPoint for VMs.

The syntax is as follows:

```
kvss.exe bookmark
bookmark=bookmark_name
  writers=writer_name writer_name
  [groups=group_name group_name]
  [consolidation_policy=never|survive_daily|survive_weekly|
survive_monthly|always]
  [type=[FULL|COPY]]
[ip=RecoverPoint_cluster_management_ip_address]
[user=RecoverPoint_username]
[password=RecoverPoint_password]
```

NOTICE

Parameters that are surrounded by square brackets [] are optional. Using the `-version` flag prints out the KVSS version number.

Where:

Table 2 KVSS syntax

Option	Description
<code>writers</code>	A VSS-aware host application
<code>groups</code>	RecoverPoint consistency group
<code>bookmark</code>	Name by which you can identify the applied bookmark
<code>consolidation_policy</code>	<p>Consolidation policy to set for this snapshot. Valid values are:</p> <ul style="list-style-type: none"> never. Snapshot is never consolidated. survive_daily. Snapshot remains after daily consolidations, but is consolidated in weekly, monthly, and manual consolidations. survive_weekly. Snapshot remains after daily and weekly consolidations, but is consolidated in monthly and manual consolidations. survive_monthly. Snapshot remains after daily, weekly, and monthly consolidations, but is consolidated in manual consolidations. always. Snapshot is consolidated in every consolidation process, whether manual or automatic. <p>Default = always. If the <code>consolidation_policy</code> parameter is not specified, the snapshot is consolidated in both automatic and manual consolidation processes.</p>
<code>type</code>	<p>The shadow copy type:</p> <ul style="list-style-type: none"> • <i>FULL</i> • <i>COPY</i> <p>This setting is optional. Default = <code>COPY</code>. The writer application determines the settings. Generally, when <code>type = full</code>, backup logs are truncated. When <code>type = copy</code>, backup logs are not truncated.</p>
<code>ip</code>	vRPA cluster management IP
<code>user</code>	RecoverPoint for VMs username

Table 2 KVSS syntax (continued)

Option	Description
password	RecoverPoint for VMs password

Procedure

- To create a bookmark for a Microsoft Exchange application for the first time:

```
kvss.exe set_credentials
  ip="10.10.0.145"
  user="admin"
  password="admin"

kvss.exe bookmark
  writers="Microsoft Exchange Writer"
  groups="exchange\comp1" "exchange\comp2"
  bookmark="exchange hourly snapshot"
  consolidation_policy="survive_daily"
```

- To create a bookmark every subsequent time for a Microsoft Exchange application after defining the *ip*, *user*, and *password* through the `kvss.exe set_credentials` command:

```
kvss.exe bookmark
  writers="Microsoft Exchange Writer"
  groups="exchange\comp1" "exchange\comp2"
  bookmark="exchange hourly snapshot"
  consolidation_policy="survive_daily"
```

Adding a copy

To add a copy to an existing consistency group:

Procedure

1. Select **Protection > Virtual Machines**.
2. Click the **Add a copy** () icon.
3. In the **Add a Copy** screen, type a descriptive name for the copy, select the vRPA cluster to manage the copy, and click **Next**.
4. In the **Configure copy settings** screen:
 - Define the minimum **Journal Size** for the copy journal.
 - Either allow RecoverPoint for VMs to automatically select a registered datastore for the specified journal size or manually select one from the table.
 - If the table does not contain the required datastore, click **Register Datastore** and select the datastore to register.
 - Load the replication policy for this copy from a template or manually define it. By default, the replication policy is set to *Synchronous* mode and the RPO (Recovery Point Objective) is set to 25 *Seconds*. The RPO is the point in time to which you are required to recover data, for a specific application, as defined by the organization. RPO defines the maximum lag that is allowed on

a link, and is set manually in **Bytes, KB, MB, GB, TB, Writes, Seconds, Minutes, Hours**.

5. In the **Select copy resources** screen, select how to protect the VM:
 - To **Automatically create new copy VM(s)**, expand the tree, select an ESX host, cluster, or resource pool. Note that ESX clusters are only displayed after [Registering ESX clusters](#) on page 17. Click **Next**. Repeat this process for all production VMs. In the **Select copy storage** screen, select the datastore where you want to place the data of the copy VM. All VMDKs are mapped to a single datastore. Click **Next**.
 - To **Manually select an existing VM to use as the VM copy**, expand the tree, select an ESX host, cluster, or resource pool, and select an existing virtual machine. Repeat this process for all production VMs, click **Next**, and skip the next step.
6. In the **Define failover networks** screen, select the networks to be used after failover, and click **Next**.

NOTICE

Skip this step or select `system-defined` to allow RecoverPoint for VMs to automatically select the networks to be used after failover.

7. In the **Ready to complete** screen, review the new copy details.
Click **Add a Copy** to add more copies to the group, **Edit...** to modify an existing copy, or **Protect** to start replication from the production to all copies.

Managing copies

This section describes how to manage the copies of a RecoverPoint for VMs system.

Editing copy policies

To edit a copy's protection policy:

Procedure

1. In the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon** > **Protection** tab and select **Consistency Groups**.
2. Expand the list of consistency groups and select the consistency group whose copy policies you want to edit.
3. Expand the consistency group and select the copy whose policies you want to edit.
4. To edit the copy policy protection settings, click the **Edit copy policy** link:
 - **Journal Compression**
Default = none
.
Compresses snapshots in the journal so that more images can be saved in the same journal capacity. Best practice is to compress the journal when forcing synchronous replication. Compression impacts the CPU resources of the target vRPA of the consistency group.
Enabling journal compression while a consistency group is enabled results in the loss of all snapshots in the journal.

- **Maximum Journal Lag**
Default = unlimited
Defines the maximum amount of snapshot data (in bytes, KB, MB, or GB) that can be held in the copy journal before distribution to the copy. In terms of RTO, this lag is the maximum amount of data that would bring the copy up to date with production.
 - **Required Protection Window**
The protection window indicates how far in time the copy image can be rolled back.

To define a required protection window and to specify the length of the required window, select this option. You will be notified if the current window is less than the required window.
 - **Enable RecoverPoint Snapshot Consolidation** - Select this option to enable automatic snapshot consolidation.
Automatic snapshot consolidation cannot be enabled for a group that is part of a group set. When enabled, the Predicted Protection Window is not calculated.
 - **Do not consolidate any snapshots for at least**
Default = 2 days
Define the period during which snapshot data is not to be consolidated. If no daily or weekly consolidations are specified, the remaining snapshots are consolidated monthly.
 - **Consolidate snapshots that are older than x to one snapshot per day for y days.**
Default = 5 days
Snapshots are consolidated every 24 hours. Select Indefinitely to consolidate all subsequent snapshots in 24-hour intervals.
 - If Indefinitely is not selected, and no weekly consolidations are specified, the remaining snapshots are consolidated monthly.
 - If Indefinitely is selected, weekly and monthly consolidations are disabled, and the remaining snapshots are consolidated daily.
 - **Consolidate snapshots that are older than x to one snapshot per week for y weeks**
Default = 4 weeks
Snapshots are consolidated every 7 days.
Select Indefinitely to consolidate all subsequent snapshots in seven-day intervals.
 - If Indefinitely is not selected, the remaining snapshots are consolidated monthly.
 - If Indefinitely is selected, monthly consolidations are disabled, and the remaining snapshots are consolidated weekly.
5. To load an existing copy policy template, click the **Load copy policy from template** link and select the template.

Configuring copy VM failover networks

To change the network that is used by a copy VM after failover:

Procedure

1. Select **Protection > Consistency Groups**.
2. Expand the list of consistency groups, expand the relevant group, and select the relevant copy.
3. Click **Modify failover networks**.
4. In the **Failover Networks of <CopyName>** screen, select a virtual machine to display its network adapters, and for each adapter, select the network to be used after failover.

Results

To use these settings, select *preconfigured failover networks* when defining the testing network for [Testing a copy](#) on page 52 and when defining the **Target Network** before [Failing over](#) on page 54.

Managing virtual machines and VMDKs

This section describes how to manage virtual machines, VMDKs, and their settings after they are in a consistency group.

