# RecoverPoint for Virtual Machines

Version 4.3

## Administrator's Guide

P/N 302-001-905

REV 04

EMC²

Published July 2016



For the most up-to-date regulatory document for your product line, go to EMC Online Support (https://support.emc.com).

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

# CONTENTS


**Tables** 5

**Preface** 7

**Chapter 1 Getting started** 9

RecoverPoint concepts ... 10
Activating your entitlements ... 10
Licensing, support, and registration ... 11
Exporting the post-deployment form to a CSV file ... 13
Registering RecoverPoint by email or phone ... 14
Registering datastores ... 14
Registering ESX clusters ... 15
Registering vCenters ... 15

**Chapter 2 Protecting your data** 17

Support for VMware actions ... 18
Protecting virtual machines ... 19
Stopping protection ... 22

**Chapter 3 Managing and monitoring the system** 25

Monitoring the system ... 26
- Reviewing recovery activities ... 26
- Monitoring consistency group replication ... 27
- Consistency group transfer states ... 27

Managing the system ... 27
- Licenses ... 27
- Support settings ... 28
- RecoverPoint for VMs registration ... 28
- Registering an external host ... 28
- Policy templates ... 29

Managing group sets ... 29
- Creating a group set ... 29
- Editing an existing group set ... 30
- Enabling a group set ... 30
- Removing a group set ... 30

Managing consistency groups ... 30
- Enabling or disabling a consistency group ... 30
- Editing group or link policies ... 31
- Creating bookmarks ... 33
- Adding a copy ... 36

Managing copies ... 38
- Editing copy policies ... 38

Managing VMs and VMDKs ... 40
- Orchestration ... 40
- Automation ... 45

Chapter 4 Testing and recovery 49

Testing and recovery ..... 50
Testing a copy ..... 50
Disabling image access ..... 52
Failing over ..... 52
Recovering production ..... 54

Chapter 5 Troubleshooting 57

Finding your vRPA cluster management IP ..... 58
Identifying a RecoverPoint for VMs system ..... 58
Recovering from a cluster disaster ..... 58
Detecting bottlenecks ..... 58
Load balancing ..... 60
System alerts ..... 60
Collecting system information ..... 61
Collecting RecoverPoint for VMs splitter logs ..... 62
Network configuration settings ..... 62

# TABLES


1 Add vCenter server ..... 16
2 KVSS syntax ..... 35
3 Glue script samples ..... 43
4 RPA and cluster bottlenecks ..... 58
5 Consistency group and link bottlenecks ..... 59
6 Network configuration file ..... 62

# Preface

As part of an effort to improve its product lines, EMC periodically releases revisions of its 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 EMC 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 EMC Online Support (https://support.emc.com) to ensure that you are using the latest version of this document.

---

### Related documentation

The following EMC publications provide additional information:

- *EMC RecoverPoint for Virtual Machines 4.3 Installation and Deployment Guide*
- *EMC RecoverPoint for Virtual Machines 4.3 Product Guide*
- *EMC RecoverPoint for Virtual Machines 4.3 Release Notes*
- *EMC RecoverPoint for Virtual Machines 4.3 Security Configuration Guide*

### Where to get help

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

#### Product information

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

#### Technical support

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

### Your comments

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

# CHAPTER 1

# Getting started

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


- RecoverPoint concepts ..... 10
- Activating your entitlements ..... 10
- Licensing, support, and registration ..... 11
- Exporting the post-deployment form to a CSV file ..... 13
- Registering RecoverPoint by email or phone ..... 14
- Registering datastores ..... 14
- Registering ESX clusters ..... 15
- Registering vCenters ..... 15

# RecoverPoint concepts

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

# Activating your entitlements

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

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

### Procedure

1. To activate your RecoverPoint for VMs entitlements, access your entitlements on emc.support.com:
   - If you have your License Authorization Code email, open it and click the **Click here** link to automatically access Powerlink Licensing on the EMC Online Support site, and search for the entitlements associated with the License Authorization Code.
   - If you do not have your License Authorization Code email but you do have your License Authorization Codes or sales order numbers, log into the EMC Online support at http://support.emc.com, and

   a. Select **Support › Service Center** from the main menu.

   b. Select **Get and Manage Licenses**.

   c. Select **RecoverPoint for Virtual Machines**.

   d. Enter the customer's License Authorization Code and click **Activate** to search for all inactive entitlements 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 your entitlement(s) and download your license file(s).

   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, enter 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, enter 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 the license is enforced for, namely, the vCenter Server ID. To find your vCenter Server ID, type **`https://<vCenterServerIP>/mob`** into your browser address bar or SSH client, and enter your credentials to log into the vCenter Server. Select **Content** › **About**. The instanceUuid is the vCenter Server (Locking) ID that your license is enforced for.

g. In the **Confirm** screen, enter the email addresse(s) of the recipients of the license file in the **Email to** field of the **Additional Email Options** section, and click **Finish**. Separate multiple email addresses with commas.

h. In the **Complete** screen, click **Save to File** to download the license file and save the file locally. The resulting license file has a `*.lic` extension and is in plain text format (can be opened in any text editor).

i. Repeat this procedure for all inactive entitlement in each License Authorization Code email.

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

# Licensing, support, and registration

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

### Before you begin

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

### Procedure

1. In the Getting Started Wizard **Welcome** screen, click **Next Add Licenses**.
2. In the **Add Licenses** screen, click **Add**. In the **Add license** dialog box, enter the location of the license file or click **Browse** to locate the file. Click **OK**. Click **Next Enable Support**.

3. To provide communication between your RecoverPoint for VMs system and the EMC System Reports database (ESRS), in the **Enable Support** screen, select **Enable pre-emptive support for RecoverPoint for VMs**.
4. Define the transfer method:
   - To transfer system notifications through an SMTP server, in the **Transfer Method** section, select **SMTP**. In the **SMTP server address** field, specify the IP address or DNS name of your dedicated SMTP server, in IPv4 format. In the **Sender address field**, specify the email address from which the system notifications will be sent. Click **OK**.
   - To transfer system notifications through RecoverPoint's built-in FTPS server, in the **Transfer Method** section, select the **FTPS radio button**. Click **OK**.
   - To transfer system notifications through an ESRS gateway, in the **Transfer Method** section, select the ESRS radio button. In the **ESRS gateway IP address** field, specify the IP address of your ESRS gateway in IPv4 format. Click **OK**.
5. Click **Test Connectivity**. Wait ten minutes. Then, create an `ssh` connection to your cluster management IP address, run the `get_events_log` command, and look for event 1020 `"Failed to send system report"`.
   - If this event does not appear in your event logs, the system notifications mechanism has been properly configured.
   - If you do receive an event 1020 `"Failed to send system report"`, check whether there is an issue with your selected method of transfer. If a problem exists, fix it, configure support, and click **Test Connectivity** again.. If the problem persists, contact EMC Customer Support.
6. Click **Next Register RecoverPoint**.
7. In the **Register RecoverPoint** screen, register or re-register each cluster in your RecoverPoint system after every RecoverPoint system installation, after connecting vRPA clusters in a RecoverPoint system, and after upgrading a RecoverPoint system.
   a. Click **Edit Settings...** to display the **Update Post-Deployment Form Details** dialog box.
   b. In the **Update Post-Deployment Form Details** dialog box, update the form information.

| Option | Description |
|---|---|
| Company name | The name of your company |
| Connect in method | The method used to allow remote connectivity to your RecoverPoint environment. Enabling this feature is recommended as it allows secure access to your RecoverPoint environment to gather logs and resolve issues as expeditiously as possible.<br>If you already have an ESRS Gateway servicing other EMC products, simply use the ESRS Config Tool (refer to the *EMC Secure Remote Support Gateway Operation Guide* for further instructions on Config Tool usage) to add your RecoverPoint devices (or have your gateway administrator do it for you) to the list of ESRS monitored environments. Once the device is added, click the request **update** button to send the new device information to EMC and contact your local EMC Customer Engineer to approve the update.<br>If you do not have an ESRS Gateway at your site, contact your EMC Account Manager to find out more about the benefits of ESRS. |

| Option | Description |
|---|---|
| Location | The city, state and country where the customer is located. |
| Sales order number(s) | 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 EMC Customer Support. |
| VCE | Indicate whether or not this RecoverPoint implementation is operating within a VCE (Vblock) environment. VCE = VMware +Cisco+EMC. |
| Activity type | Enter the type of activity you are performing (upgrade, installation) |
| Resource performing this upgrade/ installation | Enter the role of the person performing this upgrade or installation activity |
| Connect home | The method used to send configuration reports and alerts to EMC. Enabling this feature is recommended as it allows EMC to pro-actively address issues within your RecoverPoint environment, should they arise. “Configure System Reports and Alerts” to establish connectivity for your RecoverPoint environment. |

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

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

# Exporting the post-deployment form to a CSV file

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

### Procedure

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

# Registering RecoverPoint by email or phone

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

### Before you begin

- Register your RecoverPoint system after:
  - installing a RecoverPoint system
  - connecting RPA clusters in a RecoverPoint system
  - upgrading a RecoverPoint system
- The registration process will be incomplete if valid values are not provided for every field in the post-deployment form.

