

RecoverPoint for Virtual Machines

Version 5.2

Administrator's Guide

302-005-002

REV 03

Copyright © 2017-2019 Dell Inc. or its subsidiaries. All rights reserved.

Published May 2019

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

Tables		5
Preface		7
Chapter 1	Before you begin	9
	Create your license files.....	10
	Access the RecoverPoint for VMs vSphere plugin.....	11
	License and register RecoverPoint for VMs.....	12
Chapter 2	Protecting VMs	15
	Protect a virtual machine in a consistency group.....	16
	Add a copy to a consistency group.....	23
Chapter 3	Monitoring protection	25
	The RecoverPoint for VMs Dashboard.....	26
	Monitoring replication.....	26
	Monitoring system alerts and events.....	28
	Identifying a RecoverPoint for VMs system.....	29
Chapter 4	VM automation and orchestration	31
	Create a bookmark.....	32
	Create an application-consistent bookmark.....	33
	Create a group set.....	35
	Protection orchestration.....	37
	Group start-up sequence.....	37
	Virtual machine start-up sequence.....	37
	Defining user prompts.....	38
	Defining user scripts.....	39
	Automatic copy VM network configuration (Re-IP).....	39
	Protection automation.....	43
	Disabling automatic protection of new VMDKs.....	43
	Provisioning VMDKs.....	43
	Adding a new VMDK.....	43
	Removing a VMDK.....	44
	Excluding a VMDK from replication.....	44
	Automatic replication of VM hardware changes.....	45
	Automatically expanding copy VMDKs.....	45
Chapter 5	Recovering VMs	47
	Test a copy.....	48
	Fail over to a copy and fail back to production.....	50
	Recover production from a copy.....	52
	Monitor recovery activities.....	54

Appendix A	Managing RecoverPoint for VMs	55
	Managing system licenses.....	56
	Managing system registration.....	56
	Managing system support.....	57
	Managing system component registration.....	58
	Managing vCenter Server registration.....	58
	Managing ESX cluster registration.....	59
	Managing journal datastore registration.....	60
	Managing external host registration.....	61
	Managing system protection policy templates.....	61
	Managing virtual machines.....	62
	Stop protecting a virtual machine.....	62
	Managing consistency groups.....	62
	Disabling or enabling a consistency group.....	63
	Managing the protection policies of groups.....	63
	Managing copies.....	65
	Modifying copy policies.....	65
	Configuring copy VM failover networks.....	66
	Managing the protection policies of cloud copies.....	67
	Managing group sets.....	68
	Modifying a group set.....	68
	Removing a group set.....	69
 Appendix B	 Troubleshooting	 71
	Finding the vRPA cluster management IP.....	72
	Collecting system information.....	72
	Collecting RecoverPoint for VMs splitter logs.....	73
	Recovering from a cluster disaster.....	73
	RecoverPoint for VMs licensing.....	73
	Register RecoverPoint by email or phone.....	75
	Creating VMkernel ports.....	75
	Load balancing.....	76
	Manual copy VM network configuration (RE-IP).....	77
	Copy VM network configuration guidelines.....	79
	Changing the network adapter configuration of a protected VM.....	84
 Glossary		 85

TABLES

1	KVSS syntax.....	34
2	Glue script samples.....	78
3	Virtual machine network settings available through the GUI.....	79
4	Network settings only available through the CSV file.....	83

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 Installation and Deployment 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

	<ul style="list-style-type: none"> • 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

Before you begin

Before you start protecting your data in RecoverPoint for VMs, you must perform the following tasks:

Before deploying the cloud solution, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.

NOTICE

This guide provides the procedures for protecting, recovering and managing VMs with on-premises local and/or remote RecoverPoint for VMs copies. In RecoverPoint for VMs 5.2.1 and later versions, you can also protect your VMs by creating a copy of them in the Amazon Cloud. To deploy the RecoverPoint for VMs Cloud Solution and protect your virtual machines with a cloud copy, see the *RecoverPoint for VMs Cloud Solutions Guide*.

• Create your license files	10
• Access the RecoverPoint for VMs vSphere plugin	11
• License and register RecoverPoint for VMs	12

Create your license files

When a RecoverPoint for Virtual Machines sales order is approved, a License Authorization Code is automatically sent to the email addresses provided during order entry. The License Authorization Code contains your entitlements. You must activate each entitlement and then save it as a license file before it can be entered into the RecoverPoint for VMs vSphere plugin.

For a detailed description of how RecoverPoint for VMs is licensed, see [RecoverPoint for VMs licensing](#) on page 73.

Procedure

1. Access the entitlements on support.emc.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 support site, and searches for all entitlements associated with the License Authorization Code.
 - If you do not have the **License Authorization Code** email but you do have the LACs or sales order numbers, log into 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 entitlements in each License Authorization Code email.

Results

The entitlements are converted to license files.

After you finish

Transfer the license files to the computer from which you will be running RecoverPoint for VMs.

Access the RecoverPoint for VMs vSphere plugin

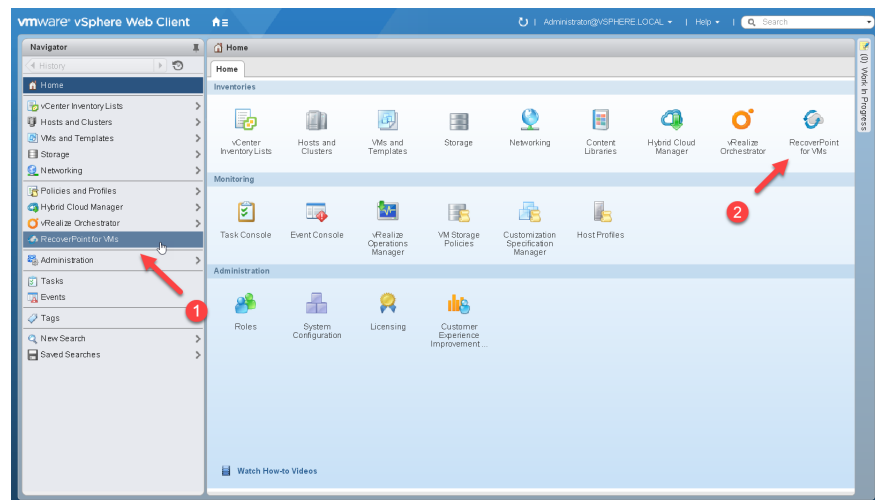
There are two ways in which you can access the RecoverPoint for VMs plugin in the vSphere Web Client.

Before you begin

Connect to the vSphere Web Client of your production site.

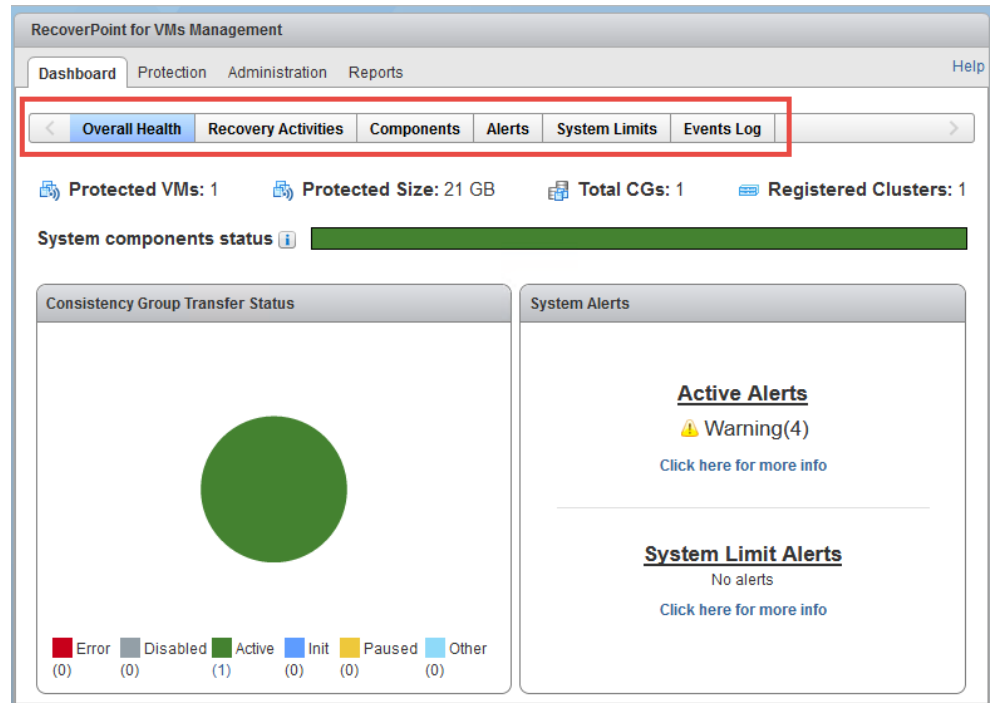
Procedure

1. Click the **RecoverPoint for VMs** menu item in your **vSphere Web Client > Navigator**.
2. Click the **RecoverPoint for VMs** icon in your **vSphere Web Client > Navigator > Inventories**.



Results

The RecoverPoint for Virtual Machines **Dashboard** is displayed.



To monitor the status of your vRPA clusters, click the **Dashboard > Components** tab and ensure that a green checkbox appears next to each vRPA cluster and that the vRPA cluster **Status** column contains an **OK**.

To monitor the environment, click the **Overall Health, Recovery Activities, Components, Alerts, System Limits, and Events Log** sub-tabs of the system Dashboard.

License and register RecoverPoint for VMs

The **Getting Started Wizard** will guide you through the process of entering a license file, registering the product, and enabling system support.

Before you begin

- 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 **RecoverPoint for VMs vSphere plug-in**, click **Administration > vCenter Servers > Licensing**.
2. Click **Add...** under the **Registered Licenses** table.
The **Getting Started Wizard** is displayed.
3. In the **Welcome** screen, click **Next**.
4. In the **Licensing** screen, click **Browse...** to locate and select the license file (*.lic extension). Click **Next**.

5. In the **Support** screen, to provide communication between the RecoverPoint for VMs system and the System Reports database, select **Enable pre-emptive support for RecoverPoint for VMs**.
 - a. 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.
 - To transfer system notifications through the FTPS server, in the **Transfer Method** section, select **FTPS**.
 - To transfer system notifications through the Secure Remote Services gateway, in the **Transfer Method** section, select **ESRS**. In the **ESRS gateway IP address** field, specify the IP address of the Secure Remote Services gateway in IPv4 format..
 - b. Click **Test Connectivity**. Wait 10 minutes. Then, click **Dashboard > Events Log** and look for event 1020: "Failed to send system report".
 - If this event does not appear in the **Events Log**, the system notifications mechanism is correctly configured.
 - If you do receive event 1020: "Failed to send system report", check whether there is an issue with the selected method of transfer. If a problem exists, fix it, re-configure support, and click **Test Connectivity** again. If the problem persists, contact Customer Support.
 - c. Click **Next**.
6. In the **Registration** screen:
 - a. Register your RecoverPoint for VMs system, at the current vRPA cluster:
 - **Company name**: The name of your company, as it appears on your sales order.
 - **Connect home method**: The method that is used to send configuration reports and alerts to Dell EMC. The connect home method allows Dell EMC to pro-actively address issues within the RecoverPoint for VMs environment, should they arise.
 - **Connect in method** The method that is used to allow remote connectivity to the RecoverPoint environment. Enabling this feature is recommended as it enables secure access to the RecoverPoint for VMs environment to gather logs and resolve issues as quickly as possible. 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. When 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 *Secure Remote Services Gateway Operation Guide* for further instructions on Config Tool usage. 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.
 - **License type**: Displays the type of the license that has been registered in RecoverPoint for VMs. Ensure the displayed license type is **RecoverPoint for VMs**.

- **Location:** The city, state, and country where your company is located.
 - **Sales order number** If you don't have your sales order, your Customer Engineer can provide it.
 - **Site (party) ID:** The unique ID of the customer site. This value is automatically retrieved from the registered license file and can only be modified by Customer Service.
- b. If your company does not have outside connectivity, click **Export to CSV** to export the registration information to a CSV file.
- c. Enter the email address to which a verification email should be sent when the registration information is updated in the Install Base in the **Send verification email to** field.
-

Note

Skip this step if your company does not have outside connectivity.

- d. Click **Next**.
7. In the **Ready to complete** screen, verify that the information is correct, and click **Finish**.

Results

If your company has outside connectivity, 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 for every vRPA cluster in the RecoverPoint for VMs system.

After you finish

If your company does not have outside connectivity, use the exported CSV file to register by email or phone, as described in [Register RecoverPoint by email or phone](#) on page 75.

CHAPTER 2

Protecting VMs

In RecoverPoint for VMs, consistency groups are used to protect virtual machines and replicate virtual machine application data to a consistent point in time. A consistency group is a logical entity that constitutes a container for virtual machines and all of their copies.

Consistency groups can protect many VMs. If this is the first time you are using RecoverPoint for VMs, protect your virtual machines by creating new consistency groups for them, or by adding them to an existing consistency group. If you already have RecoverPoint for VMs consistency groups, you can create a new copy to protect your production VMs, alongside your existing copy.

Before protecting VMs, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.