Orchestration

This section describes the RecoverPoint for VMs features for orchestrating the protection of virtual machines and VMDKs.

Group start-up sequence

The Group start-up sequence defines the order in which the consistency groups in a group set power on when image access is enabled during a recovery activity (such as testing a copy, failover, or production recovery). The group start-up sequence overrides the virtual machine start-up sequence. For more information, see [VM start-up sequence](#).

Procedure

1. In the vSphere web client home page, click the **RecoverPoint for VMs Management icon > Protection tab > Group Sets**.
2. Select a group set.
3. Click the **Edit Start-up Sequence** icon.



4. Select each group and set its **Start-up priority**.

Virtual machine start-up sequence

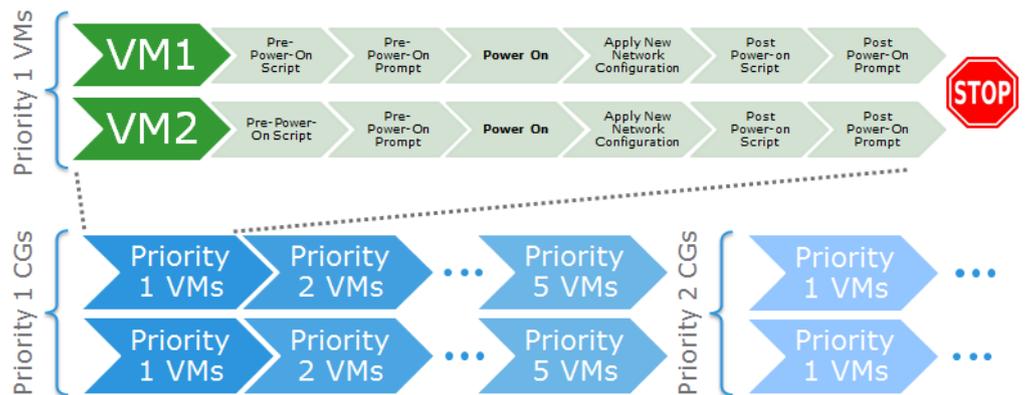
The virtual start-up sequence defines the order of the power on sequence of the virtual machines in a consistency group. The sequence is initiated when image access is enabled during a RecoverPoint for VMs recovery activity (test a copy, after failover or recover production). Virtual machines are powered on in order of priority, as defined by the user. All virtual machines with the same priority power on simultaneously. The startup-sequence can also be defined between consistency groups within the same group set. The start-up sequence can be set as **Critical**. If a critical virtual machine fails to power on, the start-up sequence pauses, and no other virtual machines power on.

Before you begin

- Best practice is to install VMware Tools on each production virtual machine.
- One user script and one user prompt can be configured to run before power on and to run after power on in a strict sequence: **script > prompt > power-up > script > prompt**.

When VMware Tools are installed on a production virtual machine, the virtual machine is considered *powered on* only after its operating system loads. When VMware Tools are not installed on a production VM, the virtual machines is considered *powered on* as soon as it is powered on. Once a virtual machine is *powered on*, the system moves to the next virtual machine in the start-up sequence that you define.

The following graphic illustrates the order of sequences:



Procedure

1. Select **Protection > Consistency Groups**.
2. Expand the consistency group tree, and select the consistency group that you are defining the start-up sequence for.
3. Click the **Edit Start-up Sequence** icon.



The **Start-up Sequence of VMs in this Group** dialog box is displayed.

4. Set the order of the power on sequence by selecting each virtual machine and setting a start-up priority for it.

Option	Description
1	The first virtual machine to power on
3	Default
5	The last virtual machine to power on

5. Optionally, select each virtual machine whose start-up sequence you want to stop if the virtual machine does not power on, and set it to **Critical**.

After you finish

See [Defining user prompts](#) on page 43 and [Defining user scripts](#) on page 43.

Defining user prompts

User prompts define a message to be displayed in the RecoverPoint Dashboard to prompt the user to perform specified tasks before continuing with the start-up sequence. The user must dismiss the prompt before the start-up sequence continues. If the user defines a time-out, the user prompt automatically dismisses if the set time-out period passes. If no time-out is defined and the user does not dismiss the start-up prompt, the start-up sequence does not continue until the user dismisses the prompt.

Before you begin

- You can define one user prompt before power on and one user prompt after power on.

Procedure

1. In the **The Start-up Sequence of VMs in this Group** dialog box, select **Prompt user**.
2. Type a logical name for the prompt.
3. Type the prompt message.
4. Optionally, type a time-out period.

Defining user scripts

A user script runs commands immediately before or after powering on virtual machines. The scripts are executed with `ssh` on the External Host that is provided by the user. Each script has a mandatory time-out period. The recovery flow is blocked until the script executes successfully. If the script does not run within the set time or the script fails or becomes stuck, the system retries the script a pre-defined number of times (set by the user). A prompt indicates if the script failed.

Before you begin

- Maximum size of the script name and parameters = 1024 bytes.
- You can define one user script before power on, and one user script after power on per VM.
- [External host](#) must be configured.
- One external host can be defined per vRPA cluster.
- An SSH server must be installed on each external host.

Procedure

1. In the **The Start-up Sequence of VMs in this Group** pane, check **Run script**.
2. Type a logical name for the script.
3. Type the script command, including parameters (separated by a space).
4. Type the time-out period (mandatory).
5. Type the number of retries.

Automatic copy VM network configuration (Re-IP)

Use one of the following procedures to change a copy VM's network settings when *testing a copy*, *failing over*, or *recovering production*. Use the GUI to change the

network configuration of a small number of VMs, or use a comma-separated values (*.CSV) file to change the network configuration of many VMs in a copy or system.

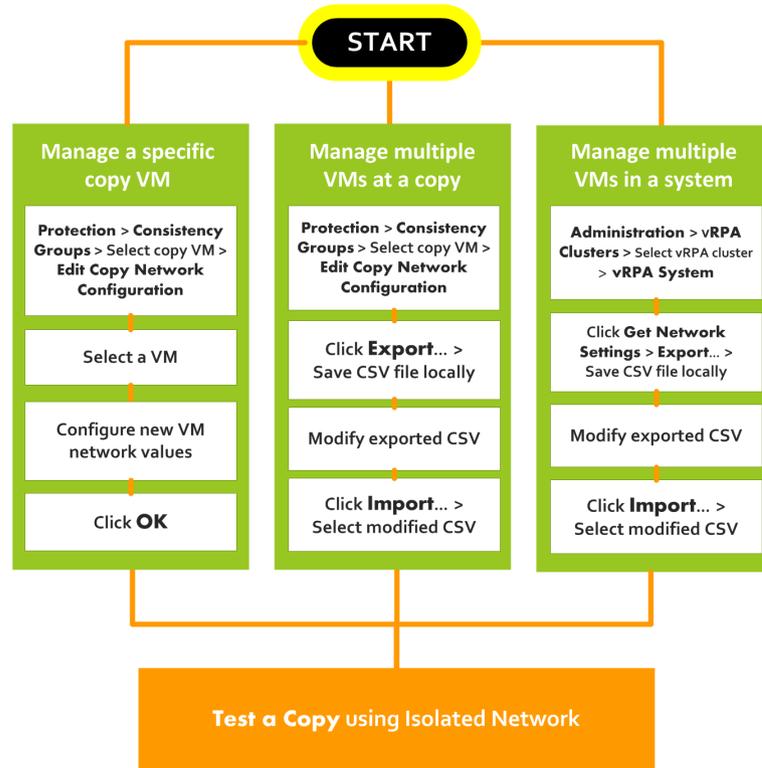
Before you begin

- Automatic network configuration is supported for VMs running MS Windows server versions 8, 10, 2008 R2, 2012, and 2016, Red Hat Linux server versions 6.5 and 7.2, and Ubuntu Studio 15.10.
- Best practice is to ensure that VMware Tools are installed on each relevant production VM.
 - For Linux CentOS 7.x, automatic network configuration is not supported unless VMware Tools version 10.1.0.57774 has been manually installed, and the value of each production VM's `ifconfig` version has been changed to 1.6 in the VM settings.
 - For Linux SLES12, automatic network configuration is not supported unless *Open VM Tools* version 9.4.0.25793 and `deployPkg` has been manually installed. See *VMWare KB article 2075048* for detailed information on how to install `deployPkg`.
 - For VMs running *Open VM Tools* versions lower than 9.10, automatic network configuration is not supported unless `deployPkg` has been manually installed. See *VMWare KB article 2075048* for detailed information on how to install `deployPkg`.
 - For all other operating systems, see [Manual copy VM network configuration \(RE-IP\)](#) on page 64
- By default, the **Network Configuration Method** is set to *Automatic*. Best practice is to leave this setting as is. However, if you want to use glue scripts (for example, because you have upgraded from a previous version of RecoverPoint for VMs and have already implemented glue scripts), ensure *Use glue scripts* is selected and see [Manual copy VM network configuration \(RE-IP\)](#) on page 64.