### Procedure

1. Access https://support.emc.com
2. Search for the term *Post-Deployment Form*
3. Download and fill out the *RecoverPoint and RP for VMs Post-Deployment Form*
4. Send the information to the EMC Install Base group:
   - EMC customers and partners: Email the post-deployment form to the Install Base group at rp.registration@emc.com.
   - EMC 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 (regular Eastern Time Zone working hours).

# Registering datastores

Registers a datastore at a vRPA cluster.

### Procedure

1. Navigate to the datastore management tab:
   - In RecoverPoint for VMs 4.3: Select **Administration** › **vRPA Clusters** › **Datastores**. Select the vRPA cluster at which you wish to register datastores, and click **Add...**.
   - In RecoverPoint for VMs 4.3.1: Select **Administration** › **vRPA Clusters** › **Related Objects**. Select the vRPA cluster at which you wish to register datastores, and click **Add...** under the **Datastores** widget.

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

### Results

The specified datastore is registered at the selected vRPA cluster.

# Registering ESX clusters

Registers ESX clusters at a vRPA cluster.

### Before you begin

- This feature is only available in RecoverPoint for VMs 4.3.1 and later versions. For instructions on how to register additional ESX clusters in RecoverPoint for VMs 4.3, see the *EMC RecoverPoint 4.3 Installation and Deployment Guide*.
- ESX clusters that host protected or copy VMs must be registered before you can protect VMs on page 19.

### Procedure

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

### Results

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

# Registering vCenters

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

### Before you begin

- This feature is only available in RecoverPoint for VMs 4.3.1 and later versions.
- Best practice is to configure the vCenter Server to require a certificate, because once RecoverPoint has read the certificate, it does not need further access to the location.
- The default certificate locations are:
  - Windows 2003 Server:
    `C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\SSL\rui.crt`.
  - Windows 2008 Server:
    `C:\Users\All Users\Application Data\VMware\VMware VirtualCenter\SSL\rui.crt`.

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

### Procedure

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

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

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

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

   **Table 1** Add vCenter server

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

4. Click **OK**.

### Results

The specified vCenter Server is registered at the specified vRPA cluster(s) with the specified details.

# CHAPTER 2

# Protecting your data

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


- Support for VMware actions ..... 18
- Protecting virtual machines ..... 19
- Stopping protection ..... 22

# Support for VMware actions

- Host and storage vMotion are both supported for vRPAs and for protected virtual machines.
- Memory and CPU changes to protected virtual machines are allowed and can be replicated to all copy virtual machines using the "Replicating hardware settings" procedure.
- When you add VMDKs to a production virtual machine, RecoverPoint for VMs can automatically provision the corresponding VMDKs at all copies (that is, create a new replication set for each new VMDK). To do so, use the procedure "Adding a VMDK." The journal history is deleted and a volume sweep occurs.
- When production VMDKs are expanded, RecoverPoint for VMs automatically expands all corresponding copy VMs in the group. Refer to "Expanding copy VMDKs automatically."
- When you remove a VMDK from a production virtual machine, the corresponding copy VMDK is not removed. Refer to "Removing a VMDK."
- When you delete a virtual machine from production, any corresponding copy virtual machines in the consistency group (and their VMDKs) are not removed. Delete copy virtual machines manually; or reinstate the production machne by creating an empty virtual machine at the production and Recovering production on page 54.
- When a copy VMDK is removed, RecoverPoint for VMs recreates the VMDK. The copy journal is deleted and a volume sweep is performed.
- VMware snapshots:
  - Taking VMware snapshots is supported only on production virtual machines and not on copy virtual machines. Using the `VM Restore` operation to restore a production virtual machine from a VMware snapshot or clone will cause a full sweep.
  - Before creating a VMware snapshot of a vRPA, it must be detached from the vRPA cluster. See the *EMC RecoverPoint for VMs Installation Guide* for the procedure for detaching vRPAs from a cluster.
  - Promoting VMware snapshots on a copy virtual machine is not supported.
- Replicating a virtual machine with fault tolerance enabled is not supported.
- To clone a copy virtual machine, you must first use the **Test a Copy** wizard to enable image access, and then power the virtual machine off. If the virtual machine is not powered off, a VMware snapshot will be created, which can cause data corruption.
- The resources (memory and CPU) of a running vRPA can be increased but not decreased. The changes will take effect only when the vRPA is rebooted.
- Any `suspend` or `resume` tasks (or hibernation-like functionality) of a running vRPA virtual machine is not supported and may lead to data corruption.
- Collecting vSphere logs from a running vRPA (for instance, using vSphere **File** › **Export** › **Export System Logs** command) may cause inconsistencies in the datastore or other unexpected behavior.
- Upgrading VMware Tools on the vRPAs is unsupported and may lead to undesired results. There is a default version of VMware tools that is installed and shipped with vRPAs, and this special version of VMware tools should never be upgraded.
- Online or offline porting of vRPA is supported without affecting high availability.

# Protecting virtual machines

Describes how to add VMs to new or existing consistency groups, and define their protection settings in RecoverPoint.

### Before you begin

- RecoverPoint for VMs must be licensed and registered on page 11.
- All ESX clusters that host protected VMs or their copies must be registered on page 15.
- All datastores to be used for copy and production journals must be registered on page 14.
- VMs in the same consistency group cannot have the same name.
- When a VM is added to an existing consistency group, if the VM image is larger than the allotted journal size, the system automatically enters one-phase distribution mode.

### Procedure

1. In the vSphere Web Client, right-click on the virtual machine you wish to protect. Select **All RecoverPoint for Virtual Machines Actions › Protect**.
2. In the **Select VM protection method** screen, choose either:
   - **Create a new consistency group for this VM**.
     a. Enter a descriptive name for the new consistency group.
     b. Select the production vRPA cluster.
     c. If you want to add additional virtual machines to protect, mark the **Protect additional VM(s) using this group** checkbox, and continue to step 3 on page 19. If you do not want to add additional virtual machines, click **Next**, and skip to step 4 on page 19.
   - **Add this VM to an existing consistency group**

     **NOTICE**

     Adding a VM to an existing consistency group results in journal loss for that group. Best practice is to protect each virtual machine in its own consistency group.

     a. Select an existing consistency group.
     b. If you want to add additional virtual machines to protect, mark the **Protect additional VM(s) using this group** checkbox, and continue to step 3 on page 19. If you do not want to add additional virtual machines, click **Next**, and skip to step 4 on page 19.
3. In the **Select other VMs to protect** screen, select the additional virtual machines to protect in the consistency group, and click **Add**.
4. In the **Configure production settings** screen, define the production journal and set the advanced production settings:

   If you are adding a virtual machine to an existing consistency group, skip to step 5 on page 20.

   a. Define the minimum journal size for the production copy.
   b. If you don't see datastores in the **Registered Datastores** table, click the **Add** button and select the datastores to register.

c. Choose to automatically select the optimal datastore for the specified journal size or to manually select a datastore, and select a datastore from the table. If you chose to manually select a datastore, select a datastore that has more available size than the specified journal size.

5. Optionally, configure the advanced options for the production copy. The advanced options are displayed per virtual machine.

| Settings | Description |
|---|---|
| **VMDK(s)** | Displays the number of included VMDKs at the relevant production copy, and their total size. |
| **Protection policy** | Default = **`Enabled`**<br>Selecting the **Automatically protect new VMDKs** checkbox ensures that all new VMDK(s) are automatically protected. |
| **Disk provisioning** | Default = **`Same as source`**<br>Defines the way in which the copy VMDKs are to be provisioned; **`Same as source`**, **`Thick provisioning`**, or **`Thin provisioning`**.<br>**Note**<br>When **`Thick provisioning`** or **`Same as source`** are selected, if the production VMDKs are thick (either eager or lazy), the copy VMDKs will be thick provisioned lazy zero. |
| **Hardware changes** | Default = **`Enabled`**<br>Automatically replicates each production VMs hardware resources (such as CPU and memory) to its copy VMs during image access. When enabled, RecoverPoint for VMs will also try to replicate the VM version to all VM copies. If the ESX at a copy doesn't support the production VM version, no hardware resources will be replicated. |

6. In the **Add a copy** screen:

   a. Enter a descriptive name for the copy.

   b. Select the copy vRPA cluster.