NOTICE

This Administrator's Guide provides the procedures for protecting, recovering and managing VMs with on-premises local and/or remote RecoverPoint for VMs copies. In RecoverPoint for VMs 5.2.1 and later versions, you can also protect your VMs by creating a copy of them in the Amazon Cloud. To deploy the RecoverPoint for VMs Cloud Solution and protect your virtual machines with a cloud copy, see the *RecoverPoint for VMs Cloud Solutions Guide*.

- [Protect a virtual machine in a consistency group](#).....16
- [Add a copy to a consistency group](#).....23

Protect a virtual machine in a consistency group

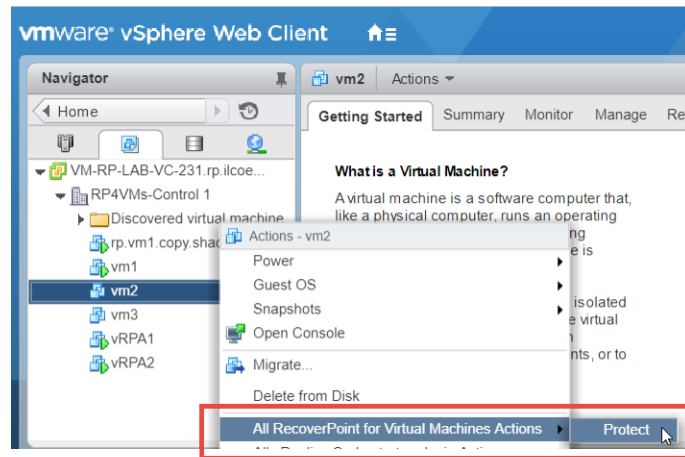
The RecoverPoint for VMs **Protect VMs Wizard** will guide you through the process of protecting your production VMs.

Before you begin

Ensure all vCenter Servers that manage production VMs and copy VMs are registered according to [Managing vCenter Server registration](#) on page 58

Procedure

1. Connect to the vSphere Web Client of your production site.
2. Select **VMs and Templates** view.
3. Power on the virtual machine that you want to protect.
4. Right-click on the virtual machine and select **All RecoverPoint for Virtual Machines Actions > Protect**.



Note

Protecting a virtual machine with fault tolerance enabled is not supported.

The **Protect VMs Wizard** is displayed.

The minimum input required in each screen of the **Protect VMs Wizard** is indicated in red, in the following screenshots.

5. In the **Select VM protection method** screen:

- **Create a new consistency group for this VM.** Type a descriptive name for the new consistency group. Best practice is to use the VM or application name as your consistency group name. Ensure the production vRPA cluster is selected. 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*.

6. In the **Configure production settings** screen:

Name	Total Size	Estimated Free Space
SharedDS1	200GB	49.9GB
datastore1	512MB	495MB

- a. Enter a name for the production copy. Best practice is to differentiate the production copy name from the replica copy name (for example, use "Production" or the production site location).
- b. If you chose to create a new consistency group in the previous step (not relevant if adding a VM to an existing group):
 - Accept or define the minimum **Journal Size** for the production copy. The default size (3GB) is the most practical size for most production journals.
 - Optionally, select a specific datastore to use for the production journal. By default, RecoverPoint automatically registers up to 15 datastores for the production journal and automatically selects the datastore with the most free space.

NOTICE

RecoverPoint for VMs will attempt to create the journal on the selected datastore. If for any reason journal creation fails, the system will attempt to create the journal on another registered datastore.

-
- If you want to select a specific datastore from the registered datastores, select **Manually select a registered datastore from the table below** and select it in the table.
 - If you want to select a datastore that hasn't been registered, click **Register Datastore**. Select the datastore and click **Register**. Ensure the required datastore is selected in the table.
- c. 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 lazy zeroed`, `Thick provisioning eager zeroed` Or `Thin provisioning`.
 - **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.

d. Click **Next**.

7. In the **Add a copy** screen, define the copy.

Protect VM(s) Wizard

1. Select VM protection method
2. Configure production settings
3. Add a copy
4. Configure copy settings
5. Select copy resources
6. Define failover networks
7. Ready to complete

Create a new copy to protect the production VMs of 'my_cg':

Enter a name for the new copy:
copy1
Copy names can contain up to 100 characters and must be unique.

Select the copy type:
AWS Cloud

Select the copy vRPA cluster:

Search

- Cluster1 (local)
- Cluster2 (remote)

Back Next Protect Cancel

- a. Enter a name for the copy. Best practice is differentiate the replica copy name from the production copy name (for example, use "Remote Copy" or the copy site location) .

NOTICE

There is no need to change the value of the **Select copy type** field unless you want to protect your VM with a copy on the AWS cloud. For the instructions on how to protect VMs on the AWS cloud, see the *RecoverPoint for VMs Cloud Solutions Guide*.

- b. Ensure the vRPA cluster that will manage the group data is selected.
- c. Click **Next**

8. In the **Configure copy settings** screen:

Protect VM(s) Wizard

✓ 1. Select VM protection method
✓ 2. Configure production settings
✓ 3. Add a copy
✓ 4. Configure copy settings
5. Select copy resources
6. Define failover networks
7. Ready to complete

Copy: copy1
Configure the journal at this copy:

Journal Size: 10 GB

Select a registered datastore for the journal:
☒ Automatically select a registered datastore from the table below
☐ Manually select a registered datastore from the table below

Name	Total Size	Estimated Free Space
Cluster1_DS	200GB	151GB

Register Datastore...

Define a protection policy for this copy:
☐ Load policy from template default remote link
☒ Manually define policy

Replication Mode:
☐ Synchronous
☒ Asynchronous

RPO: 25 Seconds

The host application initiates a write, and does not wait for an acknowledgment from the vRPA at the target site before initiating the next write. The primary advantage of asynchronous replication is its ability to provide synchronous-like replication without degrading the performance of host applications.

Back Next Protect Cancel

a. Configure the copy journal settings:

- Define or accept the minimum **Journal Size** for the copy journal. The larger the copy journal, the more history can be saved.
- By default, RecoverPoint automatically selects a datastore 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.

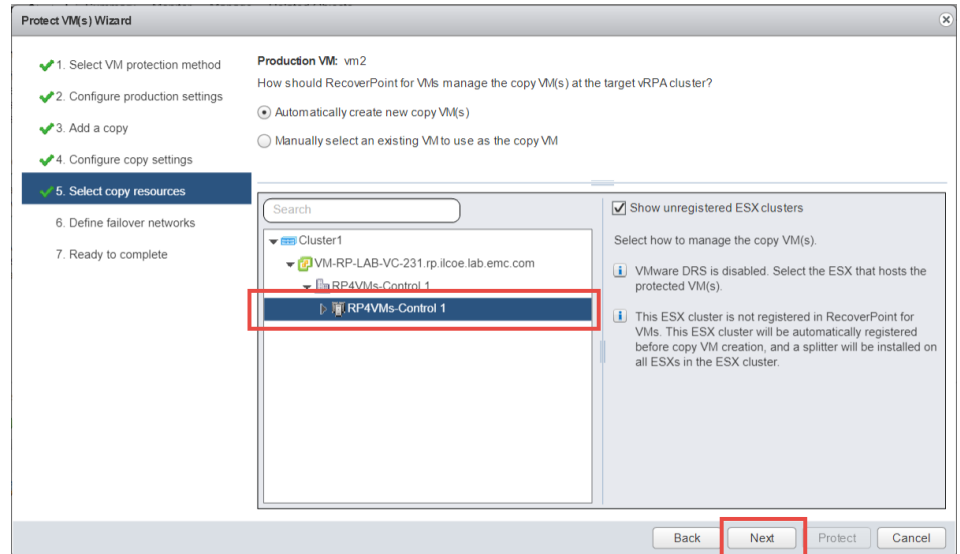
NOTICE

RecoverPoint for VMs will attempt to create the journal on the selected datastore. If for any reason journal creation fails, the system will attempt to create the journal on another registered datastore.

- 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, Or Hours**.

b. Click **Next**.

9. In the **Select copy resources** screen, select where and how to create the copy VM which will protect the virtual machine:



- To **Automatically create new copy VM(s)**, expand the tree, select the ESX host, ESX cluster, or VMware Resource Pool that will manage the copy VM. Click **Next**. Repeat this process for all production VMs.
- To **Manually select an existing VM to use as the VM copy**, expand the tree, select an ESX host, ESX cluster, or VMware Resource Pool, and select an existing virtual machine. Repeat this process for all production VMs, click **Next**, and skip the next step.

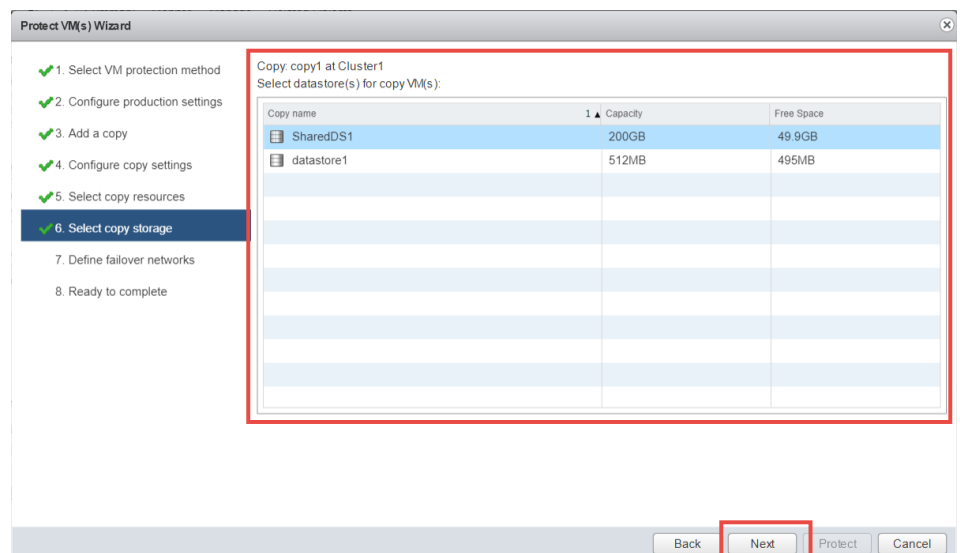
10. In the **Select copy storage** screen:

- a. Select the datastore(s) that will contain the data of the copy VM.

Note

All VMDKs are mapped to a single datastore. If the datastore you need is not displayed, ensure you have completed [Managing vCenter Server registration](#) on page 58, and have mounted the shared datastore(s) on all ESX hosts in every ESX cluster or VMware Resource Pool.

b. Click **Next**



11. (Optionally) In the **Define failover networks** screen, for each network adapter of each production VM, specify the port groups to use when [Fail over to a copy](#) and [fail back to production](#) on page 50, and when [Test a copy](#) on page 48, and click **Next**. See [Configuring copy VM failover networks](#) on page 66 for more information.

NOTICE

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

12. In the **Ready to complete** screen:

- a. Ensure your protection settings are as required:

- Expand the **Production** and **Copy** settings to ensure that they are correct.
- If need be, click **Edit...** to change a setting before clicking **Protect**.
- Note if a warning is displayed regarding a potential communications problem (no action required at this time).

- If you do not want to start replicating data from the production VM immediately, uncheck **Start replicating this group when I click Protect**.
- Click **Add a Copy** to add more copies to the group.

Note

If a warning regarding a potential communications problem is displayed, see [Creating VMkernel ports](#) on page 75.

- b. Click **Protect** to create the copy(s) and enable VM protection.

Results

The specified virtual machine(s) are protected. If you added 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 virtual machines in the consistency group.

Note

If an unregistered ESX cluster, or an ESX host or VMware Resource Pool of an unregistered ESX cluster were selected to manage the copy, the unregistered cluster is automatically registered with the specified vRPA cluster, a splitter is installed on all ESXs in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

After you finish

- Use the **Protection > Consistency Groups** screen for [Monitoring replication](#) on page 26.
- Use [The RecoverPoint for VMs Dashboard](#) on page 26 to monitor the system.
- To stop replicating a VM, see [Stop protecting a virtual machine](#) on page 62.

Add a copy to a consistency group

The RecoverPoint for VMs **Add a Copy Wizard** will guide you through the process of protecting your virtual machines with a new local and/or remote copy.

Before you begin

- [Create your license files](#) on page 10
- [License and register RecoverPoint for VMs](#) on page 12
- Ensure all vCenter Servers that manage production VMs and copy VMs are registered according to [Managing vCenter Server registration](#) on page 58

Procedure

1. Select **Protection > Consistency Groups**.
2. Click the **Add a copy** icon:



3. Follow the instructions for protecting VMs, starting from **step 7** on page 19.

CHAPTER 3

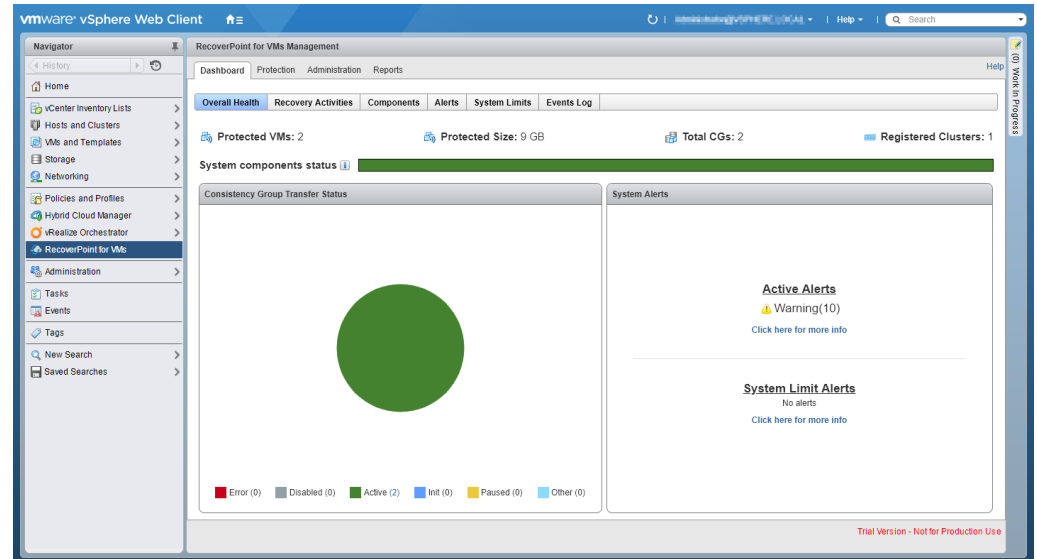
Monitoring protection

After protecting your VMs, use the RecoverPoint for VMs **Dashboard** to monitor the system. Use the **Protection > Consistency Groups** screen to monitor replication.