Note

If you are planning a temporary failover, to ensure that you don't lose your production VM network configuration when you fail back to the production, ensure that you edit the copy network configuration of your production VMs too.

The following diagram illustrates how to works in RecoverPoint for VMs version 5.0.1 and later. For how it works in previous versions, see [Manual copy VM network configuration \(RE-IP\)](#) on page 64.



Copy VM network configuration for a few virtual machines at a copy

Before you begin

You should be familiar with:

- [Testing a copy](#) on page 52
- [Failing over](#) on page 54
- [Recovering production](#) on page 56
- [Copy VM network configuration guidelines](#) on page 67

Procedure

1. Select **Protection > Consistency Groups**. Expand the consistency group, select the copy, and click the **Edit Copy Network Configuration** icon.
2. Select a virtual machine in the table.

To import the production VM setting value, select a setting in the GUI. When a relevant setting is selected, a **Get Value from Production** button is displayed next to it.

- To retrieve a specific setting value from the production VM. Click **Get Value from Production**.
 - To retrieve all relevant setting values from the production VM, click **Get All Values From Production**.
3. Type new network values for the copy VM.
 4. To apply the new values, click **OK**.
 5. Repeat for each virtual machine at the copy.

Results

The new copy VM network configuration is used when testing a copy, failing over, or recovering production.

Copy VM network configuration for many virtual machines at a copy

Before you begin

You should be familiar with:

- [Testing a copy](#) on page 52
- [Failing over](#) on page 54
- [Recovering production](#) on page 56
- [Copy VM network configuration guidelines](#) on page 67

Procedure

1. Select **Protection > Consistency Groups**. Expand the relevant consistency group, select the relevant copy, and click the **Edit Copy Network Configuration** icon.
2. To save the current network configuration of all virtual machines at the selected copy to a local *.csv file, click **Export....**
3. Open the CSV file and modify the network configuration of relevant virtual machines.
4. To apply the new network configuration, click **Import...** and select the modified CSV file .

Results

The new copy VM network configuration is used when testing a copy, failing over, or recovering production.

Copy VM network configuration for many virtual machines in a system

Before you begin

You should be familiar with:

- [Testing a copy](#) on page 52
- [Failing over](#) on page 54
- [Recovering production](#) on page 56
- [Copy VM network configuration guidelines](#) on page 67

Procedure

1. Select **Administration > vRPA Clusters > vRPA System**, and select the vRPA of the relevant RecoverPoint for VMs system.
2. To retrieve the network configuration of all copy VMs at all vRPA clusters in the system, in the **Network Configuration** section, click **Get Network Settings**.
3. To save the current network configuration to a local CSV file, click **Export....**
4. Open the CSV file and modify the network configuration of relevant virtual machines.
5. To apply the new network configuration to the system, click **Import...** and select the modified CSV file.

Results

The new copy VM network configuration is used when testing a copy, failing over, or recovering production.

Automation

This section describes the RecoverPoint for VMs features for automating the protection of virtual machines and VMDKs.

Replicating VM hardware changes

By default, **Replicate hardware changes** is `Enabled` and the production VM version, MAC address, CPU, memory, resource reservations, and network adapter status and type are all replicated to any corresponding copy VMs whenever an image on the copy VM is accessed. Use this procedure to change the default behavior.

Note

Replication of the SR-IOV network adapter type is not supported. If the ESX at a copy does not support the production VM version, no hardware resources are replicated.

Procedure

1. Select **Protection > Virtual Machines**.
2. Select a copy VM.
3. In the **Hardware Settings** widget, click **Edit...** and uncheck **Replicate hardware changes**.

VMDK provisioning**Procedure**

1. Select **Protection > Virtual Machines**.
2. Select the virtual machine.
3. In the **Hardware Settings** widget, click **Edit...**
4. In the **Disk provisioning** drop-down, select `Same as source`, `Thick provisioning`, or `Thin provisioning`.

Results

By default, VMDKs are provisioned `Same as source`. When `Thick provisioning` or `Same as source` are used, if the production VMDKs are thick (either eager or lazy), the copy VMDKs are thick provisioned lazy zero.

Enabling and disabling automatic protection of new VMDKs**Procedure**

1. To enable or disable the automatic protection of newly included VMDKs, browse to one of the following:
 - **Protection > Virtual Machines screen > Protected VMDKs widget.**
 - In the vCenter Server inventory, **Summary > RecoverPoint for VMs widget > Protected VMDK(s) section.**
2. Click **Edit...** and mark or clear the **Automatically protect new VMDKs** checkbox.

Removing a VMDK

VMDKs are removed in vSphere client when a hard disk is removed from a virtual machine using the **Virtual Machine Properties**.

- Removing VMDKs from the production does not delete their copies.
- After removing a VMDK from the production, do one of the following:
 - Add the missing VMDK to the production VM with the same port type, ID, and size as the copy VMDK. This action causes a volume sweep on the added VMDK and a short initialization on all other VMDKs in the group.
 - Exclude these VMDKs from replication.

Excluding a VMDK from replication

If required, you can mark individual VMDKs for exclusion from replication. For example, virtual machines containing shared or non-persistent VMDKs cannot be replicated. You can, however, change the VMDK type, or use this feature to mark those VMDKs to be excluded from replication and replicate the virtual machines without them.

Before you begin

- Changing the disk type of an excluded shared or non-persistent VMDK to a supported type (such as non-shared or persistent) does not automatically include the VMDK, regardless of the value of the **Automatically protect new VMDKs** setting.
- The excluded production VMDKs are not replicated, but the corresponding copy VMDKs are not deleted.
- If you remove VMDKs before protecting a virtual machine, the VMDK copies are not created.
- If there is no connectivity between the vCenter Server and the vRPA cluster, some VMDK information, such as size, type, and excluded VMDKs is lost.

Procedure

1. To exclude VMDKs from replication, browse to one of the following:
 - **Protection > Virtual Machines > Protected VMDKs**
 - In the vCenter Server inventory: **Summary > RecoverPoint for VMs > Protected VMDK(s)**
2. Click **Edit...** and clear the checkbox next to the VMDKs you want to exclude from replication.

Results

The excluded VMDKs are deleted from the system and are not recoverable when accessing any point in time, even one previous to the VMDK removal.

Including a VMDK in replication

Procedure

1. Browse to:
 - **Protection > Virtual Machines > Protected VMDKs**
 - vCenter Server inventory **Summary > RecoverPoint for VMs > Protected VMDK(s)**

2. Click **Edit...** and mark the checkbox next to the VMDKs you want to include.

Results

A volume sweep starts on the included VMDK, and a short initialization starts on all other VMDKs in the group.

Adding a VMDK

VMDKs are added in vSphere client when a hard disk is added to a virtual machine using the **Virtual Machine Properties**.

- When adding VMDKs to the production, RecoverPoint automatically creates the relevant copy VMDKs.
- If you add a VMDK of type `shared` to a production VM, RecoverPoint does not automatically replicate the VMDK. Manually change each copy VMDKs type back to `shared` after VMDK addition. Otherwise, group transfer is paused and an alert is displayed in the system **Dashboard**.
- After a VMDK is added, a volume sweep occurs on the added VMDK, and a short initialization occurs on all other VMDKs in the group.