   When RecoverPoint for VMs automatically creates a virtual machine at the target, the following limitations apply:

   - All VMDKs will be mapped to a single datastore. You can select the datastore, but you cannot select different datastores for different VMDKs.
   - You cannot use the Protection wizard to select different networks for each network interface card in the non-production copy virtual machine. Virtual machine network connections can be changed afterwards in the vSphere GUI.
   - If changes are made to the memory or CPU resources of a protected (production) virtual machine, those changes will not be replicated to the copy virtual machine. The same changes can be made manually at the copy virtual machine.

   If you chose to have RecoverPoint for VMs create virtual machines automatically, skip to Protecting virtual machines on page 19 .

7. In the **Configure copy settings** screen, configure the copy journal:
   a. Define the minimum journal size for the production copy.
   b. If you don't see datastores in the **Registered Datastores** table, click the **Add** button and select the datastores to register.
   c. Choose to automatically select the optimal datastore for the specified journal size or to manually select a datastore, and select a datastore from the table. If you chose to manually select a datastore, select a datastore that has more available size than the specified journal size.
8. In the **Configure copy settings** screen, define the replication policy for this copy:

| Setting | Description |
|---|---|
| **Load replication policy from template** | Policy templates allow you to re-use pre-defined protection settings. Select link policy |
| **Manually define replication policy** | **`Synchronous`**<br>The host application initiates a write, and then waits for an acknowledgment from the remote vRPA before initiating the next write. This is not the default replication mode, and must be specified by the user. Replication in synchronous mode produces a copy that is always one hundred percent up-to-date with its production source.<br>**`Asynchronous`**<br>The host application initiates a write, and does not wait for an acknowledgment from the vRPA at the copy before initiating the next write. The data of each write is stored in the RPA at the production site, and acknowledged by the vRPA cluster at the production. **RPO**: Default = **`25 Seconds`**.<br>The Recovery Point Objective (or RPO) is the point in time to which you are required to recover data, for a specific application, as defined by your organization. RPO defines the maximum lag allowed on a link, and is set manually in **`Bytes`**, **`KB`**, **`MB`**, **`GB`**, **`TB`**, **`Writes`**, **`Seconds`**, **`Minutes`**, **`Hours`**. |

9. In the **Select copy resources** screen, select how to protect the virtual machine at the target vRPA cluster:

**Note**

Only registered ESX clusters on page 15 will be displayed.

- Automatically create a new virtual machine at the target vRPA cluster:
  a. Expand the tree, and select either an ESX host or an ESX cluster.
  b. Click **Next** and go to step 10 on page 22.
- Select an existing VM to use as the VM copy at the target vRPA cluster:
  a. Expand the tree, and select either an ESX host or an ESX cluster.
  b. Select a virtual machine in the right pain

c. Click **Next** and go to step 10 on page 22.

- If you are configuring multiple virtual machines, configure the target resources for all the production virtual machines:
  a. Select **Automatically create new copy VM(s)**.
  b. Expand the tree in the lower pane, and select the ESX host or the EXS cluster for the target virtual machine.
  c. Click **Next**
  d. In the upper pane, select the production VMs to manually configure the target resources.
  e. In the lower pane, select the datastore to use for the target virtual machine.
  f. Repeat this process for all production virtual machines.
  g. Click **Next** and go to step 10 on page 22.
- Select a production virtual machine to configure the target resource:
  a. In the upper pane, select the production VMs to manually configure the target resources.
  b. In the lower pane, expand the tree, and select an existing virtual machine to use as the copy virtual machine at the target vRPA cluster.
  c. Repeat this process for all production virtual machines.
  d. Click **Next** and go to step 10 on page 22.

10. In the **Select copy storage** screen, elect a datastore where you want to place the data of the copy virtual machine, and click **Next**.
11. In the **Ready to complete**, screen, the consistency group details are displayed:

| Option | Description |
|---|---|
| **Edit...** | Click to view and edits the settings of available copies. |
| **Add a Copy** | Enables you to create multiple copies before starting replication. Returns to the **Add a copy** screen. |
| **Finish** | To complete the wizard, and start replication from the production to all copies. |
| **Cancel** | Cancels the protection wizard. Note that settings will not be saved. |

### Results

The consistency group begins an initialization process and its progress is displayed. The initialization may take some time. After initialization, the consistency group becomes active.

# Stopping protection

To stop replication of a production virtual machine:

### Procedure

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

### Results

Replication stops and the virtual machine is removed from its consistency group. The copy virtual machine 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 will be lost.

# CHAPTER 3

# Managing and monitoring the system

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


- Monitoring the system ..... 26
- Managing the system ..... 27
- Managing group sets ..... 29
- Managing consistency groups ..... 30
- Managing copies ..... 38
- Managing VMs and VMDKs ..... 40

# Monitoring the system

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

### Procedure

1. To access the **Dashboard**, in the **vSphere Web Client** home page, click on the **RecoverPoint for VMs** icon.
2. Note the information in each of the Dashboard widgets:

| Option | Description |
|---|---|
| **Alerts** | Provides a summary of system health categorized by alert level and alert category. |
| **Consistency Group Transfer Status** | Displays a graphical overview of the transfer state of all consistency groups in the RecoverPoint for VMs system. See Consistency group transfer states on page 27 for more information. |
| **Recovery Activities** | Displays all of the accessed copies in the system. To return to the **Recovery** wizard, click on the **Back to Wizard** link. To cancel image access and the recovery activity, click on the **Cancel** link. |

## Reviewing recovery activities

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

### Before you begin

Before you begin, note the following about activity reports:

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

### Procedure

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

   The Recovery Activities will be displayed in the right-hand pane.
4. Select the desired activity.
   - To export the selected activity report, click on the **Export to CSV** button.

- If you wish to remove an activity report from the list, click the **Remove** button.
- If you with to change the timezone, click **Change to GMT**or **Change to local time**.

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

## Monitoring consistency group replication

### Procedure

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

### Results

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

## Consistency group transfer states

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

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

# Managing the system

## Licenses

### Before you begin

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

### Procedure

1. Navigate to the license management tab:
   - In RecoverPoint for VMs 4.3: Select **Administration** › **Licensing**.
   - In RecoverPoint for VMs 4.3.1: Select **Administration** › **vRPA Clusters** › **Licensing**.
2. Click **Add..** to add a new license to the system or click **Remove** to delete an existing license form the system.

## Support settings

### Procedure

1. Select **Administration** › **vRPA Clusters** › **Support** tab.
2. Select a vRPA cluster.
3. Click **Edit Settings...** to edit the cluster support settings.
4. See licensing, support, and registration on page 11 for details on how to configure your support settings.

## RecoverPoint for VMs registration

### Procedure

1. Select **Administration** › **vRPA Clusters** › **Registration** tab.
2. Select a vRPA cluster.
3. Click **Edit Settings...** to manage the cluster registration settings.
4. See licensing, support, and registration on page 11 for details on how to register RecoverPoint for VMs.

## Registering an external host

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

### Before you begin

- SSH must be installed on the external host.
- Only one external host can be configured per vRPA cluster.
- You must define the external host before defining VM start-up scripts in a VM start-up sequence on page 40.

### Procedure

1. Navigate to the external host management tab:
   - In RecoverPoint for VMs 4.3: Select **Administration** › **External Host**. Select the vRPA cluster for which you want to define an external host, and click **Edit...**.
   - In RecoverPoint for VMs 4.3.1: 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, enter the **Name**, **IP**, **User**, and **Password** of the external host for the selected vRPA cluster.
3. Optionally:
   - Click **Check Connectivity** to verify connectivity with the external host.

- Click **Remove** to unregister the external host from the specified vRPA cluster.

## Policy templates

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

**Procedure**

1. Navigate to the policy template management tab:
   - In RecoverPoint for VMs 4.3: Select **Administration** › **vRPA Clusters** › **Templates**.
   - In RecoverPoint for VMs 4.3.1: Select **Administration** › **vRPA Clusters** › **vRPA System** .
2. Manage your **Policy Templates**:
   - Click **Add** to add a new policy template.
   - Select a policy template and click **Edit...** to modify the settings of an existing policy template.
3. In the **Add/Edit Policy Template** dialog box:
   - To configure a link policy template, enter a name for the policy template and define the group or link policy settings on page 31.
   - To configure a copy policy template, enter a name for the policy template and define the copy policy settings on page 38.
4. Optionally:
   - Select a policy template and click **Remove** to delete a policy template.
   - Click **Import**, and select a policy template to import to all RPA clusters in your system.