• The RecoverPoint for VMs Dashboard	26
• Monitoring replication	26
• Monitoring system alerts and events	28
• Identifying a RecoverPoint for VMs system	29

The RecoverPoint for VMs Dashboard

Use the RecoverPoint for VMs **Dashboard** to attain a high-level overview of the RecoverPoint for VMs system. The RecoverPoint for VMs **Dashboard** and its sub-tabs present important system information to help you analyze and monitor your RecoverPoint for VMs system.



Procedure

1. In your **vSphere Web Client** > **Navigator**, select **RecoverPoint for VMs** to access the system **Dashboard**.
2. To monitor your RecoverPoint for VMs system, click the **Overall Health**, **Recovery Activities**, **Components**, **Alerts**, **System Limits**, and **Events Log** sub-tabs.

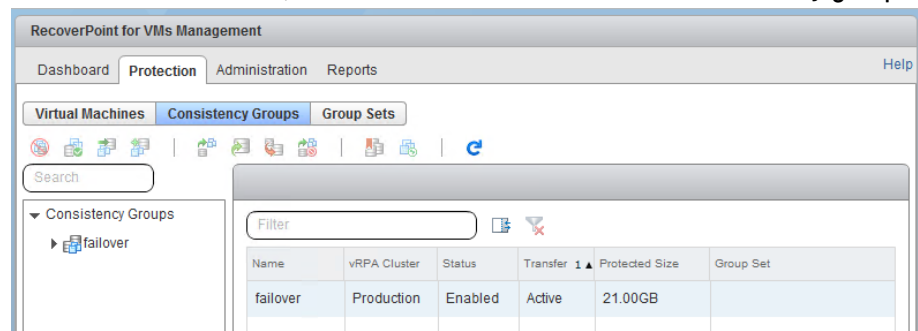
Monitoring replication

Use the RecoverPoint for VMs **Protection** tab to monitor all aspects of replication.

Procedure

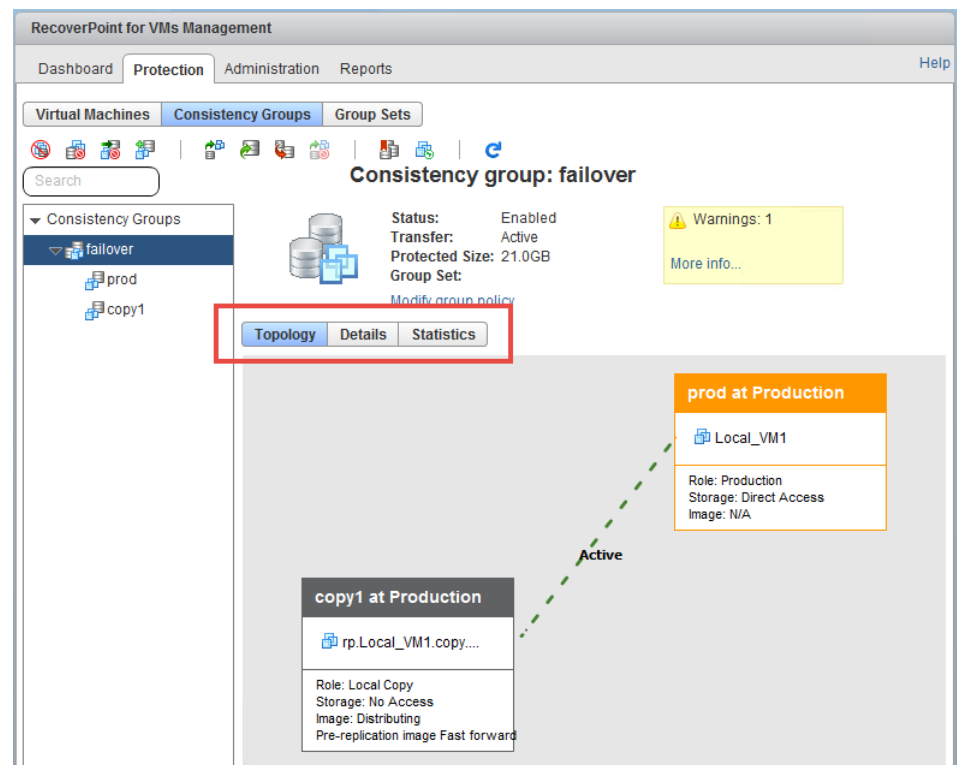
1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection** > **Consistency Groups**.

In the **Transfer** column, note the **Transfer State** of each consistency group.



- **Active:** Data is being transferred asynchronously to a copy.

- **Active (Synchronized)**: Data is being transferred synchronously to a copy.
 - **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.
 - **Paused**: Data is not being transferred to a copy, because transfer has been paused by the user.
 - **Paused by System**: Data is not being transferred to a copy, because transfer has been paused by the system. If this state occurs for long periods of time, check the system alerts and events in the **Dashboard** for more information.
 - **N/A**: Data is not being transferred to a copy, because the copy has been disabled by the user.
2. Select a specific consistency group. The state of transfer to each of the copies in the group, and other important information about the replication process, is displayed in the group **Topology** diagram.



3. Select the **Details** and **Statistics** sub-tabs for more detailed information about your replication environment, configuration, and performance.

Results

If after protecting a VM, the transfer state of its group does not become **Active**, see [Creating VMkernel ports](#) on page 75.

Monitoring system alerts and events

Use RecoverPoint for VMs events and alerts to understand and troubleshoot events in your RecoverPoint for VMs environment.

To monitor your system events, use the **RecoverPoint for VMs vSphere plugin > Dashboard > Events Log** tab.

An event is a notification that a change has occurred in the state of a system component. In some cases, the change indicates an error or warning condition for the component. Multiple events can occur simultaneously on a single component. A single incident can generate events across multiple system components.

In RecoverPoint for VMs, events have a:

- **Level:** Info, Warning, Or Error
- **Scope:** Normal, Detailed, Or Advanced
- **Topic:** All, vRPA Cluster, vRPA, Group, Splitter, Or Management

RecoverPoint for VMs Management

Dashboard Protection Administration Reports

Overall Health Recovery Activities Components Alerts System Limits **Events Log**

Events

Search Refresh Last updated: 02/01/2019 21:59:47 Warnings and errors only from the last 24 hours

Level	Scope	Point in Time	vRPA Cluster	ID	Topic	Summary
Cleared Warning	Root cause	02/01/2019 16:53:25	winterfell	3000	RPA	RPA successfully communicating with ...
Cleared Warning	Root cause	02/01/2019 17:11:12	N/A	3000	RPA	RPA successfully communicating with ...
Cleared Warning	Normal	02/01/2019 17:06:24	winterfell	3022	RPA	LAN link to RPA restored
Cleared Warning	Normal	02/01/2019 16:53:47	N/A	3007	RPA	RPA is up
Cleared Warning	Normal	02/01/2019 16:59:43	winterfell	3022	RPA	LAN link to RPA restored
Cleared Warning	Normal	02/01/2019 17:00:11	N/A	3007	RPA	RPA is up
Error	Normal	02/01/2019 17:27:45	winterfell	4003	Consistency group	Group capabilities problem

Total Items: 63 (Showing Items: 20)

1 - 20 of 63 items

Details

Group capabilities problem (Group(s): [KATE_CG_1546442831531 - prod])

[Read more](#)

Trial Version - Not for Production Use

In the events log:

1. Select an event to display the event **Details**.
2. If there is detailed event information that can help you to troubleshoot, a **Read more** link is displayed. Click this link to display additional information about the selected event.
3. Click the **Event Filter** icon to configure which events are displayed in the **Events Log**.


NOTICE

For more event monitoring and troubleshooting options, log into the RecoverPoint for VMs **Command Line Interface (CLI)** and run the `get_events_log` command. For more information, see the *RecoverPoint for Virtual Machines CLI Command Reference Guide*.

To monitor your system alerts, click the **RecoverPoint for VMs vSphere plugin > Dashboard > Alerts** tab. System alerts are a mechanism that allows vRPAs to send

events about system components in real-time, to a specified email, or the system reports database, via SMTP.

- To manage your system alerts settings, see [Managing system support](#) on page 57.
- You can also monitor the alerts of specific consistency groups from the **Protection > Consistency Groups** screen, when you select a specific consistency group. Click the **More info...** link under the alert box for more details about these alerts.



Identifying a RecoverPoint for VMs system

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

Procedure

1. In the **RecoverPoint for VMs vSphere plugin**, select **Administration > vRPA Clusters > vRPA System**
2. Select a vRPA cluster.
3. Note the value of **Other vRPA clusters in system**.

CHAPTER 4

VM automation and orchestration

RecoverPoint for VMs provides the following features that automate and orchestrate the protection of your VM copies.

- [Create a bookmark](#)..... 32
- [Create an application-consistent bookmark](#)..... 33
- [Create a group set](#)..... 35
- [Protection orchestration](#)..... 37
- [Protection automation](#)..... 43

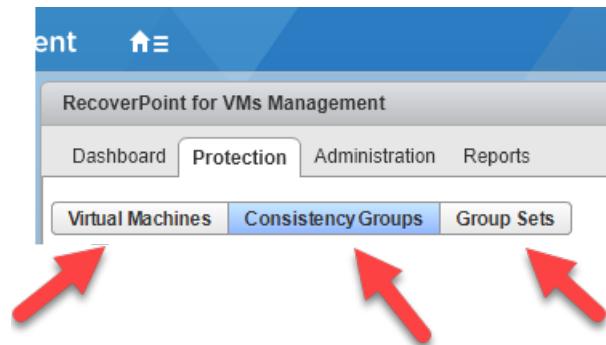
Create a bookmark

Create a snapshot of a virtual machine, a consistency group, or a group set, and label it for easy identification during testing and recovery.

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.

Procedure

1. In the **RecoverPoint for VMs vSphere plugin**, click **Protection**.



- To bookmark a virtual machine, select the **Virtual Machines** tab.
 - To bookmark a consistency group, select the **Consistency Groups** tab.
 - To bookmark a group set, select the **Group Sets** tab.
2. Select the consistency group or group set that you want to bookmark.
 3. Click the **Create Bookmark** button at the top of the screen:



4. In the **Create Bookmark** dialog box:

 A screenshot of the 'Create Bookmark' dialog box. It has a title bar with 'Create Bookmark' and a close button. The form contains:

- 'Bookmark Name:' followed by a text input field.
- 'Label Bookmark as:' followed by a dropdown menu currently showing 'Crash-Consistent'.
- 'Consolidation Policy:' with two radio button options:
 - ☒ Never consolidate this snapshot
 - ☐ This snapshot must survive [Daily] consolidations
- An information icon (i) and a note: 'The snapshot consolidation policy is not relevant for cloud copies.'
- 'OK' and 'Cancel' buttons at the bottom right.

- **Name** - Type a name for the snapshot. This is the bookmark. The bookmark is the name that will be used to identify the snapshot during testing and recovery.
- **Label Bookmark As** - Select:

- **Crash-Consistent:** Labels the snapshot as crash-consistent.
- **Application-Consistent:** Labels the snapshot as application-consistent. Selecting this option does not create an application-consistent snapshot, it only labels the snapshot as application-consistent.
- **Consolidation Policy** - Specifies 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 will survive daily consolidations, but is consolidated weekly and monthly.
 - **Weekly** - Snapshot will survive daily and weekly consolidations, but is consolidated monthly.
 - **Monthly** - Snapshot will survive daily, weekly, and monthly consolidations.

5. Click OK.

Results

A crash-consistent snapshot is created with the specified name for the specified virtual machine, consistency group, or group set.

Create an application-consistent bookmark

KVSS bookmarks are created using the `kvss.exe bookmark` command. The working folder for running KVSS commands is `%SystemDriver%/EMCRecoverPointVSSProvider/`.