Automatically expanding copy VMDKs

When you use VMware to expand a production VMDK, RecoverPoint for VMs automatically expands all corresponding copy VMDKs, with the following limitations:

- VMDKs can be expanded, but they cannot be shrunk.
- When a production VMDK is expanded, the system pauses replication of the consistency group while the system is busy resizing the corresponding copy VMDK.
- Automatic VMDK expansion fails if:
 - The datastore does not contain enough free space. In this case, free up space in the copy VM datastore.
 - A snapshot has been taken of the virtual machine containing the copy VMDK. In this case, enable image access to the copy VM containing the VMDK and then use the vCenter snapshot manager to delete all snapshots before disabling image access.
 - The version of the file system that you are running does not support the VMDK size. In this case, consider upgrading the file system version.

After fixing any of these issues, wait 15 minutes for the automatic expansion process to restart and the error to resolve itself. If the problem persists, try manually resizing the copy VMDKs or contact Customer Support.

- The system pauses replication of the consistency group if:
 - The user accesses a copy containing a VMDK marked for automatic expansion.
 - A production VMDK is smaller than the size registered in the system settings (because the production VMDK has been removed and re-added with a smaller size). Ensure that the size of all of the VMDKs in the consistency group is the same. If problem persists, contact Customer Support.
 - One or more copy VMDKs has been marked for automatic expansion, but the system cannot automatically resize a raw device. In this case, enable image access to the copy VM with the problematic VMDK and manually expand it before disabling image access. If problem persists, contact Customer Support.

- If the size of a copy VMDK is larger than the size of its corresponding production VMDK, to begin the automatic VMDK expansion process, you must manually expand the production VMDK. This manual expansion might be required if you failed over while automatic expansion was in progress, or if the copy VMDK was manually expanded.

Adding a virtual machine to a consistency group

When adding a virtual machine to an existing consistency group, a volume sweep occurs on the newly added virtual machine and a short initialization on all other existing virtual machines in the consistency group. When there are three copies or more in a group, there is journal loss to the non-production copies on failback to the original production.

CHAPTER 4

Testing and recovery

This chapter includes procedures for testing, failing over, and recovering your data.

- [Testing and recovery](#)..... 52
- [Testing a copy](#).....52
- [Failing over](#).....54
- [Recovering production](#).....56

Testing and recovery

The following testing and recovery activities are available in RecoverPoint for VMs:

- [Testing a copy](#) on page 52
- [Recovering production](#) on page 56
- [Failing over](#) on page 54

The **Recovery Wizard** directs these recovery activities the on consistency groups or group sets. The initial steps of the wizard, including selecting the image to access and testing the network, are the same for all recovery activities. The actual recovery activities are performed at the end of the wizard, after testing the copy.

Before performing recovery activities on a group set, note that during image access on a group set:

- When you select the latest image, the latest available image of each relevant group copy is accessed.
- When you select an image from the image list, the images of the first group in the group set are displayed. When you select an image from the list, RecoverPoint for VMs constructs a search query containing the parameters of the first group's image. RecoverPoint for VMs then queries the rest of the groups in the group set according to those parameters.
- When you select a specific point in time by entering advanced search options, the same image search query is sent for all groups in a group set.

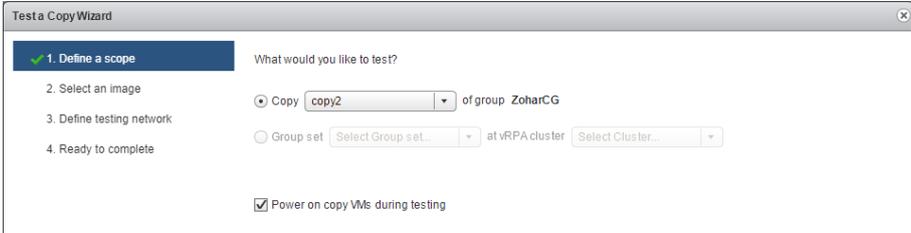
Note

In environments containing multiple RecoverPoint for VMs systems, to lessen the load on back-end storage arrays, best practice is to fail over the consistency groups of up to seven systems concurrently.

Testing a copy

Procedure

1. Select the **Protection** tab and click the **Test Copy** icon  to open the **Test a Copy Wizard**.
2. In the **Define a scope** screen, select whether you want to test the consistency group or the group set.

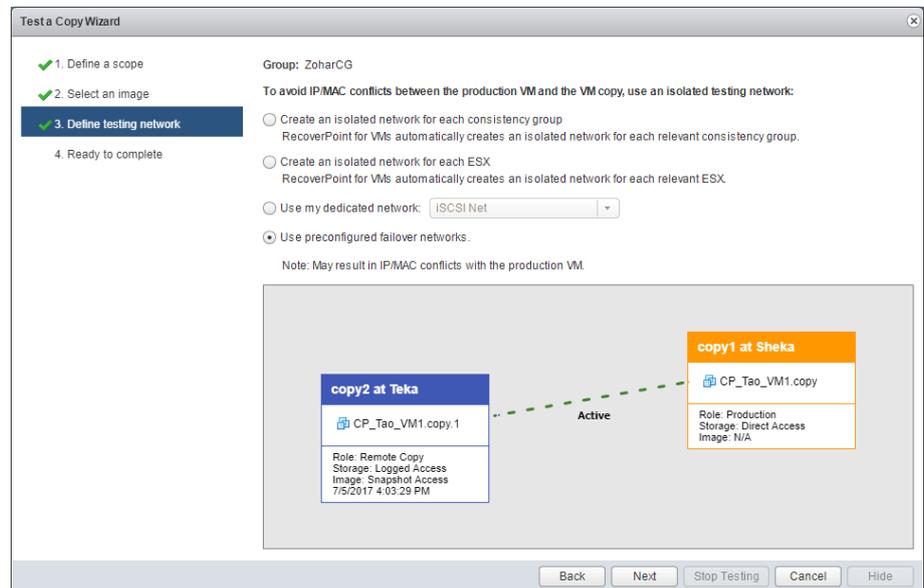


Uncheck the **Power on copy VMs during testing** checkbox if you want the copy VMs powered off during image access. The default, checked, powers on the virtual machine during image access. Unselecting this option skips the post-power up steps (including copy VM network reconfiguration) in the start-up sequence.

3. In the **Select an image** screen, select the image to access. You may want to start with the last image that is known to be valid. You can select:

Option	Description
The latest image	The last snapshot that was created at the production, and transferred to the copy journal.
An image from the image list	An image from the list. The number of snapshots available in the image list is limited. Snapshots that are not in the image list may still be selected by specifying a Point in Time .
A specific point in time or bookmark	Allows you to perform a customized search: <ul style="list-style-type: none"> • Point in Time - Searches for a bookmark that was created at the specified date and time. • Max Range - Searches for a bookmark that was created between the specified number of minutes/hours before and after the specified Point in Time. • Bookmark - Searches for bookmarks with the specified text in the bookmark name. • Exact - Searches for bookmarks that contain the exact text that was entered in the Bookmark field. • Image Type - Searches for the specified image type with the specified bookmark name.