# Managing group sets

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

## Creating a group set

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

**Procedure**

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

3. In the **Add Group Set** dialog box, enter a name for the group set.
4. Choose the vRPA cluster from which to select consistency groups.
5. Select one or more consistency groups to add to the group set.
6. To enable parallel bookmarking, select **Enable Parallel Bookmarking** and set the bookmarking frequency value. Click **OK**.

You cannot enable parallel bookmarking for a group set if any of the groups in the group set are part of another group set that has parallel bookmarking enabled.

## Editing an existing group set

**Before you begin**

You cannot enable parallel bookmarking for a group set if any of the groups in the group set are part of another group set that has parallel bookmarking enabled.

**Procedure**

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

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

## Enabling a group set

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

**Procedure**

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

## Removing a group set

**Procedure**

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

# Managing consistency groups

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

## Enabling or disabling a consistency group

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

### Procedure

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

 or 

## Editing group or link policies

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

### Procedure

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

| Option | Description |
|---|---|
| Name | The name of the consistency group. |
| Primary RPA | The vRPA that you prefer to replicate the consistency group. When the primary vRPA is not available, the consistency group will switch to another vRPA in the vRPA cluster. When the primary vRPA becomes available, the consistency group will switch back to it. |
| Group priority | Only relevant for remote replication when two or more consistency groups are using the same Primary RPA.<br>Default = Normal<br>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. Click on the **Edit link policy** link to edit the link policy protection settings:

| Option | Description |
|---|---|
| Replication Mode | |
| Dynamic by Latency | Only relevant for synchronous replication mode.<br>Default = Disabled<br>When Enabled, RecoverPoint for VMs alternates between synchronous and asynchronous replication modes, as necessary, according to latency conditions.<br>**Start async replication above**<br>When the specified limit is reached, RecoverPoint for VMs automatically starts replicating asynchronously |

| Option | Description |
|---|---|
| | **Resume sync replication below**<br>When the specified limit is reached, RecoverPoint goes back to replicating synchronously. |
| **Dynamic by Throughput** | Only relevant for synchronous replication mode.<br>Default = Disabled<br>When enabled, RecoverPoint for VMs alternates between synchronous and asynchronous replication modes, as necessary, according to throughput conditions.<br>**Start async replication above**<br>When the specified limit is reached, RecoverPoint for VMs automatically starts replicating asynchronously<br>**Resume sync replication below**<br>When the specified limit is reached, RecoverPoint goes back to replicating synchronously. |
| **RPO** | 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.<br>Default = None<br>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.<br>Enabling and disabling compression causes a short pause in transfer and a short initialization. |
| **Enable Deduplication** | Only relevant for asynchronous remote replication.<br>Default = Disabled<br>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.<br>Enabling and disabling deduplication causes a short pause in transfer and a short initialization. |
| **Snapshot Granularity** | Default = fixed per second.<br>**Fixed per write**<br>Creates a snapshot for every write operation.<br>**Fixed per second**<br>Creates one snapshot per second. Use this for local replication.<br>**Dynamic**<br>The system determines the snapshot granularity according to available resources. Use this for remote replication. |

# Creating bookmarks

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

## Creating crash-consistent bookmarks

### Procedure

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

3. In the **Apply a Bookmark** dialog box, define the bookmark settings:
   a. Enter a bookmark name.
   b. Specify your **Mark as** choice:

| Option | Description |
|---|---|
| **Bookmark name** | Enter a name for the bookmark. |
| **Mark as** | **Crash-Consistent**<br>Labels bookmarks as Crash-Consistent.<br>**Application-Consistent**<br>Labels bookmarks as Application-Consistent. Selecting Application-Consistent does not create an application-consistent snapshot, it only labels the snapshot as application-consistent. |
| **Consolidation policy** | To specify how the consolidation policy will be managed the next time the process runs.<br>• Never consolidate this bookmark<br>• This bookmark snapshot must survive Daily/Weekly/Monthly consolidations:<br>**Daily**<br>Snapshot survives daily consolidations, but is consolidated weekly and monthly.<br>**Weekly**<br>Snapshot survives daily and weekly consolidations, but is consolidated monthly. |

| Option | Description |
| --- | --- |
| | **Monthly**<br>Snapshot survives daily, weekly, and monthly consolidations. |

## Creating application-consistent bookmarks

### Before you begin

The *EMC® RecoverPoint Replicating Microsoft SQL Server Technical Notes* and the *EMC® RecoverPoint Replicating Microsoft Windows File Systems Technical Notes* contain additional detailed information on how to download, install and use RecoverPoint KVSS to create application-consistent bookmarks. KVSS is supported on Windows 2012 R2 and Windows 2008.
KVSS bookmarks are created using the `kvss.exe bookmark` command.

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

When using KVSS to apply bookmarks:

- Parameter values should be surrounded by quotation marks.
- You can use the `vssadmin list writers` command to obtain a list of registered writers on the host 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 enter these values again.
- Multiple writers and groups can be entered simultaneously if they are separated by a space.
- Only the application on which KVSS is run will be application consistent, and only when run on the same VM. Best practice is to name the bookmark accordingly, and to ensure the name of the bookmark contains both the name of the application and the VM.
- Upgrading KVSS requires upgrading your vRPA clusters first. An older version of KVSS will work with a vRPA cluster running a newer version of RecoverPoint for VMs, but a newer version of KVSS will not work with 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_monthy|always]
    [type=[FULL|COPY]]
[ip=<RecoverPoint_cluster_management_ip_address>]
    [user=<RecoverPoint_username>]
    [password=<RecoverPoint_password>]
```

**NOTICE**

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

Where:

**Table 2** KVSS syntax

| Option | Description |
|---|---|
| *writer* | a VSS-aware host application |
| *version* | the KVSS version |
| *group* | RecoverPoint consistency group |
| *bookmark* | name by which you can identify the applied bookmark |
| *consolidation_policy* | consolidation policy to set for this snapshot. Valid values are:<br>*never*; Snapshot is never consolidated.<br>*survive_daily*; Snapshot remains after daily consolidations, but is consolidated in weekly, monthly, and manual consolidations.<br>*survive_weekly*; Snapshot remains after daily and weekly consolidations, but is consolidated in monthly and manual consolidations.<br>*survive_monthly*; Snapshot remains after daily, weekly, and monthly consolidations, but is consolidated in manual consolidations.<br>*always*; Snapshot is consolidated in every consolidation process, whether manual or automatic.<br>Default = *always*<br>**NOTICE**<br>If the consolidation_policy parameter is not specified, the snapshot is consolidated in both automatic and manual consolidation processes. |
| *type* | the shadow copy type: either *FULL* or *COPY*. This setting is optional. Default = *COPY*. The settings are implemented by the writer application. Generally, when type = *full*, backup logs are truncated; when type = *copy*, backup logs are not truncated. |
| *ip* | vRPA cluster management IP |
| *user* | RecoverPoint for VMs username |
| *password* | RecoverPoint for VMs password |

## Procedure

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

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

kvss.exe bookmark
    writers="Microsoft Exchange Writer"
    groups="exchange\comp1" “exchange\comp2”
```

```
    bookmark="exchange hourly snapshot"
    consolidation_policy=”survive_daily”
```

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

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

## Adding a copy

To add a copy to an existing consistency group:

**Procedure**

1. In the RecoverPoint for VMs plug-in, go to the **Protection** tab and click **Virtual Machines**.
2. Click the **Add a copy** icon:

3. In the **Add a Copy** screen:
   a. Enter a descriptive name for the copy.
   b. Select the copy vRPA cluster.
   c. Click **Next**.
4. In the **Configure copy settings** screen, configure the copy journal:
   a. Define the minimum journal size for the copy.
   b. If you don’t see datastores in the registered datastores table, click the **Add** button and select the datastores to register now.
5. In the **Configure copy settings** screen, choose to automatically select the optimal datastore for the specified journal size, or to manually select a datastore, and select a datastore from the table. If you chose to manually select a datastore, select a datastore that has more available size than the specified journal size. Define the replication policy:

| Setting | Description |
|---|---|
| **Load replication policy from template** | Policy templates allow you to re-use pre-defined protection settings.<br>Select link policy. |
| **Manually define replication policy** | **`Synchronous`**<br>The host application initiates a write, and then waits for an acknowledgment from the remote vRPA before initiating the next write. This is not the default replication mode, and must be specified by the user. Replication in synchronous mode |

| Setting | Description |
|---|---|
| | produces a copy that is always one hundred percent up-to-date with its production source.<br><br>**Asynchronous**<br>The host application initiates a write, and does not wait for an acknowledgment from the vRPA at the copy before initiating the next write. The data of each write is stored in the RPA at the production site, and acknowledged by the vRPA cluster at the production. **RPO**: Default = **25 Seconds**.<br><br>The Recovery Point Objective (or RPO) is the point in time to which you are required to recover data, for a specific application, as defined by your organization. RPO defines the maximum lag allowed on a link, and is set manually in Bytes, KB, MB, GB, TB, Writes, Seconds, Minutes, Hours. |

6. In the **Select copy resources** screen, select how to protect the virtual machine at the target vRPA cluster:
   - Automatically create a new virtual machine at the target vRPA cluster:
     a. Expand the tree, and select either an ESX host or an ESX cluster.
     b. Click **Next** and go to step 7 on page 37.
   - Select an existing VM to use as the VM copy at the target vRPA cluster:
     a. Expand the tree, and select either an ESX host or an ESX cluster.
     b. Select a virtual machine in the right pane.
     c. Click **Next** and go to step 7 on page 37.
   - If you are configuring multiple virtual machines, configure the target resources for all the production virtual machines:
     a. Select **Automatically create new copy VM(s)**.
     b. Expand the tree in the lower pane, and select the ESX host or the EXS cluster for the target virtual machine.
     c. Click **Next**
     d. In the upper pane, select the production VMs to manually configure the target resources.
     e. In the lower pane, select the datastore to use for the target virtual machine.
     f. Repeat this process for all production virtual machines.
     g. Click **Next** and go to step 7 on page 37.
   - Select a production virtual machine to configure the target resource:
     a. In the upper pane, select the production VMs to manually configure the target resources.
     b. In the lower pane, expand the tree, and select an existing virtual machine to use as the copy virtual machine at the target vRPA cluster.
     c. Repeat this process for all production virtual machines.
     d. Click **Next** and go to step 7 on page 37.
7. In the **Ready to complete** screen, the consistency group details are displayed. Select the desired action:

| Option | Description |
|---|---|
| **Edit...** | Click to view and edits the settings of available copies. |
| **Add a Copy** | Enables you to create multiple copies before starting replication. Returns to the **Add copy screen**. |
| **Finish** | To complete the wizard, and start replication from the production to all copies. |
| **Cancel** | Cancels the protection wizard. Note that settings will not be saved. |

**Results**

The consistency group begins an initialization process and its progress is displayed. After initialization, the consistency group becomes active.

# Managing copies

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

## Editing copy policies

To edit a copy's protection policy:

**Procedure**

1. In the **vSphere Web Client** home page, click the **RecoverPoint for VMs Management icon** › **Protection** tab › **Consistency Groups**.
2. Expand the list of consistency groups and select the consistency group whose copy policies you wish to edit.
3. Expand the consistency group and select the copy whose policies you wish to edit.
4. Click on the **Edit copy policy** link to edit the copy policy protection settings:

| Option | Description |
|---|---|
| **Journal Compression** | Default = none.<br>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.<br>Enabling journal compression while a consistency group is enabled will result in the loss of all snapshots in the journal. |
| **Maximum Journal Lag** | Default = unlimited<br>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 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. |

| Option | Description |
| --- | --- |
| | Select this option to define a required protection window and specify the length of the required window. You will be notified if the current window is less than the required window. |
| **Enable RecoverPoint Snapshot Consolidation** | Select this to enable automatic snapshot consolidation.<br>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<br>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<br>Snapshots are consolidated every 24 hours. Select Indefinitely to consolidate all subsequent snapshots in 24-hour intervals.<br>• If Indefinitely is not selected, and no weekly consolidations are specified, the remaining snapshots are consolidated monthly.<br>• 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<br>Snapshots are consolidated every seven days.<br>Select Indefinitely to consolidate all subsequent snapshots in seven-day intervals.<br>• If Indefinitely is not selected, the remaining snapshots are consolidated monthly.<br>• If Indefinitely is selected, monthly consolidations are disabled, and the remaining snapshots are consolidated weekly. |

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

# Managing VMs and VMDKs

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

## Orchestration

This section describes the RecoverPoint for VMs features for orchestrating the protection of VMs 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 VM start-up sequence on page 40.

#### Procedure

1. n 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*.

### VM start-up sequence

The VM start-up sequence defines the order of the power-on sequence of the virtual machines in a consistency group 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-up in order of priority, as defined by the user. All virtual machines with the same priority will power-up 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**. When a virtual machine is defined as critical, if it fails to power-up, the start-up sequence will be paused, and no other virtual machines will power-up.

#### Before you begin

- The VMware Tools plug-in must be installed on each production VM that you want to protect.
- One user script and one user prompt can be configured to execute before power-up and to execute after power- up in a strict sequence: **script** › **prompt** › **power-up** › **script** › **prompt**.

The following graphic illustrates the order of sequences:

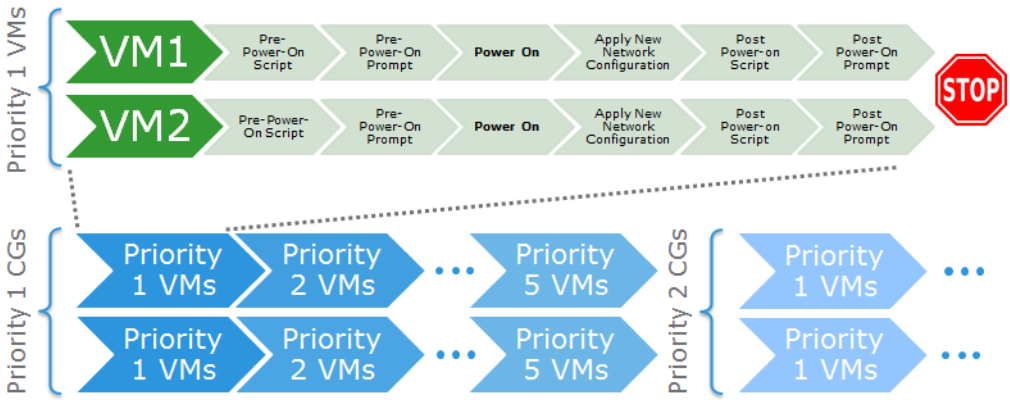


### Procedure

1. In the **RecoverPoint for VMs** plugin, navigate to the **Protection** tab and click **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 will be displayed.
4. Set the order of the power-up sequence by selecting each VM and setting a start-up priority for it.

| Option | Description |
|---|---|
| 1 | The first virtual machine to power-up |
| 3 | Default |
| 5 | The last virtual machine to power-up |

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

### After you finish

See Defining user prompts on page 41 and Defining user scripts on page 42.

## 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 will continue. If the user defines a time-out, the user prompt will automatically dismiss 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 will not continue until the user dismisses the prompt.

### Before you begin

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

### Procedure

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

## Defining user scripts

A user script runs scripts immediately before or after powering on virtual machines. The scripts are executed with `ssh` on the External Host provided by the user. Each script has a mandatory time-out period. The recovery flow will be blocked until the script executes successfully. If the script does not execute within the set time or the script fails or becomes stuck, the system will retry the script a pre-defined number of times (set by the user). The user will receive a prompt indicating 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-up, and one user script after power-up per VM.
- External host on page 28 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. Enter a logical name for the script.
3. Enter the script command, including parameters (separated by a space).
4. Enter time-out period (mandatory).
5. Enter the number of retries.

## Managing VM network configuration (Re-IP)

Virtual network settings can be re-configured for a specific VM, for all VMs at a copy, or for all VMs in a RecoverPoint for VMs system.

The network settings of specific VMs at each copy can be re-configured using the RecoverPoint for VMs GUI. All of the network settings of all VMs at a copy or system can be re-configured by exporting, modifying, and then, re-importing a comma-separated value (CSV) file with the new network settings into the system.

**Note**

VMware Tools must be installed to enable network re-configuration.

The following diagram illustrates how it works.