Before you begin

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 1 KVSS syntax

Option	Description
writers	A VSS-aware host application
groups	RecoverPoint consistency group
bookmark	Name by which you can identify the applied bookmark
consolidation_policy	<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>
type	The shadow copy type:

Table 1 KVSS syntax (continued)

Option	Description
	<ul style="list-style-type: none"> <i>FULL</i> <i>COPY</i> <p>This setting is optional. Default = <i>COPY</i>. The writer application determines the settings. Generally, when <i>type</i> = <i>full</i>, backup logs are truncated. When <i>type</i> = <i>copy</i>, backup logs are not truncated.</p>
<i>ip</i>	vRPA cluster management IP
<i>user</i>	RecoverPoint for VMs username
<i>password</i>	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"
```

Create a group set

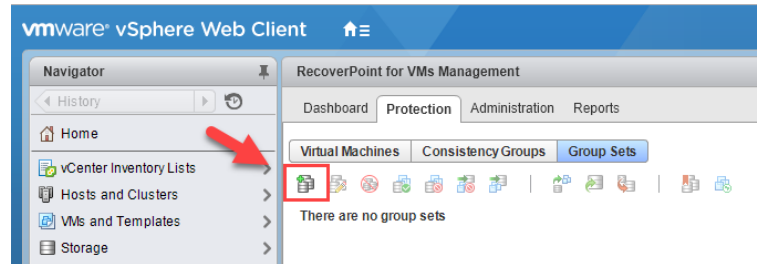
A group set is a collection of consistency groups to which the system applies parallel bookmarks at a user-defined frequency. Group sets are useful for consistency groups that are dependent on one another or that must work together as a single unit. If any of the groups in the group set are part of another group set that has parallel bookmarking enabled, you cannot enable parallel bookmarking for that group set.

NOTICE

If a group set contains a consistency group with both an on-premises copy and a cloud copy, parallel bookmarks can be enabled even though this feature is not supported for cloud copies. In this case, parallel bookmarks are applied only to the on-premises copy.

Procedure

1. Click **Protection > Group Sets**.
2. Click the **Add Group Set** icon:



3. In the **Add Group Set** dialog box:

	Consistency Group	Replicating From	Copies in Group
<input checked="" type="checkbox"/>	Consistency Group 1	My vRPA Cluster	4
<input checked="" type="checkbox"/>	Consistency Group 2	My vRPA Cluster	4
<input type="checkbox"/>			
<input type="checkbox"/>			
<input type="checkbox"/>			

☐ Select All

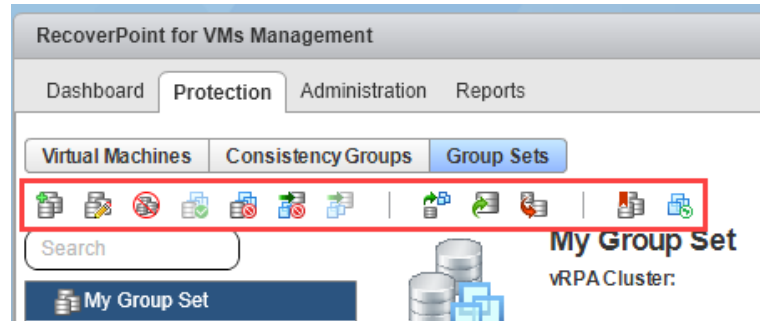
- a. Enter a descriptive name for the group set.
 - b. Choose the vRPA cluster from which to select consistency groups.
 - c. Select the consistency groups to add to the group set.
 - d. To create bookmarks for all of the consistency groups in this group set, at the same pre-defined intervals, select **Enable Parallel Bookmarking** and set the **Frequency** value.
4. Click **OK**.

Results

The group set is created.

After you finish

Select a group set and use the buttons at the top of the screen to bookmark, enable/disable, modify the start-up sequence, and pause/start replication for multiple groups, simultaneously.



Protection 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.

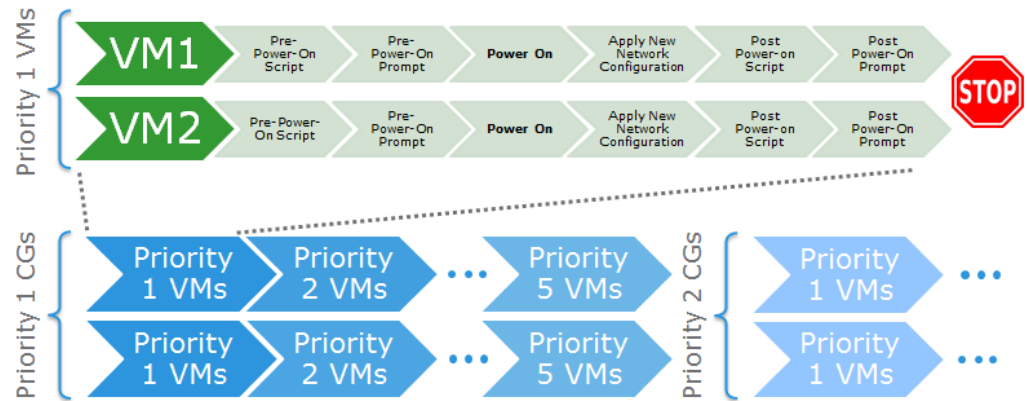
Before you begin

- 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 38 and [Defining user scripts](#) on page 39.

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, Red Hat Enterprise Linux (RHEL) server version 7.1, 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 77
- 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 77.

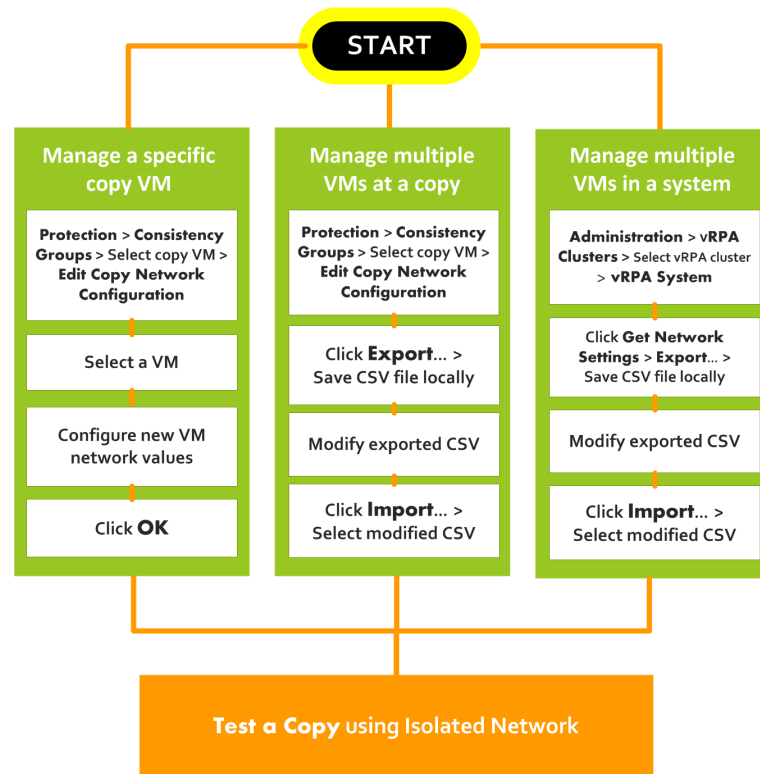
Note

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 it works.

Note

If you have upgraded from a RecoverPoint for VMs version prior to 5.0.1, see [Manual copy VM network configuration \(RE-IP\)](#) on page 77 for how it works.



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

Before you begin

You should be familiar with:

- [Test a copy](#) on page 48
- [Fail over to a copy and fail back to production](#) on page 50
- [Recover production from a copy](#) on page 52
- [Copy VM network configuration guidelines](#) on page 79

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:

- [Test a copy](#) on page 48
- [Fail over to a copy and fail back to production](#) on page 50
- [Recover production from a copy](#) on page 52
- [Copy VM network configuration guidelines](#) on page 79

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:

- [Test a copy](#) on page 48
- [Fail over to a copy and fail back to production](#) on page 50
- [Recover production from a copy](#) on page 52
- [Copy VM network configuration guidelines](#) on page 79

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.

Protection automation

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

Disabling automatic protection of new VMDKs

By default, all newly added VMDKs are automatically protected. Use this procedure to change the default behavior.

Procedure

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection > Virtual Machines**.
2. Select the virtual machine for which you want to disable the automatic protection of any newly added VMDKs, in the future.
3. Under the **Protected VMDKs** widget, click **Edit...**
4. In the **Edit VMDK Protection Policy** dialog, clear the **Automatically protect new VMDKs** checkbox to disable automatic protection of any newly added VMDKs to this VM.
5. Click **OK**.

Results

Any VMDKs added to this VM in the future will not be automatically protected.

Provisioning VMDKs

By default, VMDKs are provisioned `Same as source`.

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 lazy zeroed`, `Thick provisioning eager`, or `Thin provisioning`.

Results

The VMDK is provisioned according to the specified settings.

Adding a new VMDK

RecoverPoint for VMs automatically detects when a hard disk is added to a protected virtual machine using the **vSphere client > Virtual Machine Properties**, and orchestrates VMDK addition.

When adding VMDKs to a protected VM, RecoverPoint for VMs automatically starts protecting each added VMDK by creating copies of it, as specified by the VM's consistency group link and copy settings.

- To disable automatic protection for all VMDKs added to a protected VM in the future, see [Disabling automatic protection of new VMDKs](#) on page 43.
- To exclude specific VMDKs of a protected VM from protection, see [Excluding a VMDK from replication](#) on page 44.

RecoverPoint for VMs does not automatically protect VMDKs of type `shared` when they are added to a protected VM.

Results

A volume sweep occurs on the added VMDK(s) and a short initialization occurs on all other VMDKs in the consistency group, but no history is lost.

Removing a VMDK

RecoverPoint for VMs automatically detects when a hard disk is removed from a protected virtual machine using the **vSphere client > Virtual Machine Properties**, and orchestrates VMDK removal.

If you don't want to replicate a specific VMDK of a protected VM to the copy, you can remove the VMDK from the production VM through the vSphere client **Virtual Machine Properties**, or exclude it from replication, as described in [Excluding a VMDK from replication](#) on page 44.

Note

Removing VMDKs from a protected VM does not delete their copies and does not remove their history from the copy journal.

Results

During recovery, the VM snapshots will not include the removed VMDK, even when recovering a snapshot from a point in time that precedes the removal of the VMDK. Future VM snapshots will not include the removed VMDK.

After you finish

The removed VMDKs can be recovered by [Recovering VMs](#), and selecting a snapshot that precedes the removal of the VMDK.

After removing VMDKs from a production VM in vSphere, you can:

- Add the missing VMDKs 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.
- Follow the instructions for [Excluding a VMDK from replication](#) on page 44.

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 in vSphere, or mark these VMDKs to be excluded from replication in the **RecoverPoint plugin for vSphere**, and replicate the virtual machines without them. You can also use the following procedure to include an excluded VMDK.

- The excluded production VMDKs are not replicated, but the corresponding copy VMDKs are not deleted.

Note

VMDKs that are excluded are created at the copy and attached to the copy VM, but any writes going to excluded VMDKs are not replicated to the copy.

- 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.

Procedure

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection > Virtual Machines**.
2. Under the **Protected VMDKs** widget, click **Edit....**
3. Clear the checkbox next to the VMDKs that you want to exclude from replication.
4. Click **OK**.

Results

In the future, the excluded VMDKs are not displayed in the list of snapshots that you can select when [Recovering VMs](#) on page 47 from a previous point in time, even when recovering snapshots from a time previous to the VMDK removal.

Automatic replication of VM hardware changes

By default, RecoverPoint for VMs automatically detects any hardware changes that are made to a protected VM (through the **vSphere client > Virtual Machine Properties**), and orchestrates the replication of hardware changes to all copy VMs. 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.

RecoverPoint for VMs replicates the protected VM's **version, MAC address, CPU, memory, resource reservations, network adapter status and network adapter type** to all copy VMs in the consistency group, upon image access.

Procedure

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

Results

RecoverPoint for VMs stops replicating changes to the protected VM's hardware.

Automatically expanding copy VMDKs

When a production VMDK is expanded, 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:
 - Replicating to RDM. After expanding the production VMDK/RDM, you will need to manually expand the copy RDM.
 - 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:
 - Replicating to RDM. After expanding the production VMDK/RDM, you will need to manually expand the copy RDM for replication to resume.
 - 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 (for example, because the production VMDK has been removed and re-added with a smaller size). To mitigate this situation, 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. In this case, 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.

CHAPTER 5

Recovering VMs

Periodically test copy images. In a disaster, fail over to a copy, or recover production to an earlier point-in-time.

Before recovering VMs, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.


NOTICE


This Administrator's Guide provides the procedures for protecting, recovering and managing VMs with on-premises local and/or remote RecoverPoint for VMs copies. In RecoverPoint for VMs 5.2.1 and later versions, you can also protect your VMs by creating a copy of them in the Amazon Cloud. To deploy the RecoverPoint for VMs Cloud Solution and protect your virtual machines with a cloud copy, see the *RecoverPoint for VMs Cloud Solutions Guide*.

• Test a copy	48
• Fail over to a copy and fail back to production	50
• Recover production from a copy	52
• Monitor recovery activities	54

Test 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.



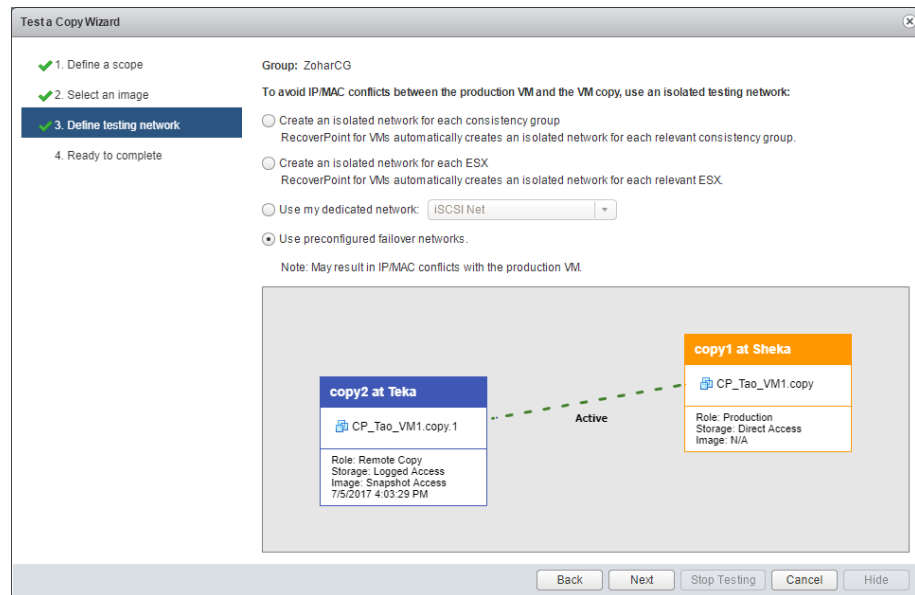
Uncheck the **Power on copy VMs during testing** checkbox if you want the copy VMs powered off during image access. The default, checked, powers on the virtual machine during image access. Unselecting this option skips the post-power up steps (including copy VM network reconfiguration) in the start-up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

When selecting the image, you have the following options:

- **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:** When selecting an image from the list, 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:** Allows you to perform a customized search.
 - **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.

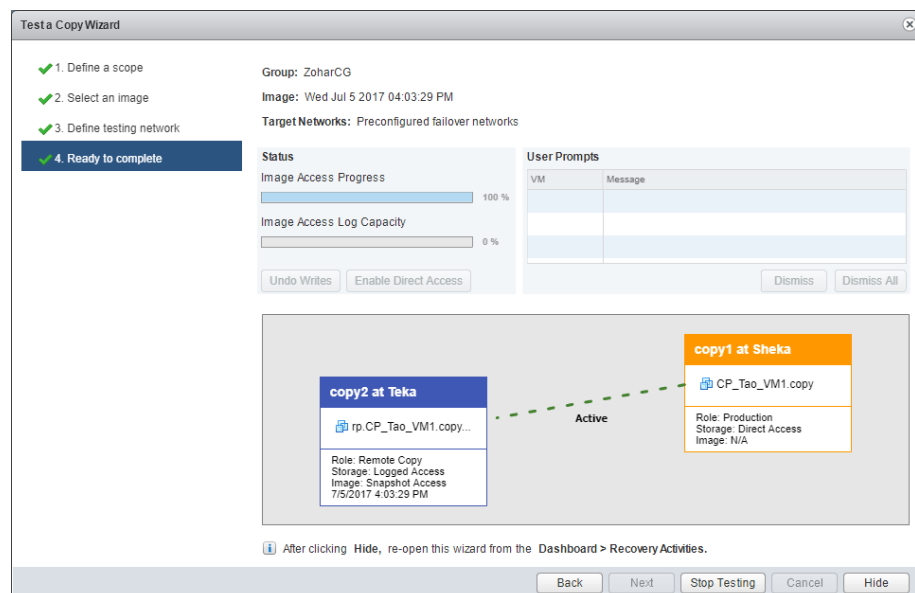
4. In the **Define testing network** screen, to avoid IP conflicts between the production and copy VMs, best practice is to use a dedicated testing network.



When defining the testing network, you can:

- **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 *group sets*.
- **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 66. Not relevant for *group sets*.

5. 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**.
- To roll back all of the writes made to the copy during image access, click **Undo Writes**.
- To write directly to the copy storage click **Direct Access**. Any changes made to the copy storage in direct access mode cannot be automatically undone, because when an image is directly accessed, the journal at the copy is deleted. However, direct access mode does not impose a limit to the amount of data that you can write to a copy storage.

Fail over to a copy and fail back to production

Use the **Failover Wizard** to guide you through the process of selecting a copy image, testing it, failing over, and (if need be) failing back to the production.

Before you begin

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.

Procedure

1. Select **Protection > <Select a Consistency Group, Group Set, or VM> > 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.

When selecting the image, you have the following options:

- **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:** When selecting an image from the list, 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:** Allows you to perform a customized search.
 - **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.
- When defining the testing network, you can:
- **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 *group sets*.
 - **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 66. Not relevant for *group sets*.