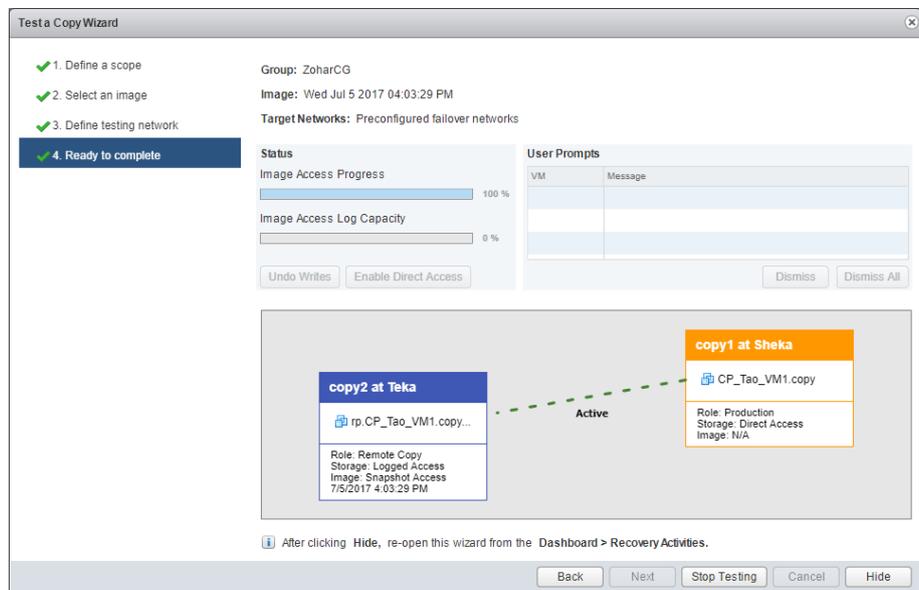
4. In the **Define testing network** screen, to avoid IP conflicts between the production and copy VMs, best practice is to use a dedicated testing network.



Option	Description
Create an isolated network for each consistency group	RecoverPoint for VMs auto-provisions an isolated network for virtual machines in this consistency group or group set to avoid IP conflicts between the production VMs and the tested virtual machine.

Option	Description
Create an isolated network for each ESX	RecoverPoint for VMs automatically creates an isolated network for each ESX splitter.
Use my dedicated network	Manually select a configured network. Not relevant for group sets.
Use preconfigured failover networks	RecoverPoint for VMs uses the configured failover networks for each copy VM. To view or edit failover networks, see Configuring copy VM failover networks on page 40. Not relevant for <i>group sets</i> .

- In the **Ready to complete** screen, verify that the displayed image access details are correct.



The **Image Access Progress Bar** indicates the progress of image access. After image access is enabled, the **Image Access Log Capacity** progress bar indicates how long you can access the copy image before the image access log is full and all writes to the copy fail.

- To exit the **Test a Copy** wizard and start testing the image, click **Hide**.
- When testing is complete, re-open the **Test a Copy** wizard to disable access to the copy image. Select the **Recovery Activities** widget in the system **Dashboard**, select the activity and click **Next Action > Back to Wizard**. In the **Test a Copy** wizard, click **Stop Testing**.

Failing over

Guides you through the process of selecting a copy image, testing it, and failing over to the image at the copy or failing back to the production.

Before you begin

- The **Failover Wizard** screens contain the following options:

- **Hide:** Keeps access to the image enabled at the specified copies and exits the wizard.

Note

After clicking **Hide**, you can re-open the wizard through the **Recovery Activities** widget in the system **Dashboard**, by selecting **Back to Wizard** in the relevant recovery activity bar.

- **Cancel:** Disables access to the image at the specified copies and exits the wizard.
- **Fail Over:** Starts failing over to the image at the specified copies.
- After failover, the production and copy VM change roles, but the names do not change. Therefore, after failover, the new production VM will still be *YourVMName.copy* and the new copy VM name is still named *YourVMName*.
- The marking information in the production journal is deleted, the copy journal is deleted, and the consistency group undergoes a full sweep synchronization.
- Before failing back to the production, use the Recovery wizard to select an image at the production that predates the failover. You can also verify the image before permanently selecting it as the image you want to fail back to.

Procedure

1. Select **Protection tab > Fail Over** icon:



The **Failover Wizard** is displayed.

2. In the **Define a Scope** screen, select whether you want to test the consistency group or the group set. If there are no group sets, the option is grayed out.
3. In the **Select Image** screen, select the image to access. You may want to start with the last image that is known to be valid.

Option	Description
Current image	The current image, as displayed in the wizard.
The latest image	The last snapshot that was created at the production, and transferred to the copy journal.
An image from the image list	<p>Select an image from the list.</p> <ul style="list-style-type: none"> • The number of snapshots available in the image list is limited. You can still select snapshots that are not in the image list by specific Point in Time. • During snapshot dilution, priority is given to bookmarked images.
A specific point in time or bookmark	<p>Allows you to perform a customized search.</p> <ul style="list-style-type: none"> • Point in Time - Searches for a bookmark that was created at the specified date and time. • Max Range - Searches for a bookmark that was created between the specified number of minutes/hours before and after the specified Point in Time.

Option	Description
	<ul style="list-style-type: none"> • Bookmark - Searches for bookmarks with the specified text in the bookmark name. • Exact - Searches for bookmarks that contain the exact text that was entered in the Bookmark field. • Image Type - Searches for the specified image type with the specified bookmark name.

4. In the **Define testing network** screen, to avoid IP conflicts between the production and copy VMs, best practice is to use a dedicated testing network.

Option	Description
Create an isolated network for each group	RecoverPoint for VMs auto-provisions an isolated network for virtual machines in this group or group set to avoid IP conflicts between the production VM and the tested virtual machine.
Create an isolated network for each ESX	RecoverPoint for VMs automatically creates an isolated network for each ESX splitter.
Use my dedicated network	Manually select a preconfigured network. Not relevant for <i>group sets</i> .
Use preconfigured failover networks	RecoverPoint for VMs uses the preconfigured failover networks for each copy VM. To view or edit failover networks, see Configuring copy VM failover networks on page 40. Not relevant for <i>group sets</i> .

5. In the **Ready to complete** screen:
- To ensure that the failover is configured correctly, review the displayed summary information.
 - To define the **Target Networks** for failover, click **Edit....**
 - Once image access is enabled, click **Fail Over** to start failover.

Recovering production

Corrects file or logical corruption by rolling the production back to a previous point-in-time. Guides you through the process of selecting a copy image, testing it, and recovering the production from the selected image.

Before you begin

The **Recover Production Wizard** screens contain the following options:

- **Hide:** Keeps access enabled to the image at the specified copy(s) and exits the wizard.

Note

After clicking **Hide**, you can re-open the wizard through the **Recovery Activities** widget in the system **Dashboard**, by selecting **Back to Wizard** in the relevant recovery activity bar.

- **Cancel:** Disables access to the image at the specified copy(s) and exits the wizard.
- **Recover Production:** Starts failing over to the image at the specified copy(s).

Procedure

1. To recover production, in the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon > Protection** tab. Click the **Recover Production** icon:



The Recovery Wizard appears.

2. In the **Define a Scope** screen, select whether you want to test the consistency group or the group set. If there are no group sets, the option is grayed out.
3. In the **Select an Image** screen, select the image to access. You may want to start with the last image known to be valid.

Option	Description
The latest image	The last snapshot that was created at the production, and transferred to the copy journal.
An image from the image list	<p>Select an image from the list.</p> <ul style="list-style-type: none"> • The number of snapshots available in the image list is limited. Snapshots that are not in the image list may still be selected by specific <i>Point in Time</i> • During snapshot dilution, priority is given to bookmarked images.
A specific point in time or bookmark	<p>Allows you to perform a customized search.</p> <ul style="list-style-type: none"> • Point in Time - Searches for a bookmark that was created at the specified date and time. • Max Range - Searches for a bookmark that was created between the specified number of minutes/hours before and after the specified Point in Time. • Bookmark - Searches for bookmarks with the specified text in the bookmark name. • Exact - Searches for bookmarks that contains the exact text entered in the Bookmark field. • Image Type - Searches for the specified image type with the specified bookmark name.

4. In the **Define testing network** screen, define the testing environment by specifying Testing Network options. Best practice to avoid IP conflicts between the production VM and the copy VM, is to use a dedicated testing network

Option	Description
Create an isolated network for each group	RecoverPoint for VMs auto-provisions an isolated network for virtual machines in this consistency group or group set to avoid IP conflicts between

Option	Description
	the production VMs and the tested virtual machines.
Create an isolated network for each ESX	RecoverPoint for VMs automatically creates an isolated network for each ESX splitter.
Use my dedicated network	Manually select a preconfigured network. Not relevant for <i>group sets</i> .
Use preconfigured failover networks	RecoverPoint for VMs uses the preconfigured failover networks for each copy VM. To view or edit failover networks, see Configuring copy VM failover networks on page 40. Not relevant for <i>group sets</i> .