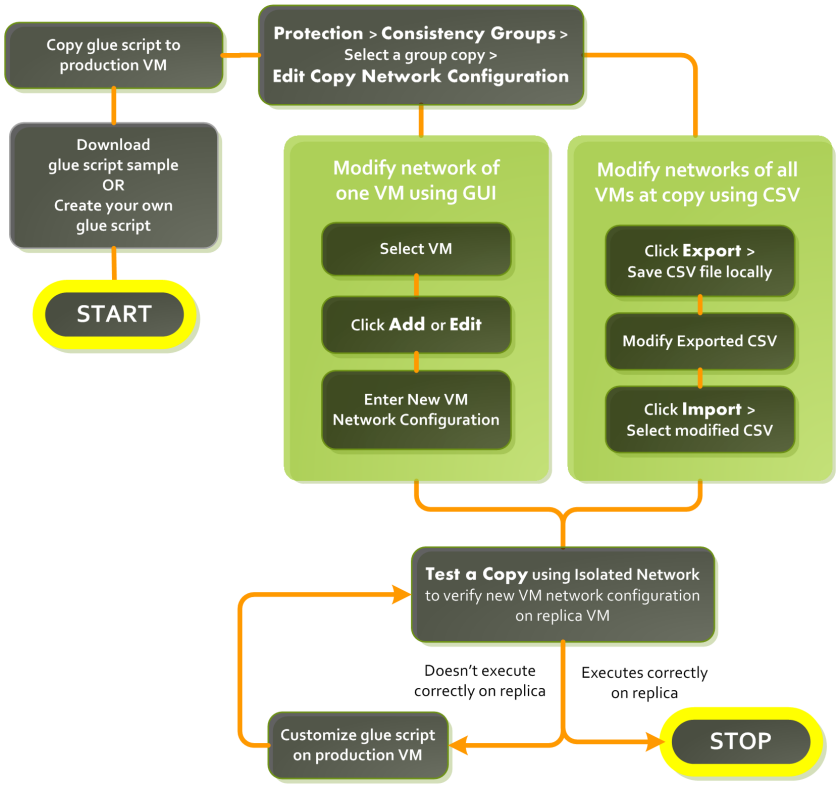


To re-configure the network settings, a *glue script* is required on each protected production VM. RecoverPoint for VMs customers can use their own glue scripts or download the following glue script samples from https://download.emc.com/downloads/DL66792 and customize them to their purposes.

**Table 3** 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 | VMware Tools installed on each protected VM<br>Rename `glue_script_win_this_bat_file.txt`<br>to<br>`glue_script_win.bat` |
| `glue_script_win.py` | Python | Microsoft Windows 2008 and 2012 | Modification of IPv4, IPv6, Subnet Mask, Gateway, DNS servers, DNS Suffix | VMware Tools installed on each protected VM<br>Python 2.7 |

**Table 3** Glue script samples (continued)

| Name | Language | Target OS | Capabilities | Prerequisites |
|---|---|---|---|---|
| `glue_script_rhel.bash` | BASH | LINUX | Modification of IPv4, Subnet Mask, Gateway | VMware Tools installed on each protected VM<br>**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 | VMware Tools installed on each protected VM<br>Python 2.7 |

## Re-configuring VM networks

The following is an example of how a new virtual network configuration can automatically be applied to replica VMs running Windows 2008, 2012, or Linux operating systems.

### Before you begin

1. Ensure that you have VMware Tools installed on each relevant production VM.
2. Create your own glue scripts or download the relevant glue script samples from https://download.emc.com/downloads/DL66792.
3. Copy the relevant glue scripts to the relevant production VMs:
   - In Windows:
     - Place the glue script in a directory that is accessible by all authorized users.
     - To configure the script to run upon 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 your glue script, set the **Start in** directory to the directory into which 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 upon login, add the glue script path to the registry by creating a string 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`**.
   - In Linux, add the relevant glue script to the **`rc.local`** file under `/etc/rc.d`.

### Procedure

1. Select the relevant VM(s):
   - To change the settings of all VMs in a RecoverPoint for VMs system:

  - In RecoverPoint for VMs versions earlier than 4.3.1.1, navigate to **Administration** › **vRPA Clusters** › **Networking** and select the vRPA of the relevant RecoverPoint for VMs system.
  - In RecoverPoint for VMs 4.3.1.1 and later, navigate to **Administration** › **vRPA Clusters** › **vRPA System**, select the vRPA of the relevant RecoverPoint for VMs system, and use the options under the **Network Configuration** widget.
- To modify the network settings of a specific VM or all VMs at a copy, navigate to **Protection** › **Consistency Groups**, expand the relevant consistency group, select the relevant copy, and click the **Edit Copy Network Configuration** icon.

2. Modify the relevant VM network settings.
   - To modify the network settings of one VM, in RecoverPoint for VMs 4.3.1.1 and later:
     a. Select a VM in the table.
     b. click **Add** or **Edit**.
     c. Enter a new network configuration in the **VM Network Configuration** dialog box according to the guidelines in Network configuration settings on page 62.
   - To modify the network settings of all VMs at a copy, or all copies of a RecoverPoint for VMs system:
     a. Click **Save as...** or **Export** to export the current network configuration to a CSV file.
     b. Define a new network configuration for each relevant VM according to the guidelines in Network configuration settings on page 62.
     c. Click **Browse...** or **Import** to import the new network configuration from the modified CSV file.

**Note**

Your new IP configuration will be applied when Testing a copy on page 50.

#### After you finish

Perform the Testing a copy on page 50 procedure to use the new network configuration. Customize the glue scripts on the production VMs until they execute correctly on the replica VMs. Keep in mind:

- By default, RP4VMs does not replicate the MAC address from production to replicas. To enable MAC address replication to remote replicas, contact EMC Customer Service.
- After failover, if you applied the new IP settings to a copy (as opposed to the system), you may need to re-apply the new IP settings to the new production.

## Automation

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

### Replicating virtual machine hardware settings to copies

The hardware resource settings (such as CPU and memory) of each production virtual machine can, during image access, be automatically replicated to its copy virtual machines. By default, **Replicate hardware changes** = **`Enabled`**.

### Procedure

1. To change the replication of hardware settings, navigate to **Protection** › **Virtual Machines screen** screen. Select the virtual machine whose settings you wish to modify.
2. Under the **Hardware Settings** widget, click **Edit...** to change the **Replicate hardware changes.**

## VMDK provisioning

### Before you begin

- The default provisioning method is `Same as source`.
- When **Replicate hardware changes** = `enabled`, RecoverPoint will also try to replicate the source virtual machine version. If the target ESX does not support the source virtual machine version, no hardware resources are replicated.
- When `Thick provisioning` or `Same as source` are selected, if the production VMDKs are thick (either eager or lazy), the copy VMDKs will be thick provisioned lazy zero.

### Procedure

1. Navigate to **Protection** › **Virtual Machines**.
2. Select the virtual machine whose settings you want to modify, and under the **Hardware Settings** widget, click **Edit...**
3. In the **Disk provisioning** drop-down box, select your desired provisioning method; `Same as source`, `Thick provisioning`, or `Thin provisioning`.

## Enabling and disabling automatic protection of new VMDKs

### Procedure

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

## Removing a VMDK

- Removing VMDKs from the production will not delete their copies.
- After removing a VMDK from the production, do one of the following:
  - Add the missing VMDK to the production virtual machine with the same port type, ID, and size as the copy VMDK; which causes a volume sweep.
  - Exclude these VMDKs from replication.

## Excluding a VMDK from replication

If required, you can mark individual VMDKs for exclusion from replication. For example, virtual machines containing shared or non-persistent VMDKs cannot be replicated. You

can, however, change the VMDK type, or use this feature to mark those VMDKs to be excluded from replication and replicate the virtual machines without them.

**Before you begin**

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

**Procedure**

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

## Including a VMDK in replication

**Before you begin**

- VMDKs are included without journal loss.
- Changing the disk type of an included VMDK to an unsupported type (such as **`shared`** or **`non-persistent`**) will cause the same system behavior as removing the VMDK ("Removing a VMDK").

**Procedure**

1. To include a VMDK in replication, navigate to one of the following:
   - In the **Protection** › **Virtual Machines screen** › **Protected VMDKs** widget
   - In the vCenter Server inventory **Summary tab** › **RecoverPoint for VMs widget** › **Protected VMDK(s) section**
2. Click **Edit...** and mark the checkbox next to the VMDKs you wish to include.

## Adding a VMDK

When adding VMDK(s) to the production, RecoverPoint automatically creates the relevant copy VMDK(s). When VMDKs are added to a production virtual machine:

- If the VM image is larger than the allotted journal size, the system automatically enters one-phase distribution mode.
- If you add a VMDK of type **`shared`**, RecoverPoint does not automatically replicate the VMDK. You must manually change each copy VMDK(s) type back to shared after VMDK addition; otherwise, group transfer will be paused and an alert will be displayed in the system **Dashboard**.