5. In the **Ready to complete** screen:
- Clicking **Hide** keeps access to the image enabled at the specified copies and exits the wizard. 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.
- a. To ensure that the failover is configured correctly, review the displayed summary information.
 - b. To define the **Target Networks** for failover, click **Edit....**
 - c. To roll back all of the writes made to the copy during image access, click **Undo Writes**.
 - d. To write directly to the copy storage click **Direct Access**. Any changes made to the copy storage in direct access mode cannot be automatically undone, because when an image is directly accessed, the journal at the copy is deleted. However, direct access mode does not impose a limit to the amount of data that you can write to a copy storage.
 - e. Once image access is enabled, click **Fail Over** to start failover.

Results

- 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 production journal becomes the copy journal and the copy journal becomes the production journal.

- The marking information in the production journal is deleted, the copy journal is deleted, and the consistency group undergoes a full sweep synchronization.

After you finish

- You may want to add VMDKs to the new copy journal to ensure that you have ample space for copy testing. To add VMDKs, see [Adding a new VMDK](#) on page 43.
- To fail back to the production, use the **Failover Wizard** to select an image at the production that predates the failover, and test the image before permanently selecting it as the image you want to fail back to.
- After failing back to the production, if you added VMDKs to the production journal after failover, to reset the production journal to its original size (by default, 3GB) without triggering a full sweep click **Protection > Consistency Groups**, select the group's production copy, and click **Reset Size** in the **Journal Volumes** section.

Recover production from a copy

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.

When selecting the image, you have the following options:

- **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:** When selecting an image from the list, 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:** Allows you to perform a customized search.
 - **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, 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.

When defining the testing network, you can:

- **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 *group sets*.
 - **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 66. Not relevant for *group sets*.
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.

- To roll back all of the writes made to the copy during image access, click **Undo Writes**.
- To write directly to the copy storage click **Direct Access**. Any changes made to the copy storage in direct access mode cannot be automatically undone, because when an image is directly accessed, the journal at the copy

is deleted. However, direct access mode does not impose a limit to the amount of data that you can write to a copy storage.

- Once image access is enabled, click **Recover Production**.

Results

- Host access to storage is blocked.
- 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.

Monitor recovery activities

Recovery Activity Reports display each of the steps in a recovery activity, the time that each step took, and the completion status of each step.

Before you begin

Every vRPA clock must be synchronized within its time zone to prevent inconsistencies in the report timestamps.

Procedure

1. Click the **Reports** tab.
2. Expand the consistency group tree in the left pane, and select the group that you want to monitor.

The **Recovery Activities** of the selected group are displayed in the right pane.

3. Select the recovery activity that you want to monitor.
 - 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 the selected activity is displayed. Expand the report to view each step. Up to 10 reports per consistency group are displayed.

APPENDIX A

Managing RecoverPoint for VMs

This section describes how to use the **RecoverPoint for VMs vSphere plugin** to manage the components of the RecoverPoint for VMs system, after initial system configuration.

- [Managing system licenses](#).....56
- [Managing system registration](#).....56
- [Managing system support](#).....57
- [Managing system component registration](#).....58
- [Managing system protection policy templates](#).....61
- [Managing virtual machines](#).....62
- [Managing consistency groups](#).....62
- [Managing copies](#).....65
- [Managing group sets](#).....68

Managing system licenses

You can remove a RecoverPoint for VMs license from the system or add a new license.

Before you begin

- For a detailed description of RecoverPoint for VMs licensing, see [RecoverPoint for VMs licensing](#) on page 73.
- To add a new license, you must first [Create your license files](#) on page 10.

Procedure

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration > vCenter Servers > Licensing**.
2. Modify your system license configuration:
 - To remove an existing license, select the license file and click **Remove**.
 - To add a new license file, click **Add**. The **Getting Started Wizard** is displayed to guide you through the process. Follow the instructions in [License and register RecoverPoint for VMs](#) on page 12 to add the licence to the system, and register the system.

Managing system registration

Register your RecoverPoint for VMs system whenever you complete a RecoverPoint system installation, connect vRPA clusters in a RecoverPoint system, or upgrade a RecoverPoint system.

Before you begin

A permanent RecoverPoint for VMs license must exist in the system, see [Managing system licenses](#) on page 56. System registration does not work with a temporary license.

Procedure

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Administration > vRPA Clusters > Support**.
2. Select a vRPA cluster.
3. In the **Registration** widget, enter the new registration settings for your vRPA cluster:
 - **Connect in method:** The method that is used to allow remote connectivity to the RecoverPoint environment. Enabling this feature is recommended as it enables secure access to the RecoverPoint environment to gather logs and resolve issues as quickly as possible. 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. When 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 *Secure Remote Services Gateway Operation Guide* for further instructions on Config Tool usage. 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.

- **Location:** The city, state, and country where the customer is located.
- **Sales order number:** 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.
- **Activity type:** The kind of activity you are performing (upgrade, installation).
- **Resource performing this upgrade/installation:** The role of the person performing this upgrade or installation activity.
- **Connect home method:** The method that is used to send configuration reports and alerts to Dell EMC. Enabling this feature is recommended as it allows Dell EMC to pro-actively address issues within the RecoverPoint environment, should they arise.

4. Click **Register**.

Results

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.

After you finish

If your company does not have outside connectivity, export the registration information to a CSV file and register by email or phone, as described in [Register RecoverPoint by email or phone](#) on page 75.

Managing system support

Before you begin

- A permanent RecoverPoint for VMs license must exist in the system, see [Managing system licenses](#) on page 56. System 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.
- 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 **RecoverPoint for VMs vSphere plug-in**, select **Administration > vRPA Clusters > Support**.
2. Select a vRPA cluster.
3. Select **Enable pre-emptive support for RecoverPoint for VMs** to provide communication between the RecoverPoint for VMs system and the System Reports database.
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.

- To transfer system notifications through the FTPS server, in the **Transfer Method** section, select **FTPS**.
- To transfer system notifications through the Secure Remote Services gateway, in the **Transfer Method** section, select **ESRS**. In the **ESRS gateway IP address** field, specify the IP address of the Secure Remote Services gateway in IPv4 format..

5. Click **Test Connectivity**.

After you finish

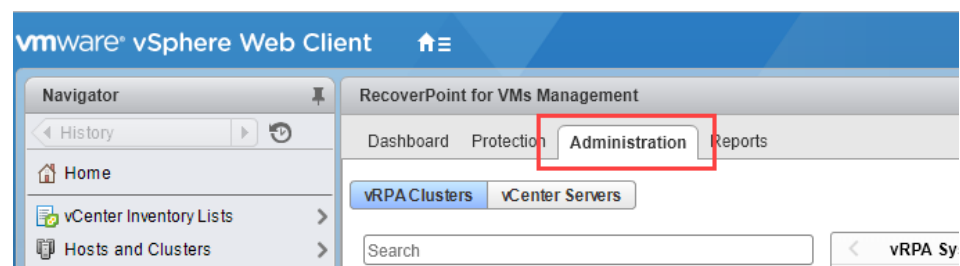
Wait 10 minutes. Then, click **Dashboard > Events Log** and look for event 1020: "Failed to send system report".

- If this event does not appear in the **Events Log**, the system notifications mechanism is correctly configured.
- If you do receive event 1020: Failed 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.

Managing system component registration

This section describes how to manage the registration of the components of your RecoverPoint for VMs system, after the system has already been configured.

After initial system configuration, manage the system configuration through the **RecoverPoint for VMs vSphere plugin > Administration** tab. Select a vRPA cluster to display the management options for that vRPA cluster.



For a detailed description of how to deploy the RecoverPoint for VMs system, see the *RecoverPoint for Virtual Machines Installation and Deployment Guide*..

For a detailed description of how to configure the RecoverPoint for VMs system, see [Before you begin](#) on page 9.

Managing vCenter Server registration

Registers the vCenter Servers used to manage your production VMs and copy VMs, at a vRPA cluster.

Before you begin

- All vCenters that manage production VMs and copy VMs must be registered at the relevant vRPA cluster before you protect VMs.
- When a vCenter is registered, all ESX clusters hosted by the vCenter are automatically registered, and a splitter is installed on all ESXs in the cluster.

- 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 **RecoverPoint for VMs vSphere plug-in**, select the **Administration** tab.
2. Access the vCenter registration information:
 - 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 vRPA 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 **Add....**
3. Click **OK**.

Results

The specified vCenter Server is registered at the specified vRPA cluster. All ESX clusters hosted by the vCenter are automatically registered with the specified vRPA cluster, a splitter is installed on all ESXs in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

Managing ESX cluster registration

Registers the ESX cluster of a production VM or copy VM, at a vRPA cluster.

By default, ESX clusters are automatically registered in RecoverPoint for VMs during VM protection and copy addition. Use this procedure to register ESX clusters in the rare case that the system cannot automatically register an ESX cluster.

Procedure

1. In the **RecoverPoint for VMs vSphere plug-in**, select **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 specified vRPA cluster.

Note

When an ESX cluster of an unregistered vCenter Server is registered with a vRPA cluster, a splitter is installed on all ESXs in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

Managing journal datastore registration

Register the datastores that are to contain the history of the data that you want to protect, at each vRPA cluster. Up to 15 shared datastores of ESX clusters running vRPAs are automatically registered in RecoverPoint for VMs. Use this procedure to register a datastore in the rare case that a datastore that you need is not automatically registered.

NOTICE

When you [Protect a virtual machine in a consistency group](#) on page 16, the **Protect VMs Wizard** will automatically select a datastore from the list of registered datastores, unless you specify a specific registered datastore to use. RecoverPoint for VMs will attempt to create the journal on the selected datastore. If for any reason journal creation fails, the system will attempt to create the journal on a different registered datastore.

Procedure

1. Select **Administration > vRPA Clusters > Related Objects**.
2. Select the vRPA cluster at which you want to register datastores, and click **Add...** under the **Journal Datastores** widget.
The **Register Journal Datastores** dialog box is displayed.
3. In the **Register Journal 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 datastore is registered at the specified vRPA cluster.

Managing external host registration

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**.

Managing system protection policy templates

You can add a new protection policy template, or edit, import, or remove 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 virtual machines

This section describes how to manage the protection of virtual machines, after they are initially protected.

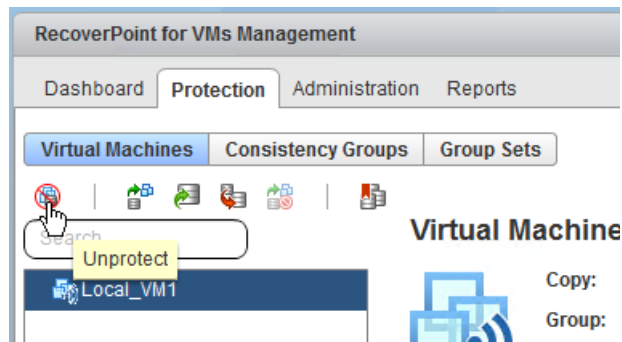
After initial protection, virtual machines are managed through the **RecoverPoint for VMs vSphere plugin > Protection > Virtual Machines** tab. Select a virtual machine to display the management options for that machine. For a detailed description of how to protect virtual machines, see [Protecting VMs](#) on page 15.

Stop protecting a virtual machine

Unprotect a VM to stop replication and remove it from its consistency group.

Procedure