5. In the **Verify image to access** screen, verify that the image access details displayed are correct, and click **Next**.
6. In the **Ready to complete** screen, detailed information about the selected copy is displayed.

The **Image Access Progress** bar will indicate the progress of image access. You can close the wizard without interfering with the process. You can reopen the wizard from the **Recovery Activities** widget on the **Dashboard**. After image access is enabled, the buffer progress bar indicates how long you can access the copy image before the image access log is full and all writes to the copy fail. Once image access is complete, click **Recover Production**. During production recovery, host access to storage is blocked.

Results

- The marking information in the production journal is deleted, the copy journal is deleted, and the consistency group undergoes a full sweep synchronization.
- The group undergoes a short initialization process to synchronize the new production data at the copy.

CHAPTER 5

Troubleshooting

This chapter provides procedures for mitigating issues that may arise when using RecoverPoint for VMs.

- [Finding the vRPA cluster management IP](#)..... 60
- [Identifying a RecoverPoint for VMs system](#)..... 60
- [Recovering from a cluster disaster](#)..... 60
- [Detecting bottlenecks](#)..... 60
- [Load balancing](#)..... 62
- [System alerts](#)..... 63
- [Collecting system information](#)..... 63
- [Collecting RecoverPoint for VMs splitter logs](#)..... 64
- [Manual copy VM network configuration \(RE-IP\)](#)..... 64
- [Copy VM network configuration guidelines](#)..... 67
- [Changing the network adapter configuration of a protected virtual machine](#)..... 72

Finding the vRPA cluster management IP

Displays the vRPA cluster management IP of a specific vRPA cluster.

Procedure

1. Select **Administration > vRPA Clusters > vRPA System**
2. Select the vRPA cluster.
3. Note the **vRPA cluster management IP** of the selected vRPA cluster.

Identifying a RecoverPoint for VMs system

When a vRPA cluster is selected, the GUI displays all other vRPA clusters (besides the selected one) that constitute a RecoverPoint for VMs system.

Procedure

1. Select **Administration > vRPA Clusters > vRPA System**
2. Select a vRPA cluster.
3. Note the value of **Other vRPA clusters in system**.

Recovering from a cluster disaster

After a full cluster disaster or a switch disaster, it may take 10 minutes or more for all the components of the vRPA system to restart, reconnect, and restore full operation.

Detecting bottlenecks

Bottleneck detection returns statistics about RecoverPoint for VMs system performance, by consistency group, vRPA, and vRPA cluster. Bottleneck detection analyzes the system data to detect the existence of any of the predefined problem types called *bottlenecks*. The types of bottlenecks are presented in [Table 3](#) on page 60 and [Table 4](#) on page 61.

Procedure

1. To detect bottlenecks, use an `ssh` client to connect to the vRPA management IP address, and type the RecoverPoint *username* and *password* to log in to the CLI.
2. To check for any bottlenecks, run the `detect_bottlenecks` command. To view command parameters that can refine the search, run:
`detect_bottlenecks ?`

Table 3 RPA and cluster bottlenecks

Bottleneck type detected	System output and comments
RPA balance	<pre>RPAs are not balanced.</pre> <p>Includes data on the load that is handled by each vRPA at the vRPA cluster.</p>

Table 3 RPA and cluster bottlenecks (continued)

Bottleneck type detected	System output and comments
	<p>Note</p> <p>vRPA balance is checked only if the time period defined is > 30 minutes.</p>
Compression	<p>Compression level is too high. The RPA resources cannot handle the current level.</p>
SAN target	<p>RPA may be regulating the application. Consider reducing RPA load.</p> <p>Includes data on the total amount of incoming data, the number of writes, and the amount of incoming data per write.</p>
RPA utilization	<p>RPA utilization reached ##%.</p>

Table 4 Consistency group and link bottlenecks

Bottleneck type detected	System output and comments
Slow production journal	<p>Writing to the local journal volume was slow during this period.</p> <p>Includes data on the delay factor.</p>
Journal phase 1	<p>Journal is unable to handle the incoming data rate.</p> <p>Includes the required I/O rates for the journal and the replication volumes at local or remote copies, for both normal and fast-forward distribution modes.</p>
Journal phase 2	<p>Journal and replication volumes are unable to handle the incoming data rate.</p> <p>Includes data on the required I/O rates for the journal and the replication volumes at local or remote copies, for both normal and fast-forward distribution modes.</p>
Journal regulation	<p>Remote storage is too slow to handle incoming data rate and regulate the distribution process.</p> <p>Includes data on the required I/O rates for the journal and the replication volumes at local or remote copies, for both normal and fast-forward distribution modes.</p>
Unknown distribution problem	<p>Target cluster cannot handle the incoming data rate.</p>

Table 4 Consistency group and link bottlenecks (continued)

Bottleneck type detected	System output and comments
	Includes data on the required I/O rates for the journal and the replication volumes at local or remote copies, for both normal and fast-forward distribution modes.
Slow WAN	<p>WAN is too slow.</p> <p>Includes data on total throughput for the vRPA cluster, the identity of the RPAs at which the problem appeared, and the throughput of that RPA (or RPAs). A slow WAN bottleneck is detected based on the group, but generates data based on the vRPA cluster and vRPA.</p>
Slow read source	<p>Reading rate from the source replication volume(s) during synchronization is too slow.</p> <p>Includes the reading rate.</p>
Link utilization	Link utilization reached ##%.

Results

The output from the system analysis is written to `/home/kos/statistics/bottlenecks.csv`.

Load balancing

Load balancing is the process of assigning preferred vRPAs to consistency groups so that the preferred vRPA performs data transfer for that group. This is done to balance the load across the system and to prevent the system from entering a high-load state.

Perform load balancing:

- When a new consistency group is added to the system. Wait 1 week after the new group is added to accumulate enough traffic history before performing load balancing.
- When a new vRPA is added to a vRPA cluster. Perform load balancing immediately after the vRPA is added.
- If the system enters high load frequently. When load balancing is required, the event logs display a message indicating so. When you see this message, perform load balancing.
- If the bottleneck detection tool recommends it. When load balancing is required, the `detect_bottlenecks` CLI command returns "RPAs are not balanced." When you see this message, perform load balancing.
- Periodically, to ensure that the system is always handling distributing loads evenly. A script can be created to periodically perform load balancing.

Procedure

1. To balance the load on the vRPAs, use an `ssh` client to connect to the vRPA management IP address, and type the RecoverPoint `username` and `password` to log in to the CLI.

2. Run the `balance_load` command to balance the load. To view command parameters that can refine the search, run: `balance_load ?`

System alerts

To view system errors and warnings, In the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon > Dashboard tab > Alerts** widget.

Collecting system information

Collects system information for support purposes.

Before you begin

- This procedure is only relevant in support cases, and should only be performed when instructed to do so by Customer Support.
- Errors occur in the following cases:
 - If connection with a vRPA is lost while info collection is in process, no information is collected. In this case, run the process again. If the collection from the remote site failed because of a WAN failure, run the process locally at the remote site.
 - If a simultaneous info collection process is being performed on the same vRPA, only the collector that established the first connection can succeed.
 - If an FTP failure occurs, the entire process fails.

Procedure

1. In the vSphere Web Client home page, select **Administration > vRPA Clusters > Log Collection**.
2. Under **Collection Period**, define a date and time for the start and end of the collection process.
3. Optionally, click **Change to GMT** to change the collection time display to GMT.

Note

GMT is not adjusted for daylight savings time. Although the system information of the past 30 days is available for collection, only 3 days of system information can be collected at a time.

4. Optionally, select **Include core files**.

Note

Core files might be large. Subsequently, including these files in the collection process could substantially increase collection time.

5. By default, **Full system log collection** is selected. If you are instructed to do so by Customer Support, use **Advanced** to select the specific logs that you want to collect.
6. Optionally, select **Copy the output file(s) to an FTP server** and define the FTP server settings.
7. Click **Start**.

Note

Be patient. The collection process can take awhile, depending on the amount of data being collected.