## Automatically expanding copy VMDKs

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

- VMDKs can be expanded, but they cannot be shrunk.
- When a production VMDK is expanded, replication of the consistency group is paused by system while the system is busy resizing the corresponding copy VMDK.
- Automatic VMDK expansion will fail if:
  - The datastore does not contain enough free space. In this case, free up space in the copy VM datastore.
  - A snapshot has been taken of the VM 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 your 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 your copy VMDK(s) or contact EMC Customer Support.

- Replication of the consistency group is paused by system if:
  - A copy containing a VMDK marked for automatic expansion is being accessed by the user.
  - A production VMDK is smaller than the size registered in the system settings (because the production VMDK has been removed and re-added with a smaller size). Ensure that the size of all of the VMDKs in the consistency group is the same. If problem persists, contact EMC Customer Support.
  - One or more copy VMDK(s) 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 EMC Customer Support.
- If the size of a copy VMDK is larger than the size of its corresponding production VMDK (for example, because you failed over while automatic expansion was in progress or the copy VMDK was manually expanded), you will have to manually expand the production VMDK to initiate the automatic VMDK expansion process.

## Adding a VM to a consistency group

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

# CHAPTER 4

# Testing and recovery

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


- Testing and recovery ..... 50
- Testing a copy ..... 50
- Failing over ..... 52
- Recovering production ..... 54

# Testing and recovery

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

- Testing a copy on page 50
- Failing over on page 52
- Recovering production on page 54

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

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

- When you select the latest image, the latest available image of each relevant group copy will be accessed.
- When you select an image from the image list, the images of the first group in the group set are displayed. When you select an image from the list, RecoverPoint for VMs constructs a search query containing the parameters of the first group's image and queries the rest of the groups in the group set according to those parameters.
- When you select a specific point in time by entering advanced search options, the same image search query is sent for all groups in the group set.
- During recovery activities, if a point in time is specified for image access at which one or more of the copy VMs either does not exist or does not contain all of the VMDKs that it contains at the current point in time, all missing VMs will be unregistered from RecoverPoint and all missing VMDKs will be detached from the VM until image access is disabled.

# Testing a copy

Guides you through the process of selecting an image and testing it.

**Before you begin**

- When image access is enabled, there is a pause in distribution while the system rolls to the specified image.
- When in image access mode, data is not distributed from the journal to the copy. The length of the delay depends on your storage capabilities and how far the selected image is from the current image. You can close the wizard without interfering with the process.
- The **Test a Copy Wizard** screens contain the following options:
  - **Save&Close**: Keeps access enabled to the image at the specified copy(s) and exits the wizard.

    **Note**

    After clicking **Save&Close**, you can disable image access through the **Recovery Activities** widget in the system **Dashboard**, by selecting **Finish Testing** in the relevant recovery activity bar.
  - **Cancel**: Disables access to the image at the specified copy(s) and exits the wizard.

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

**Note**

After clicking **Finish**, you can disable image access through the **Recovery Activities** widget in the system **Dashboard**, by selecting **Finish Testing** in the relevant recovery activity bar.

### Procedure

1. Select the **Protection tab** and click the **Test Copy** icon:

   The **Test a Copy Wizard** appears.
2. In the **Define a scope** screen, select whether you want to test the consistency group or the group set.
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. Specify the desired select image options.

| Option | Description |
|---|---|
| **The latest image** | The last snapshot that was created at the production, and transferred to the copy journal. |
| **An image from the image list** | Select an image from the list.<br>• The number of snapshots available in the image list is limited. Snapshots that are not in the image list may still be selected by specific Point in Time.<br>• During snapshot dilution, priority is given to bookmarked images. |
| **A specific point in time or bookmark** | Allows you to perform a customized search:<br>**Point in Time**<br>Searches for a bookmark that was created at the specified date and time.<br>**Max Range**<br>Searches for a bookmark that was created between the specified number of minutes/hours before and after the specified Point in Time.<br>**Bookmark**<br>Searches for bookmarks with the specified text in the bookmark name.<br>**Exact**<br>Searches for bookmarks that contains the exact text entered in the Bookmark field.<br>**Image Type**<br>Searches for the specified image type with the specified bookmark name. |

4. In the **Define testing network** screen, define the testing environment. Best practice to avoid IP conflicts between the production virtual machine and the copy virtual machine is to used a dedicated testing network:

| Option | Description |
|---|---|
| **Create an isolated network for each consistency group** | RecoverPoint for VMs will auto-provision an isolated network for virtual machines in this consistency group or group set in order to avoid IP conflicts between the production virtual machines 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. |
| **Skip this step** | Select this option if DHCP is in use or if the production and copy networks are already isolated. |

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.

6. Click **Finish** to exit the wizard and start testing the image at the specified copy(s).

### After you finish

When testing is completed, disable image access.

## Disabling image access

### Procedure

1. Select the **Protection tab** and click **Virtual Machines.**
2. Select the relevant VM.
3. Click the **Disable image** access icon:

# Failing over

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

### Before you begin

- The **Failover Wizard** screens contain the following options:
  - **Save&Close**: Keeps access enabled to the image at the specified copy(s) and exits the wizard.

**Note**

After clicking **Save&Close**, 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.
- **Fail Over**: Starts failing over to the image at the specified copy(s).

- After failover, the production and copy virtual machine change roles, but the names do not change. Therefore, after failover, the new production virtual machine will still be `<YourVMName>.copy` and the new copy virtual machine name will still be `<YourVMName>`.
- The marking information in the production journal is deleted, the copy journal is deleted, and the consistency group undergoes a full sweep synchronization.
- Before you fail back to the production, the Recovery wizard is displayed to enable you to select an image at the production that predates your failover, and to verify the image before permanently selecting it as the image you want to fail back to.

### Procedure

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

   

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

| Option | Description |
|---|---|
| **Current image** | The current image, as displayed in the wizard. |
| **The latest image** | The last snapshot that was created at the production, and transferred to the copy journal. |
| **An image from the image list** | Select an image from the list.<br>• The number of snapshots available in the image list is limited. Snapshots that are not in the image list may still be selected by specific Point in Time<br>• During snapshot dilution, priority is given to bookmarked images. |
| **A specific point in time or bookmark** | Allows you to perform a customized search.<br>**Point in Time**<br>Searches for a bookmark that was created at the specified date and time.<br>**Max Range**<br>Searches for a bookmark that was created between the specified number of minutes/hours before and after the specified `Point in Time`. |

| Option | Description |
|---|---|
| | **Bookmark**<br>Searches for bookmarks with the specified text in the bookmark name.<br>**Exact**<br>Searches for bookmarks that contains the exact text entered in the **Bookmark** field.<br>**Image Type**<br>Searches for the specified image type with the specified bookmark name. |

4. In the **Define testing network** screen, define the testing environment. Best practice to avoid IP conflicts between the production virtual machine and the copy virtual machine. is to use a dedicated testing network

| Option | Description |
|---|---|
| **Create an isolated network for each group** | RecoverPoint for VMs will auto-provision an isolated network for VMs in this consistency group or group set in order to avoid IP conflicts between the production virtual machine and the tested virtual machine. |
| **Create an isolated network for each ESX** | RecoverPoint automatically creates an isolated network for each ESX splitter. |
| **Use my dedicated network** | Manually select a preconfigured network. |
| **Skip this step** | Select this option if DHCP is in use or if the production and copy networks are already isolated. |

5. In the **Ready to complete** screen:
   a. Review the displayed summary information to ensure the failover is configured correctly.
   b. Once image access is enabled, click **Finish** to start failover.