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection > Virtual Machines**.
2. Select the production VM that you want to stop protecting.
3. Click the **Unprotect** button to the top of the screen:



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.

Managing consistency groups

This section describes how to manage consistency groups, after they are created.

After initial creation, consistency groups are managed through the **RecoverPoint for VMs vSphere plugin > Protection > Consistency Groups** tab. Select a consistency group to display the management options for that group. For a detailed description of consistency groups and how to create them, see [Protecting VMs](#) on page 15.

Disabling or enabling a consistency group

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

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:

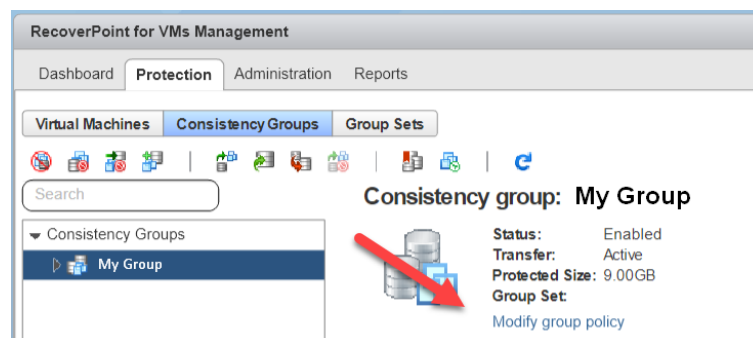


Managing the protection policies of groups

Change the protection policies of groups.

Procedure

1. In the **RecoverPoint for vSphere plugin**, select **Protection** > **Consistency Groups**.
2. Expand the list of consistency groups, and select the consistency group whose policies you want to change.
3. Click the **Modify group policy** link:



Modify the policy settings as required:

- **Name:** The name of the consistency group.
- **Primary vRPA:** 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.

NOTICE

If your vRPA cluster is the only vRPA cluster in the system, it is a single point of failure in cases of disaster. Consider adding additional vRPAs to this cluster to ensure high availability.

- **Priority:** Only relevant for remote replication when two or more consistency groups are using the same Primary vRPA.
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.

4. Expand the group, select a group copy and click the **Modify link policy** link to edit the protection policy of the link between the production and the selected copy:

- **Replication Mode**

- **Dynamic by Latency:**

Note

Only relevant for synchronous replication mode.

Default =disabled

When Enabled, RecoverPoint for VMs alternates between synchronous and asynchronous replication modes, as necessary, according to latency conditions.

- **Start async replication above:** When the specified limit is reached, RecoverPoint for VMs automatically starts replicating asynchronously
- **Resume sync replication below:** When the specified limit is reached, RecoverPoint goes back to replicating synchronously.

- **Dynamic by Throughput**

Note

Only relevant for synchronous replication mode.

Default = disabled

When enabled, RecoverPoint for VMs alternates between synchronous and asynchronous replication modes, as necessary, according to throughput conditions.

- **Start async replication above:** When the specified limit is reached, RecoverPoint for VMs automatically starts replicating asynchronously
- **Resume sync replication below :** When the specified limit is reached, RecoverPoint goes back to replicating synchronously.

- **RPO:** Defines the maximum lag allowed on a link, and is set manually in MB, GB, writes, seconds, minutes, or hours.

- **Compression:** Only relevant for asynchronous remote replication.

Default = enabled

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.

Enabling and disabling compression causes a short pause in transfer and a short initialization.

- **Enable Deduplication**

Note

Only relevant for asynchronous remote replication.

Default =enabled

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.

Enabling and disabling deduplication causes a short pause in transfer and a short initialization.

- **Snapshot Granularity**

For local replication: default = *fixed per second*.

For remote replication: default = *dynamic*.

- **Fixed per write:** Creates a snapshot for every write operation.
- **Fixed per second:** Creates one snapshot per second. Use this for local replication.
- **Dynamic:** The system determines the snapshot granularity according to available resources. Use this for remote replication.

5. Click **OK**.

Results

The group protection policies are updated.

Managing copies

This section describes how to manage copies, after they are initially created.

After initial creation, copies are managed through the **RecoverPoint for VMs vSphere plugin > Protection > Consistency Groups** tab. Select a copy to display the management options for that copy. For a detailed description of copies and how to create them, see [Protecting VMs](#) on page 15.

Modifying copy policies

To edit the protection policy of a consistency group copy:

Procedure

1. In the **vSphere Web Client** home page, go to the **RecoverPoint for VMs vSphere plugin > Protection** tab and select **Consistency Groups**.
2. Expand the list of consistency groups and select the consistency group whose copy protection policy you want to modify.
3. Expand the consistency group and select the copy whose protection policy you want to modify.
4. To change the copy protection policy, click the **Modify 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

Use the following procedure to automatically associate the VM network adapters (vNICs) of a copy VM with specific port groups upon failover. Failover networks can be configured during VM protection, through the **Protect VMs Wizard**, or after VM protection, using this procedure.

Configured failover networks are made available for selection when [Fail over to a copy and fail back to production](#) on page 50, and when [Test a copy](#) on page 48.

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

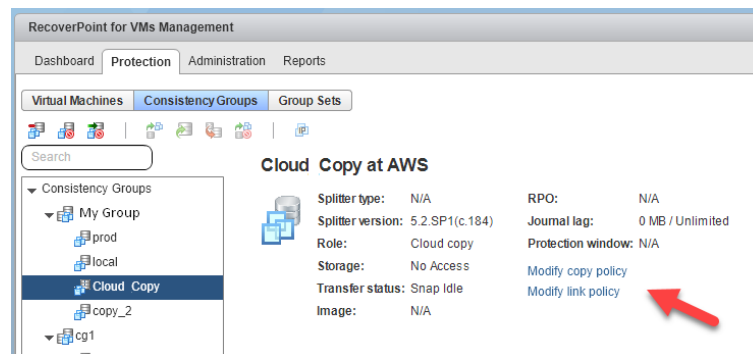
To use these settings, select *preconfigured failover networks* when defining the testing network for [Test a copy](#) on page 48 and when defining the **Target Network** before [Fail over to a copy and fail back to production](#) on page 50.

Managing the protection policies of cloud copies

Modify the protection policies of cloud copies.

Procedure

1. In the **RecoverPoint for VMs vSphere plug-in**, select **Protection > Consistency Groups**.
2. Expand the list of consistency groups.
3. Expand the consistency group whose copy protection policies you want to edit.
4. Select the cloud copy.
5. Click the **Modify copy policy** link.



Edit the copy policy settings, as required:

- **Copy Name:** Default = *copy<num>*
The name of the cloud copy. Best practice is to differentiate the cloud copy name from the production copy name.
- **Retention Policy:** Default = *5 days*
Defines the period of time that snapshots will be retained in the cloud. Snapshots can be retained for a minimum of 1 day, and a maximum of 90 days. Once every 24 hours, a temporary **Retention Service** EC2 instance is launched in AWS to consolidate the snapshots of every copy whose retention policy has expired.

6. Click the **Modify link policy** link, and edit the settings, as required:

- **RPO:** Default = *90 minutes*
The maximum data lag that is required between the production copy and the latest snapshot uploaded to the S3 bucket. If the specified RPO is exceeded,

a warning is displayed in the [The RecoverPoint for VMs Dashboard](#) on page 26. RPO can be defined in **Minutes, Hours, Or Days**.

Note

Specify an **RPO** value that is higher than the specified snap replication **Interval**. Best practice is to specify an **RPO** value that is 1.5 times the specified snap replication **Interval**. For example, if you require an RPO of *1 hour*, specify a snap replication interval of *90 minutes*.

- **Snap Replication:** Default = *Periodic* at *1 hour* intervals
Sets the periodic **Interval** between snapshots in **Minutes, Hours, Or Days**. The minimum interval value is *15 minutes* and the maximum interval value is *7 days*. A new snapshot starts after the specified interval has passed since the previous snapshot was started. If the time interval has passed and the previous snapshot is incomplete, the next snapshot will start as soon as the previous one has completed.
-

Note

Specify a snap replication **Interval** value that is lower than the specified **RPO** value. Best practice is to specify an **RPO** value that is 1.5 times the specified snap replication **Interval**. For example, if you require an RPO of *1 hour*, specify a snap replication interval of *90 minutes*.

After you finish

See [Managing the protection policies of groups](#) on page 63 for protection policies specific to groups.

Managing group sets

This section describes how to manage group sets, after they are created.

After initial creation, group sets are managed through the **RecoverPoint for VMs vSphere plugin > Protection > Group Sets** tab. Select a group set to display the management options for that group set. For a detailed description of group sets and how to create them, see [Create a group set](#) on page 35.

Modifying a group set

Procedure

1. Click **Protection > Group Sets**.
2. Select the group set that you want to modify.
3. Click the **Edit Group Set** icon:



4. In the **Edit Group Set** dialog box.
 - a. If required, modify the name of the group set.
 - b. Select the consistency groups to add or remove from the group set.
 - c. Enable or disable parallel bookmarking using the **Enable Parallel Bookmarking** and **Frequency** fields. 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.

5. Click **OK**.

Removing a group set

Procedure

1. Click **Protection > Group Sets**.
2. Select the group set to remove.
3. Click the **Remove Group Set** icon:



APPENDIX B

Troubleshooting

Use the following information, features and tools to troubleshoot your RecoverPoint for VMs system.

- [Finding the vRPA cluster management IP](#) 72
- [Collecting system information](#) 72
- [Collecting RecoverPoint for VMs splitter logs](#) 73
- [Recovering from a cluster disaster](#) 73
- [RecoverPoint for VMs licensing](#) 73
- [Register RecoverPoint by email or phone](#) 75
- [Creating VMkernel ports](#) 75
- [Load balancing](#) 76
- [Manual copy VM network configuration \(RE-IP\)](#) 77
- [Copy VM network configuration guidelines](#) 79
- [Changing the network adapter configuration of a protected VM](#) 84

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.

Collecting system information

Collecting system information is only relevant in support cases, and should only be performed when instructed to do so by Customer Support. Use the **RecoverPoint for VMs vSphere plugin** to collect system information, and retrieve it from the vRPA cluster, or a specified FTP server.

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.
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. Select the **vRPA Clusters** from which to collect the logs.
5. Optionally, select **Include core files**.
Core files might be large. Subsequently, including these files in the collection process could substantially increase collection time.
6. 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.
7. Optionally, select **Copy the output file(s) to an FTP server** and define the FTP server settings.
8. Click **Start**.

Results

Be patient. The collection process can take awhile, depending on the amount of data being collected. 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, click the relevant link in the **Output Files on vRPA Cluster** column to retrieve the vRPA cluster log files:

1. At the login prompt, type `boxmgmt` as the **User Name** and enter the **Password** for the `boxmgmt` user (Default password is `boxmgmt`).
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 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

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.

RecoverPoint for VMs licensing

RecoverPoint for VMs supports two types of licensing models; VM-based licensing and socket-based licensing .

VM-based licensing

With VM-based licensing, licenses are based on the number of supported VMs per vCenter server. Only production VMs are counted in the number of supported VMs 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 in the RecoverPoint for VMs Deployer UI. Refer to the *RecoverPoint for VMs Installation and Deployment Guide* for more information.

When you reach the maximum number of VMs that the license supports for each vCenter server, you cannot protect new VMs or enable disabled consistency groups. However, replication of existing VMs and consistency groups continues.

Failover has no effect on the license.

Socket-based licensing

With socket-based licensing, licenses are based on the number of physical CPU sockets in the ESXi servers that host the production VMs. A VM does not 'belong' to a specific socket.

When you reach the maximum number of sockets that the license supports for each vCenter server, you cannot protect new VMs or enable disabled consistency groups. However, replication of existing VMs and consistency groups continues.

As with VM-based licensing, failover does not affect the socket-based license. However, vMotion of production VMs does affect the license and may cause a license violation due to an increase in the number of sockets being used. ESXi servers that host the production VMs are the ones that count in a socket-based license. To avoid license violations, it is a best practice to license all ESXi servers of the ESXi cluster.

Adding a socket-based license to a system with VM-based licenses

When using VM-based licensing, license capacity is measured by the number of VMs. For example, when you view the license capacity in the UI, it may be listed as:

```
Capacity = 30 VMs
```

When using socket-based licensing, license capacity is measured by the number of sockets. For example, the license capacity may be listed as:

```
Capacity = 2 sockets
```

When a socket-based license is installed on a RecoverPoint for VMs system that has VM-based licenses, the system automatically converts VM-based licenses to socket-based licenses at a ratio of 15 VMs per socket. In this case, the license capacity would be listed as:

```
Capacity = 30 VMs (2 sockets)
```

In cases where the ratio does not result in an even conversion, the value is rounded up. For example:

```
Capacity = 31 VMs (3 sockets)
```

Since licenses are applied per vCenter, and not per vRPA cluster, multiple vRPA clusters with VMs or CPU sockets may count towards the same license.

License subscriptions

VM- and socket-based licenses may be installed as subscriptions. Unlike a permanent license, a subscription license has a start date and an end date. The system sends an alert beginning 30 days before license expiration to indicate the number of days remaining. Subscription and permanent licenses may coexist.

You can install a subscription license before its start date. It automatically becomes active on the start date.

Register RecoverPoint by email or phone

If your company is without external connectivity, and you cannot register your RecoverPoint for VMs system online, you can also register by phone.

- 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. Gather the required information.
 - Download the post-deployment form:
 - a. Access <https://support.emc.com>
 - b. Search for the term *Post-Deployment Form*
 - c. Download and fill out the RecoverPoint and RecoverPoint for VMs Post-Deployment Form, for every vRPA cluster
 - Export the RecoverPoint registration information, for every vRPA cluster:
 - a. Select **Administration > vRPA Clusters**.
 - b. Select the vRPA cluster for which you want to export a post-deployment form, and then click **Support**.
 - c. In the Registration pane, click the **Export to CSV** button and save the file to the computer.
2. 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).

Creating VMkernel ports

If before clicking **Protect** in the **Protect VMs Wizard** you received a warning regarding a potential communications problem, and after clicking **Protect**, transfer for the consistency group does not eventually reach the **Active** status, you may need to create VMkernel ports for all ESXi hosts in the cluster.

Procedure

1. Select **Administration > vRPA Clusters > ESX Clusters**, and click the **Settings** icon for an ESXi cluster.

2. In the **Create VMkernel Ports** dialog box, specify the settings, including a range of available IPs, for creating VMkernel ports for all ESXi hosts in the cluster.

3. Click **OK**

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.
- Periodically, to ensure that the system is always handling distributing loads evenly. A script can be created to periodically perform load balancing.

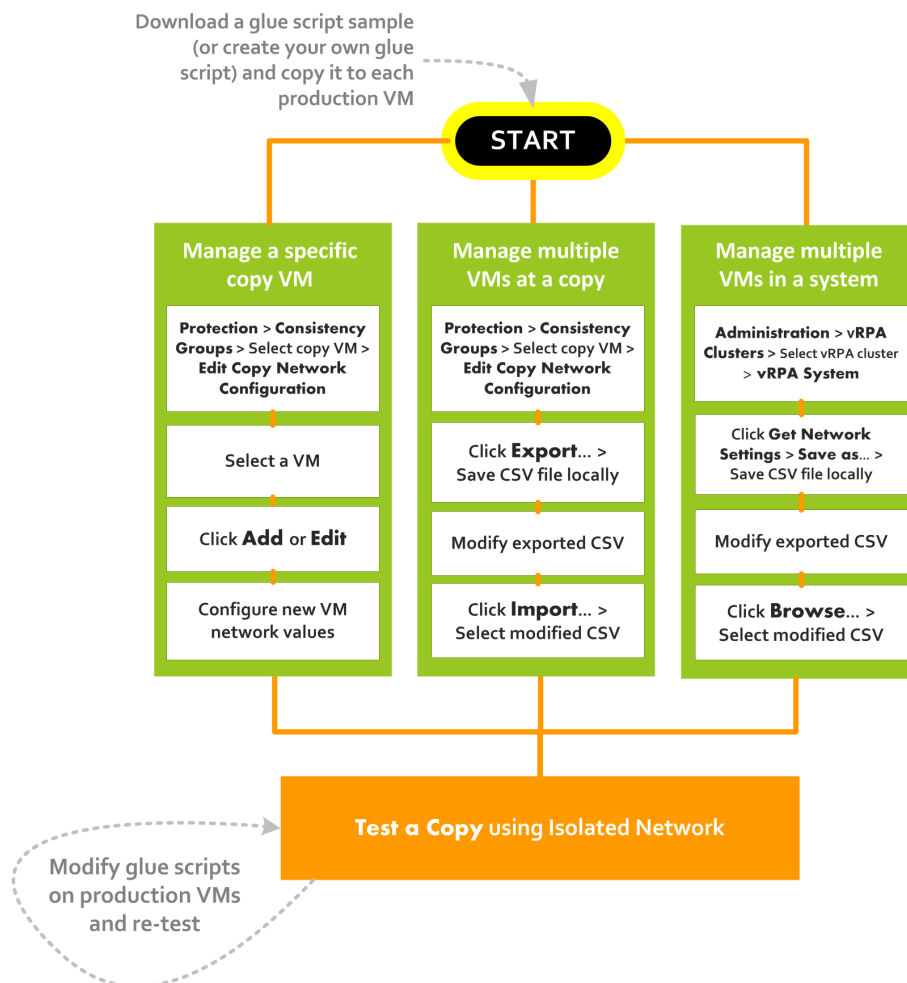
Procedure

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 ?`

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 39.

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 `deployPkg` has been manually installed. See *VMware KB article 2075048* for detailed information on how to install `deployPkg`.

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 2 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 Rename glue_script_w in_this_bat_file.txt to glue_script_w in.bat
glue_script_win.py	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
glue_script_rhel.bash	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
glue_script_rhel.py	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 39.
 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 39
- [Manual copy VM network configuration \(RE-IP\)](#) on page 77

Table 3 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 example.company.com.
(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 (;).

Table 3 Virtual machine network settings available through the GUI (continued)

Setting	Description	Guidelines
		<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) DNS Suffix(s)	The global settings of the suffixes for the DNS servers of all adapters on both Windows and Linux virtual machines.	<ul style="list-style-type: none"> Customizable. Can be left blank. 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 77.
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 IDX 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)

Table 3 Virtual machine network settings available through the GUI (continued)

Setting	Description	Guidelines
		<p>of the adapter, and NOT the NIC number.</p> <p>For example, eth0 = 1, eth1 = 2. If you have eth2 and eth3, and want to update the network settings of the second one, set <code>Adapter ID = 2</code>.</p>
(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.

Table 3 Virtual machine network settings available through the GUI (continued)

Setting	Description	Guidelines
		<ul style="list-style-type: none"> 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.
(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.

Table 3 Virtual machine network settings available through the GUI (continued)

Setting	Description	Guidelines
		<ul style="list-style-type: none"> 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 4 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	<ul style="list-style-type: none"> Customizable.

Table 4 Network settings only available through the CSV file (continued)

Setting	Description	Guidelines
	adapters (NICs) in the virtual machine settings of the vCenter web client.	<ul style="list-style-type: none"> • 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 VM

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 39. If you are already using glue scripts, perform [Manual copy VM network configuration \(RE-IP\)](#) on page 77.

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 [Test a copy](#) on page 48.

GLOSSARY

A

account ID	Part of a customer's account settings. The account ID is a unique identifier of a RecoverPoint for VMs customer account.
account settings	<p>The details that comprise a RecoverPoint customer account. The account settings are comprised of:</p> <ul style="list-style-type: none">• One account ID• One installation ID per installed RecoverPoint for VMs system• One software serial ID per vRPA cluster• A RecoverPoint for VMs license key
ACK	<p>To send an acknowledgment, which is a signal to confirm receipt of data. In synchronous replication, an acknowledgment must be sent before the next host write can be made to the production storage.</p> <p>See also acknowledgment</p>
action regulation	<p>A state that a copy can be placed into when the system quickly changes between two states for a set period of time.</p> <p>See also control action regulation</p>
activated license	A RecoverPoint license for which an activation key has been defined. An activated license can be a temporary license or a permanent license.
activation key	<p>A code that is used to activate or re-activate the RecoverPoint license, and is generated per installation ID. You get this code when you type the account ID and RecoverPoint license key into the RecoverPoint licensing server, and then request an activation code.</p> <p>See also activation code</p>
alert rule	<p>The rules that specify the events on which the system alert mechanism should send notifications, and the desired sending frequency. For each alert rule, you can specify the following:</p> <ul style="list-style-type: none">• Event topic (Site, RPA, Group, Splitter, or Management)• Event level (Info, Warning, or Error)• Event scope (Detailed or Advanced)• Event type (Immediate, Daily)• Email address (of alert recipient)
array throttling	The act of controlling the write activity. RecoverPoint's array throttling mechanism enables users to limit the storage read-rate of RPAs in a RecoverPoint cluster, allowing the storage to handle the I/O rate during initialization.

asynchronous replication A replication mode that uses ACKs from the local RPA to confirm data transfer.

See also async, async replication

B

bookmark A label that is applied to a snapshot (PIT) so that the snapshot can be explicitly called (identified) during recovery processes (for example, during image access).

bottleneck One of a predefined set of potential types of RecoverPoint performance-related problems resulting from a high write-rate (high-load), poor storage or journal performance, or problems in communication between RPAs.

C

call home events A proactive online service capability that is built into RPAs to enable them to continuously monitor their own health, and the health of the RecoverPoint system, using a pre-defined set of event-filtering rules. If a serious problem arises, the call home event mechanism automatically opens a service request with Customer Support. The service request enables Customer Support to proactively engage the relevant personnel, start working with the relevant parties, or use a configured Secure Remote Services gateway to resolve the issue as soon as possible.

CLI The RecoverPoint Command Line Interface. Using the RecoverPoint CLI, management and monitoring activities can be run textually, interactively, or through scripts. For information about the command line interface, see the *RecoverPoint Command Line Interface Reference Guide*.

cluster control The process that manages an RPA cluster.

cluster management IP A virtual, floating IP address assigned to the vRPA that is currently active (runs the cluster control).

compression The process of encoding data to reduce its size. RecoverPoint uses lossless compression (compression using a technique that preserves the entire content of the original data, and from which the original data can be reconstructed).

There are two kinds of compression in RecoverPoint:

- Journal compression – Used to compress the data in the journal of each copy, and configured through each copy's general settings.
- WAN compression – Used to compress consistency group data before transferring it over the WAN, and configured through the consistency group bandwidth reduction policy.

consistency group A logical entity that constitutes a container for virtual machines and all their copies, used to replicate virtual machine application data to a consistent point in time.

copy A logical entity that constitutes all of the data that is copied and used to protect the production data, at a specific location. A copy includes the replicated data, and a journal to hold the copy history.

current image The image that is either currently being distributed to the copy storage (typically), or that has been distributed to the copy storage (when no more writes are received by the host and all of the snapshots in the journal have been distributed to the storage). Also known as: latest image, latest snapshot, current image, current snapshot, current PIT

D

delta marking The process of writing marking information that is used when the link between the production and a copy of a consistency group is down, and host writes cannot be saved as snapshots in the copy journal. In this case, a short initialization process is triggered to synchronize the copy volumes with their production volumes, using the marking information. At least one RPA at the production RPA cluster must be available for marking to occur.

See also marking on RPA, hotspots, write-folding

direct access An image access mode in which hosts write directly to the copy storage. These changes cannot be automatically undone, because when an image is directly accessed, the journal at the copy is deleted. However, direct access mode does not impose a limit to the amount of data that you can write to a copy storage.

E

ESRS EMC Secure Remote Support (ESRS) is a server and set of services that allow customer support to remotely access vRPAs to collect system information and provide pre-emptive support.

event A notification that a change has occurred in the state of a managed device or component. In some cases, the change indicates an error or warning condition for the device or component. Multiple events can occur simultaneously on a single monitored device or service module. A single incident can generate events across multiple system components. Events in RecoverPoint have a level (*Info, Warning, Error*), scope (*Normal, Detailed, Advanced*), and a topic (*All, Cluster, RPA, Group, Splitter, Management*).

F

failback Reversal of failover, that is, when the production and copy resume their original roles in the replication process after a failover.

failover The process of changing the replication direction between the production and the copy.

flipover When a non-preferred RPA (at the same site) takes over replication responsibilities and is used to transfer the data of a consistency group, instead of the consistency group's preferred RPA.
Also known as: switchover

full sweep An initialization process that is performed on all of the volumes in a consistency group.

G

- group set** 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.
- guest operating system** The operating system running on a virtual machine. Each virtual machine can run operating systems such as Windows, Linux, Solaris, or Netware. Applications are run on the virtual machine unmodified.

H

- high-load** A system state that indicates resource depletion during replication. There are two kinds of high-loads in RecoverPoint:
- Permanent high-loads – RecoverPoint stops and waits for a user action in order to come out of high-load (not relevant in the RecoverPoint for VMs cloud solution).
 - Temporary high-load – RecoverPoint tries to recover from the high-load and keeps trying until the condition that triggered the high-load changes.
- hot spot** A single disk location to which multiple writes have been written.

I

- image** All of the snapshots that, together, constitute a specific point in time. An image consists of snapshots in the journal, and the data that has already been transferred to the copy storage.
- image access** A user-triggered operation that is performed on a copy journal to enable read/write access to a selected PIT at a copy.
- image access log** The dedicated area of the copy journal that is used to hold all of the information of the writes to the copy that occur while image access is enabled.
- initialization** The process that is used to synchronize the data of the copy volumes with their corresponding production volumes, and ensure consistency. Generally, all synchronization processes are called Initialization.
- See also** synchronization
- initialization snapshot** The first consistent snapshot in a copy journal. Whenever an initialization process completes, this initialization snapshot is created.
- Install Base** An database that is used to manage and support equipment installed at customer sites. When new equipment is installed, it is important to update the database with the new information and IDs of the newly installed equipment. The information in the installation base is also used to enable the RecoverPoint system report mechanism.

J