Results

After the collection process is complete, the results are displayed.

After you finish

If you selected the **Copy the output file(s) to an FTP server** checkbox, retrieve the output file from the specified FTP server. Otherwise, retrieve the files from the local vRPA cluster by clicking the relevant link in the **Output File (HTTPS)** column.

1. At the login prompt, type `admin` as both the **User Name** and the **Password**.
2. Right-click the file and select **Save link as...** to download the file to the local virtual machine.
3. Open the file using a data compression utility.

Collecting RecoverPoint for VMs splitter logs

RecoverPoint for VMs splitter logs are in the ESXi logs. To export the ESXi system logs, use the following procedure.

Procedure

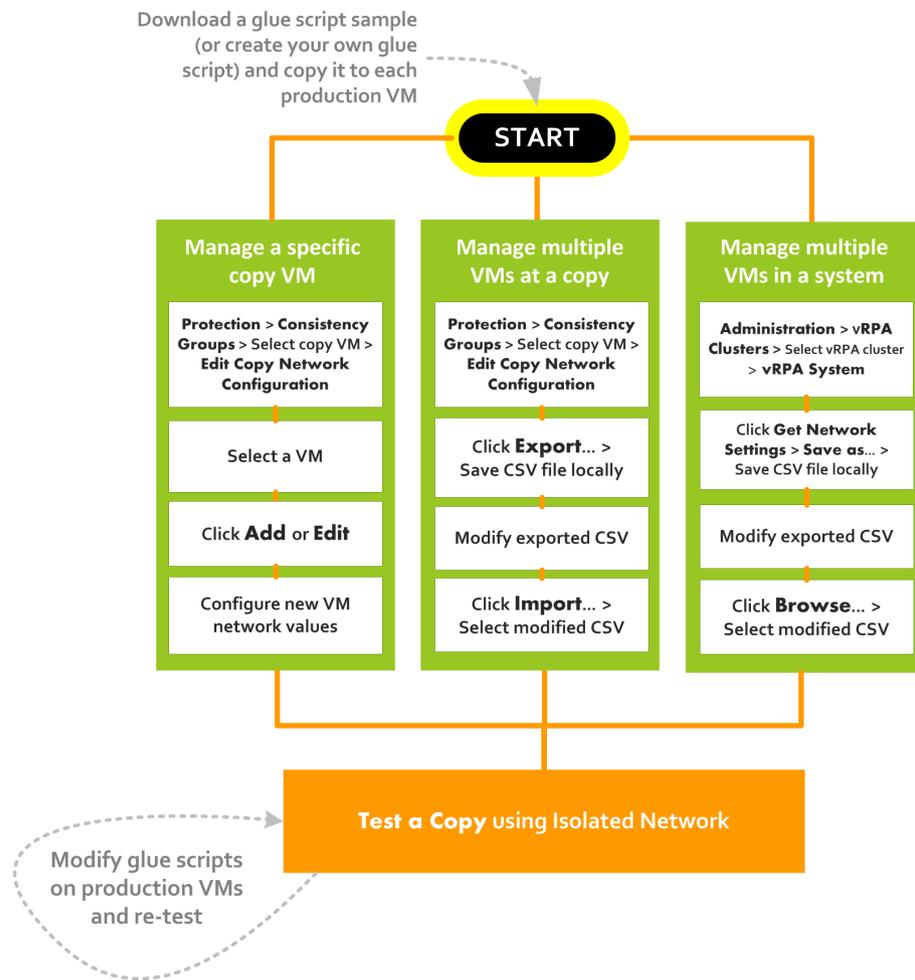
1. In the vSphere Web Client, select an ESXi host and click **Actions**.
2. Select **All vCenter Actions > Export System Logs...**
3. In the **Export Logs** screen, specify which system logs are to be exported. If required, select the `Gather performance data` option and specify a `duration` and `interval`.
4. Click **Generate Log Bundle**.
5. Click **Download Log Bundle**.
6. Upload the logs to the SFTP/FTP site.

For information on how to upload logs for VMware products, see http://kb.vmware.com/selfservice/search.do?cmd=displayKC&docType=kc&docTypeID=DT_KB_1_1&externalId=1008525

Manual copy VM network configuration (RE-IP)

Manually configure your copy VMs if you have upgraded from RecoverPoint for VMs 5.0.1 or earlier and are already using glue scripts, or when your VMs operating system does not support [Automatic copy VM network configuration \(Re-IP\)](#) on page 43.

The following diagram illustrates how it works:



Note

VMware Tools must be installed on a virtual machine copy's production VM for automatic virtual machine network re-configuration. For virtual machines running *Open VM Tools* versions lower than 9.10, network configuration is not supported unless `dep1oyPkg` has been manually installed. See *VMware KB article 2075048* for detailed information on how to install `dep1oyPkg`.

To manage a copy VM network configuration:

1. Create glue scripts or download the relevant glue script samples from <https://download.emc.com/downloads/DL66792>.

Table 5 Glue script samples

Name	Language	Target OS	Capabilities	Prerequisites
glue_script_win.bat	Windows batch	Microsoft Windows 2008 and 2012	Modification of IPv4, Subnet Mask, Gateway	<ul style="list-style-type: none"> VMware Tools that are installed on each protected VM

Table 5 Glue script samples (continued)

Name	Language	Target OS	Capabilities	Prerequisites
				<ul style="list-style-type: none"> Rename <code>glue_script_w_in_this_bat_file.txt</code> to <code>glue_script_w_in.bat</code>
<code>glue_script_win.py</code>	Python	Microsoft Windows 2008 and 2012	Modification of IPv4, IPv6, Subnet Mask, Gateway, DNS servers, DNS Suffix	<ul style="list-style-type: none"> VMware Tools that are installed on each protected virtual machine Python 2.7
<code>glue_script_rhel.bash</code>	BASH	LINUX	Modification of IPv4, Subnet Mask, Gateway	<ul style="list-style-type: none"> VMware Tools that are installed on each protected virtual machine DNS Server and Suffix are only applied in Win 2008
<code>glue_script_rhel.py</code>	Python	LINUX	Modification of IPv4, IPv6, Subnet Mask, Gateway, DNS servers, DNS Suffix	<ul style="list-style-type: none"> VMware Tools that are installed on each protected virtual machine Python 2.7

2. Copy the relevant glue scripts to the relevant production VM:

- For virtual machines running a Windows operating system:
 - Place the glue script in a directory that is accessible by all authorized users.
 - To configure the script to run on startup, open the Windows **Task Scheduler** and select **Action > Create Basic Task...**. Select **When the computer starts** as the task trigger and **Start a program** as the task action. Select the glue script, set the **Start in** directory to the directory where you want to place the glue script output, select **Open the Properties dialog for this task when I click Finish** and finish creating the task. In the **Properties** dialog box that is displayed, select **Run with highest privileges**, click **Change User or Group...**, type `SYSTEM` in the **Object name to select** field, and click **OK**.
 - To configure the glue script to run on login, add the glue script path to the registry by creating a string that is called `IP` in the `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Run` registry path containing the full path to the script. For example: `C:\my_directory\glue_script_win.bat`.

- For virtual machines running a Linux operating system, add the relevant glue script to the `rc.local` file under `/etc/rc.d`.
3. Perform [Automatic copy VM network configuration \(Re-IP\)](#) on page 43.
 4. Customize the glue scripts on the production VMs until they run correctly on the copy VMs.