journal	One or more volumes that are dedicated on the storage at each copy in a system, to hold the production data history. Journals are defined per copy, and can consist of multiple volumes.
journal history	The PIT images (or snapshots) in the journal.
journal lag	The amount of data (represented by snapshots) in the copy journal that has not yet been distributed to the copy storage.
journal loss	The loss of an entire copy journal and all of the data in it.
journal volume	The volumes that make up a journal. Each journal can be composed of one or more journal volumes.

L

lag	The current RPO of the consistency group. In RecoverPoint, lag starts being measured when a write made by the production host reaches the local RPA, and stops being measured when the write reaches either the target RPA, or the target journal (depending on the <code>transfer_by_non_preferred</code> parameter of the <code>set_policy</code> CLI command).
latency	The number of milliseconds or microseconds that it takes for data to get from the local vRPA to the vRPA or journal at the remote vRPA cluster.
latest image	The most current image available in the journal. The last snapshot that was created at the production and transferred to the journal at the copy. See also latest snapshot
link	The communication pipe through which data is transferred between the production and a copy. See also pipe
load settings	A user-triggered operation that is performed to recreate the system configuration. This operation loads a configuration file that is created from the output of the <code>save_settings</code> CLI command, which displays the current system settings in a configuration settings file.
local replication	A replication configuration that continuously captures or tracks modified data and replicates them to a copy at the local RPA cluster, storing changes independent of the primary data, and enabling recovery from any point in the past. Local replication provides fine granularities and restorations to any point in time. See also local recovery

M

maintenance mode	<p>The RecoverPoint for VMs vSphere system enters maintenance mode when undergoing any of the following operations:</p> <ul style="list-style-type: none"> • minor version upgrade • major version upgrade • replacing an RPA in an existing cluster • adding new RPAs to existing clusters • converting an environment to a RecoverPoint/SE or RecoverPoint configuration • modifying system settings <p>In maintenance mode, RecoverPoint for VMs can only monitor the system; user-initiated capabilities are disabled.</p>
management default gateway	The default gateway IP address assigned to the RPA LAN interface and used for all non-replication IP communication.
management interface	<p>An Ethernet interface that is used primarily for configuration maintenance and monitoring. The management interface is usually accessed through a virtual site-management IP address using the CLI or the GUI.</p> <p>Also known as: LAN interface.</p>
management IP	<p>The IP address assigned to the LAN interface in order to define the management interface network.</p> <p>Also known as: LAN IP</p>
management subnet mask	A network subnet mask assigned to the LAN interface in order to define the management interface network.
manual snapshot consolidation	<p>The process of consolidating snapshots manually through the CLI <code>consolidate_snapshots</code> command. This option is useful when you want to specify a specific PIT to consolidate to. For example, a script can be run once a day at a specific time to create a bookmark, and set the specific bookmark as the starting snapshot.</p>
marking mode	<p>A system state that signifies the consistency group is enabled, the splitter is replicating to the RPAs, but the RPAs are unable to transfer to the copy journal (for example, because the link is closed). When the remote RPA cluster is available again, only the disk segments or blocks that have changed are synchronized, and both transfer and replication are resumed. At least one RPA at the production RPA cluster must be available for marking to occur.</p> <p>See also marking mode, delta marking mode</p>
metadata	<p>Metadata can be defined as either data about other data or data about the container of other data. In RecoverPoint, metadata lists are used in the marking process (along with bitmaps) to ensure data consistency, and these metadata lists contain the following information:</p> <ul style="list-style-type: none"> • The GUID of the source LUN • The offset (for example, location) of the data in the target LUN • The length (for example, size) of the data

MIB Management Information Base. A (virtual) database used to manage the devices in a communications network (routers, switches, and so forth), and tune the network according to real-time requirements.

O

oldest image The snapshot that was taken the longest time ago, and is still available in the journal.

P

parallel bookmark A RecoverPoint feature that enables you to apply a bookmark with the same name to a single PIT across multiple consistency groups, while at the same time marking a consistent PIT (with the same name) across all specified consistency groups.

PIT Point In Time. A fully usable copy of a defined collection of data that contains an image of the data as it appeared at a single instant in time.

See also snapshot

predefined user The RPA is shipped with the following predefined users, each has its own initial password (which can be modified) and set of permissions:

- security-admin
- admin
- boxmgmt
- monitor
- webdownload

See also preconfigured user

preferred RPA An RPA that whenever possible, handles replication for the consistency group. If an error occurs in the preferred RPA, in most cases, another RPA at the same RPA cluster handles replication.

See also primary RPA, preferred primary RPA

pre-replication image An image that represents the state of the production data before initialization.

See also initialization snapshot

private default gateway In local replication, the IP address that the RPA uses to route any private network communications.

See also WAN default gateway

private IP In local replication, the IP address assigned to the WAN interface of an RPA.

See also WAN IP

private network	The network that is created when a single RPA cluster is deployed (in local replication). The private network is used for inter-cluster communication (exchange state of RPA nodes and cluster LEP (cluster leader arbitration)).
private subnet mask	In local replication, the IP assigned to the WAN interface to define the management interface network. See also WAN subnet mask
production	The data that is being replicated and protected. See also production source, protected copy

R

reboot regulation	A state of regulation that allows the system to detach an RPA from its RPA cluster in the event of frequent unexplained reboots or internal failures.
RecoverPoint for VMs plug-in	The vSphere web client user interface that is used for managing VM replication. The plug-in is installed automatically after the vRPA cluster has been installed.
RecoverPoint for VMs splitter	Proprietary software that is installed on every ESXi host in an ESXi cluster that is involved in RecoverPoint for VMs replication or running virtual RPAs. The splitter splits every write to the VMDK and sends a copy of the write to the vRPA and then to the designated storage volumes. It is installed automatically after you register the ESXi cluster.
Recoverpoint for VMs system	One or more vRPA clusters that have been installed using the RecoverPoint for VMs Deployer.
recover production	A user-triggered disaster recovery procedure that is used to repair the production source using the copy data.
remote replication	A system configuration where data is transferred between RPA clusters. In this configuration, the RPAs, storage, and splitters exist at both the local and the remote RPA cluster. See also two-site configuration, remote recovery
replication policy	A user-specified set of parameters driven by business objectives that control system operation during replication.
replication set	A production source and the target(s) at a copy to which it replicates.
replication set volumes	All of the sources that have been added to a replication set.
repository volume	A dedicated volume that must be allocated at the production, for each RPA cluster.
resource allocation	A consistency group policy that controls the bandwidth that RecoverPoint allocates for group replication, with regard to other groups using the same preferred RPA.

RPO Recovery Point Objective. The maximum amount of data, per application, that an organization is willing to lose if there is a disaster. For example, an RPO of 5s means that if there is a disaster, RecoverPoint ensures that no more than the last 5s of data can be lost.

RTO Recovery Time Objective. The duration of time and a service level within which a business process must be restored after a disaster (or disruption) in order to avoid unacceptable consequences associated with a break in business continuity.

S

secondary RPA When using distributed consistency groups, an additional RPA that can be the preferred RPA to transfer the data of a consistency group.

See also preferred secondary RPA

shadow VM A secondary copy VM that RecoverPoint creates, configures, and manages to allow access to copy VMDK and RDM devices. A copy shadow VM has the .copy.shadow extension at the end of the virtual machine name. User action on copy shadow VMs is not supported.

short initialization An initialization process that uses marking information to re-synchronize a copy's data with its production sources. Because this initialization process uses delta markers to synchronize the copy with the production, the initialization process is much faster and more efficient. Generally occurs when restarting transfer for a consistency group after a pause in transfer.

See also short init, short resync, short resynchronization, resynchronization, resync

snapshot consolidation A manual or automatic process that consolidates the data of multiple snapshots into a single snapshot.

snapshot dilution A process that is performed on snapshots in a copy journal to hide or display a snapshot, over others, according to a pre-set system algorithm.

software serial ID The identification that is used by the install base to support equipment that is installed at customer sites using the system reporting and Secure Remote Services mechanisms. A software serial ID is supplied per RPA cluster in a system installation.

See also SSID

source The object that RecoverPoint is replicating from. For example, the source RPA or the source copy (for example, production). After failover, the source becomes the target and the target becomes the source.

splitter Proprietary software that is installed on storage subsystems that splits application writes so that they are sent to their normally designated storage and the RPA simultaneously.

synchronous replication A replication mode in which the production host application initiates a write, and then waits for an ACK from the remote RPA before initiating the next write, ensuring that the RPO value is always zero.

See also sync, sync replication

system alerts A mechanism that allows RPAs to send system events about system components in real-time, to a specified email or the system reports database, via SMTP.

system reports A mechanism that provides one-way communication between a system installation and the system reports database. This mechanism supports two types of information: system alerts and system reports.

See also SYR

system settings The output of the `save_settings` CLI command.

See also configuration file, system configuration file, configuration settings

T

target The object at the copy that the protected data is being replicating to. For example, the target RPA, or the target storage. After failover, the target becomes the source, and the source becomes the target.

temporary failover One of four disaster recovery methods that is only relevant for configurations with more than one copy. Temporary failover is user-triggered when it is not possible to run host applications at the production site, and therefore, there is a need to temporarily fail over to a specific copy to work from the copy site until the production is repaired, or you have recovered from the disaster at the production site. Temporary failovers are not possible when there is only one copy of the production data. In these cases, every failover is a permanent failover.

testing a copy A user-triggered procedure used to ensure that a copy can be used to restore data, recover from disaster, or seamlessly take over for the production.

throughput The total amount of writes made by the production hosts and received by the local RPA.

tweak parameters A configuration parameter that only Customer Support can change. Tweak parameters enable Customer Support to change the hard-coded values of specific internal settings without requiring the RecoverPoint for VMs code to be recompiled or re-loaded onto the vRPAs.

U

user authentication The method of establishing the authenticity of users in RecoverPoint. RecoverPoint provides two independent mechanisms for authenticating users:

- Appliance-based authentication
- Authentication through LDAP (Lightweight Directory Access Protocol)

user authorization The method of establishing user access permissions to RecoverPoint resources. User authorization is identical, regardless of whether RecoverPoint or an LDAP server authenticated the user.

V

volume sweep	An Initialization process that is performed on a specific replication set in a consistency group.
vRPA	The virtual RecoverPoint Appliance that manages data replication.
vRPA cluster	A group of vRPAs that work together to replicate and protect data.
vRPA data network	The IP data path network used for traffic between the vRPA and the splitter.
vRPA LAN network	The vRPA cluster management network that is used for communication between vRPAs and to other entities such as the vCenter.
vRPA WAN network	The vRPA inter-cluster network that is used for communication between vRPA clusters. A basic configuration that uses a unified topology (a single WAN + LAN network) does not require an additional gateway to communicate between sites.

W

WAN default gateway	The default gateway IP address for replication IP communications assigned to the RPA WAN interface. In local replication, the WAN default gateway is referred to as the <i>private</i> default gateway.
WAN interface	An Ethernet interface that is used primarily for the transfer of protected data to the remote vRPA cluster as well as for inter-cluster communication (management of component states, resource discovery, and cluster leader arbitration). In local replication, the WAN interface is referred to as the <i>private</i> interface.
WAN IP	The IP address assigned to the WAN interface to define the management interface network. In local replication, the WAN IP is referred to as the <i>private</i> IP.
WAN subnet mask	A network subnet mask assigned to the WAN interface to define the WAN interface network. In local replication, the WAN subnet mask is referred to as the <i>private</i> subnet mask.