Copy VM network configuration guidelines

Use the following guidelines for:

- [Automatic copy VM network configuration \(Re-IP\)](#) on page 43
- [Manual copy VM network configuration \(RE-IP\)](#) on page 64

Table 6 Virtual machine network settings available through the GUI

Setting	Description	Guidelines
(VM) Operating System	The guest operating system of the specified VM.	<ul style="list-style-type: none"> • Not customizable. • Automatically populated by the system. • Possible values are <i>Windows</i>, <i>Linux</i>, or <i>Unknown</i>.
(VM) Host Name	The hostname of the specified VM.	<ul style="list-style-type: none"> • Only mandatory for virtual machines with a Linux operating system. • Customizable. • Value can be retrieved from production VM using Get Value from Production or Get All Values from Production.
(VM) DNS Domain	The DNS domain for the specified VM.	<ul style="list-style-type: none"> • Only relevant (and mandatory) for virtual machines with a Linux operating system. • Value should be in the format <code>example.company.com</code>.
(VM) DNS Server(s)	The global IP address that identifies one or more DNS servers for all adapters of the specified VM.	<ul style="list-style-type: none"> • Only relevant for virtual machines with a Linux operating system. • Customizable. • Can be left blank. • This setting applies to all virtual network adapters of the specified VM. • Separate multiple values with a semicolon (;). • Value can be retrieved from production VM using Get Value from Production or Get All Values from Production.
(VM) DNS Suffix(s)	The global settings of the suffixes for the DNS servers of all adapters on both	<ul style="list-style-type: none"> • Customizable. • Can be left blank.

Table 6 Virtual machine network settings available through the GUI (continued)

Setting	Description	Guidelines
	Windows and Linux virtual machines.	<ul style="list-style-type: none"> Value can be retrieved from production VM using Get Value from Production or Get All Values from Production.
(VM) Network Configuration Method	The method by which the system configures new virtual network adapters (NICs) for the specified virtual machines.	<ul style="list-style-type: none"> Possible values are Automatic and Use glue scripts. Default is Automatic. When Automatic, RecoverPoint for VMs automatically configures the virtual machine copy network. Best practice is to leave this setting as is. However, to use glue scripts (for example, because you upgraded RecoverPoint for VMs and already implemented glue scripts), ensure Use glue scripts is selected and follow the instructions for Manual copy VM network configuration (RE-IP) on page 64.
Adapter ID	ID of the virtual network adapter to customize.	<ul style="list-style-type: none"> When Network Configuration Method is set to <i>Automatic</i>, this value is not required, and should be left blank. Any entered value is ignored by the system. Mandatory when Network Configuration Method is set to <i>Use glue scripts</i>. To find this value: <ul style="list-style-type: none"> Windows - Type the interface index, which can be found by running <code>route print</code>. The adapter ID should be set according to the <code>IDX</code> value that is determined by running <code>NetSh Interface IPv4 Show Interfaces</code> on the Windows computer and determining the correct adapter. Linux - The adapter ID should be set according to the Ethernet port value. Type the sequential number (1-based) of the adapter, and NOT the NIC number. For example, <code>eth0 = 1</code>, <code>eth1 = 2</code>. If you have <code>eth2</code> and <code>eth3</code>, and want to update the network settings of the second one, set <code>Adapter ID = 2</code>.

Table 6 Virtual machine network settings available through the GUI (continued)

Setting	Description	Guidelines
(Adapter) IP Address	IPv4 address for this virtual network adapter.	<ul style="list-style-type: none"> • Can contain either a static IPv4 address or DHCP string. • Can be left blank when using IPv6. • Define one IPv4 address, one IPv6 address, or one of each, for the same virtual network adapter. Entering multiple IPv4 or IPv6 addresses for the same virtual network adapter is not supported. • Value can be retrieved from production VM using Get Value from Production or Get All Values from Production.
(Adapter) Subnet	IPv4 subnet mask for this virtual network adapter.	<ul style="list-style-type: none"> • Mandatory when an IP Address is entered. • Can be left blank when using IPv6. • Value can be retrieved from production VM using Get Value from Production or Get All Values from Production.
(Adapter) Gateway(s)	One or more IPv4 gateways for this virtual network adapter.	<ul style="list-style-type: none"> • Mandatory when an IP Address is entered. • Can be left blank when using IPv6. • Separate multiple values with a semicolon (;). • Value can be retrieved from production VM using Get Value from Production or Get All Values from Production.
(Adapter) IPv6 Address	IPv6 address for this virtual network adapter.	<ul style="list-style-type: none"> • Can contain either a static IPv6 address or it's DHCP string. • Can be left blank when using IPv4. • Define one IPv4 address, one IPv6 address, or one of each, for the same virtual network adapter. Entering multiple IPv4 or IPv6 addresses for the same virtual network adapter is not supported. • Value can be retrieved from production VM using Get Value from Production or Get All Values from Production.
(Adapter) IPv6 Subnet Prefix Length	IPv6 subnet mask for this virtual network adapter.	<ul style="list-style-type: none"> • Customizable. • Can be left blank when using IPv4. • Value can be retrieved from production VM using Get Value from Production or Get All Values from Production.

Table 6 Virtual machine network settings available through the GUI (continued)

Setting	Description	Guidelines
(Adapter) IPv6 Gateway(s)	One or more IPv6 gateways for this virtual network adapter.	<ul style="list-style-type: none"> • Customizable. • Mandatory when an IPv6 format IP Address is entered. • Can be left blank when using IPv4. • Separate multiple values with a semicolon (;). • Value can be retrieved from production VM using Get Value from Production or Get All Values from Production.
(Adapter) DNS Server(s)	IP address of one or more DNS server(s) for this virtual network adapter.	<ul style="list-style-type: none"> • Can be left blank. • Can contain one or more IPv4 DNS servers for each virtual network adapter (NIC). • Applies only to the configured adapter when a value other than Adapter ID 0 is defined. • Separate multiple values with a semicolon (;). • Value can be retrieved from production VM using Get Value from Production or Get All Values from Production.
(Adapter) NetBIOS	Whether or not to activate NetBIOS on this virtual network adapter.	<ul style="list-style-type: none"> • Cannot be left blank. • Only relevant for virtual machines running a Windows operating system. • Default is Enabled. • Net BIOS should be enabled. • Valid values are DISABLED, ENABLED, ENABLED_VIA_DHCP.
(Adapter) Primary WINS	Primary WINS server of this virtual network adapter.	<ul style="list-style-type: none"> • Relevant for windows virtual machines only. • Customizable. • Can be left blank.
(Adapter) Secondary WINS	Secondary WINS server of this virtual network adapter.	<ul style="list-style-type: none"> • Relevant for windows virtual machines only. • Customizable. • Can be left blank.

Table 7 Network settings only available through the CSV file

Setting	Description	Guidelines
CG ID	The consistency group ID in the RecoverPoint for VMs system.	<ul style="list-style-type: none"> Do not modify this field. Automatically populated by the system. Not customizable. Can be left blank.
CG Name	Name of the consistency group in the RecoverPoint for VMs system.	<ul style="list-style-type: none"> Automatically populated by the system. Must be the name associated with the specified consistency ID in RecoverPoint for VMs. Customizable. Can be left blank.
VC ID	The vCenter Server ID in VMware.	<ul style="list-style-type: none"> Do not modify this field. Automatically populated by the system. Not customizable. Can be left blank.
VC Name	The name of the vCenter Server hosting the virtual machine.	<ul style="list-style-type: none"> Customizable. Can be left blank.
VM ID	The virtual machine ID that vCenter Server uses.	<ul style="list-style-type: none"> Do not modify this field. Automatically populated by the system. Not customizable. Cannot be left blank.
VM Name	The name of the virtual machine.	<ul style="list-style-type: none"> Customizable. Automatically populated by the system. Can be left blank.
NIC Index in vCenter	The index of the adapter in the order of virtual network adapters (NICs) in the virtual machine settings of the vCenter web client.	<ul style="list-style-type: none"> Customizable. Cannot be left blank. Enter a numeric value. Enter a value of 0 to define the first virtual network adapter in the vSphere Web Client. Enter a value of 1 to define the next network adapter.

Changing the network adapter configuration of a protected virtual machine

When the virtual network adapter (NIC) configuration of a production VM changes, any pre-existing copy VM network configuration may be adversely affected and may require re-configuration before it works. After adding or removing NICs from a protected virtual machine, re-configure the copy VM network using [Automatic copy VM network configuration \(Re-IP\)](#) on page 43. If you are already using glue scripts, perform [Manual copy VM network configuration \(RE-IP\)](#) on page 64. If the NIC configuration of a production VM changes and the change is not reflected in the copy VM, ensure **Hardware changes** is enabled and *enable image access by* [Testing a copy](#) on page 52.