# Recovering production

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

### Before you begin

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

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

**Note**

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

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

### Procedure

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

   

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

| Option | Description |
|---|---|
| The latest image | The last snapshot that was created at the production, and transferred to the copy journal. |
| An image from the image list | Select an image from the list.<br>• The number of snapshots available in the image list is limited. Snapshots that are not in the image list may still be selected by specific `Point in Time`<br>• During snapshot dilution, priority is given to bookmarked images. |
| A specific point in time or bookmark | Allows you to perform a customized search.<br>**Point in Time**<br>Searches for a bookmark that was created at the specified date and time.<br>**Max Range**<br>Searches for a bookmark that was created between the specified number of minutes/hours before and after the specified `Point in Time`.<br>**Bookmark**<br>Searches for bookmarks with the specified text in the bookmark name.<br>**Exact**<br>Searches for bookmarks that contains the exact text entered in the **Bookmark** field.<br>**Image Type**<br>Searches for the specified image type with the specified bookmark name. |

| Option | Description |
|---|---|

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 virtual machine and the copy virtual machine, is to use a dedicated testing network

| Option | Description |
|---|---|
| **Create an isolated network for each group.** | RecoverPoint for VMs will auto-provision an isolated network for virtual machines in this consistency group or group set in order to avoid IP conflicts between the production virtual machines and the tested virtual machines. |
| **Create an isolated network for each ESX.** | RecoverPoint for VMs automatically creates an isolated network for each ESX splitter. |
| **Use my dedicated network** | Manually select a preconfigured network. |
| **Skip this step.** | Select this option if DHCP is in use or if the production and copy networks are already isolated.. |

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

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

### Results

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

# CHAPTER 5

# Troubleshooting

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


- Finding your vRPA cluster management IP ..... 58
- Identifying a RecoverPoint for VMs system ..... 58
- Recovering from a cluster disaster ..... 58
- Detecting bottlenecks ..... 58
- Load balancing ..... 60
- System alerts ..... 60
- Collecting system information ..... 61
- Collecting RecoverPoint for VMs splitter logs ..... 62
- Network configuration settings ..... 62

# Finding your vRPA cluster management IP

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

### Procedure

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

# Identifying a RecoverPoint for VMs system

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

### Procedure

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

# Recovering from a cluster disaster

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

# Detecting bottlenecks

Bottleneck detection returns statistics about RecoverPoint for VMs system performance, by consistency group, vRPA, and vRPA cluster. Bottleneck detection analyzes the system data to detect the existence of any of the predefined problem types called *bottlenecks*. The types of bottlenecks are presented in Table 4 on page 58 and Table 5 on page 59.

### Procedure

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

**Table 4** RPA and cluster bottlenecks

| Bottleneck type detected | System output and comments |
|---|---|
| RPA balance | `RPAs are not balanced.`<br>Includes data on the load handled by each vRPA at the vRPA cluster.<br>**Note**<br>vRPA balance is checked only if the time period defined is greater than 30 minutes. |

**Table 4** RPA and cluster bottlenecks (continued)

| Bottleneck type detected | System output and comments |
|---|---|
| Compression | `Compression level is too high. The RPA resources cannot handle the current level.` |
| SAN target | `RPA may be regulating the application. Consider reducing RPA load.`<br>Includes data on the total amount of incoming data, the number of writes, and the amount of incoming data per write. |
| RPA utilization | `RPA utilization reached ##%.` |

**Table 5** Consistency group and link bottlenecks

| Bottleneck type detected | System output and comments |
|---|---|
| Slow production journal | `Writing to the local journal volume was slow during this period.`<br>Includes data on the delay factor. |
| Journal phase 1 | `Journal is unable to handle the incoming data rate.`<br>Includes the required I/O rates for the journal and the replication volumes at local or remote copies, for both normal and fast-forward distribution modes. |
| Journal phase 2 | `Journal and replication volumes are unable to handle the incoming data rate.`<br>Includes data on the required I/O rates for the journal and the replication volumes at local or remote copies, for both normal and fast-forward distribution modes. |
| Journal regulation | `Remote storage is too slow to handle incoming data rate and regulate the distribution process.`<br>Includes data on the required I/O rates for the journal and the replication volumes at local or remote copies, for both normal and fast-forward distribution modes. |
| Unknown distribution problem | `Target cluster cannot handle the incoming data rate.`<br>Includes data on the required I/O rates for the journal and the replication volumes at local or remote copies, for both normal and fast-forward distribution modes. |
| Slow WAN | `WAN is too slow.`<br>Includes data on total throughput for the vRPA cluster, the identity of the RPAs at which the problem appeared, and the throughput of that RPA (or RPAs).<br>**NOTICE**<br>A slow WAN bottleneck is detected by group, but generates data by vRPA cluster and vRPA. |

**Table 5** Consistency group and link bottlenecks (continued)

| Bottleneck type detected | System output and comments |
|---|---|
| Slow read source | `Reading rate from the source replication volume(s) during synchronization is too slow.` Includes the reading rate. |
| Link utilization | `Link utilization reached ##%.` |

### Results

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

# Load balancing

Load balancing is the process of assigning preferred vRPAs to consistency groups so that the preferred vRPA will perform 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.

You should perform load balancing:

- When a new consistency group is added to the system. Wait one week after the new group is added to accumulate enough traffic history before you perform 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 will display a message indicating so. When you see this message, perform load balancing.
- If the bottleneck detection tool recommends it. When load balancing is required, the `detect_bottlenecks` CLI command returns "`RPAs are not balanced.`" When you see this message, perform load balancing.
- Periodically, to ensure that your 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 enter your RecoverPoint `username` and `password` to log into the CLI.
2. Run the `balance_load` command to balance the load. To view command parameters that can refine your search, run: **`balance_load ?`**

# System alerts

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

# Collecting system information

Collects system information for support purposes.

### Before you begin

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

### Procedure

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

   **Note**

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

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

   **Note**

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

5. Optionally, select **Copy the output file(s) to an FTP server** and define the FTP server settings.
6. Click **Start**.

   **Note**

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

### Results

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

### After you finish

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

1. At the login prompt, enter **admin** as both the **User Name** and the **Password**.
2. Right-click on the relevant file and select **Save link as...** to download the file to your local machine.
3. Open the file using your favorite 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 on **Actions**.
2. Select **All vCenter Actions** › **Export System Logs...**.
3. In the **Export Logs** screen, specify which system logs are to be exported. If required, select the `Gather performance data` option and specify a `duration` and `interval`.
4. Click **Generate Log Bundle**.
5. Click **Download Log Bundle**.
6. Upload the logs to the SFTP/FTP site.

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

# Network configuration settings

The following guidelines should be used to automatically re-configure the virtual networks of VMs in the system while Managing VM network configuration (Re-IP) on page 42.

**Note**

An asterisk (*) denotes a value that is automatically supplied by the system. Two asterisks (**) denote a field that is not relevant in the **VM Network Configuration** dialog box, when re-configuring the network a specific VM.

**Table 6** Network configuration file

| Setting | Description | Guidelines |
|---|---|---|
| **CG ID*** ** | The consistency group ID in the RecoverPoint for VMs system. | Do not modify this field.<br>Not customizable.<br>Can be left blank. |
| **CG Name*** ** | Name of the consistency group in the RecoverPoint for VMs system. | Must be the name associated with the specified CG ID in RecoverPoint for VMs.<br>Can be left blank. |
| **VC ID*** | The vCenter Server ID in VMware. | Do not modify this field.<br>Not customizable. |

**Table 6** Network configuration file (continued)

| Setting | Description | Guidelines |
|---|---|---|
| | | Can be left blank. |
| **VC Name**** | The name of the vCenter Server hosting the virtual machine. | Customizable.<br>Can be left blank. |
| **VM ID*** | The virtual machine ID that vCenter Server uses. | Do not modify this field.<br>Not customizable.<br>Cannot be left blank. |
| **VM Name*** | The name of the virtual machine. | Customizable.<br>Can be left blank. |
| **Adapter ID** | ID of the adapter to customize. | Customizable.<br>Can be left blank.<br>To configure the network settings for a specific NIC, define the Adapter ID as follows:<br>**Windows**<br>Enter the interface index, which can be found by running `route print`.<br>The adapter ID should be set as per the IDX value ascertained from running NetSh Interface IPv4 `Show Interfaces` on the Windows machine and determining the correct adapter.<br>**Linux**<br>The adapter ID should be set as per the Ethernet port value. Enter the sequential number (1-based) of the adapter, and NOT the NIC number.<br>For example, eth0 = 1, eth1 = 2 etc. If you have eth2 and eth3, and want to update the network settings of the second one, set Adapter ID = **2**. |
| **DNS Domain** | The DNS domain for this adapter. | Customizable.<br>Can be left blank.<br>Value should be in the format **`example.company.com`** |
| **Net BIOS** | Whether or not to activate NetBIOS on this adapter. | Net BIOS should be enabled in both Windows and Linux.<br>Valid values are **`DISABLED`**, **`ENABLED`**, **`ENABLED_VIA_DHCP`**.<br>Customizable. |

**Table 6** Network configuration file (continued)

| Setting | Description | Guidelines |
|---|---|---|
| | | Can be left blank. |
| **Primary WINS** | Primary WINS server. | Windows virtual machines only.<br>Customizable.<br>Can be left blank. |
| **Secondary WINS** | Secondary WINS server. | Windows virtual machines only.<br>Customizable.<br>Can be left blank. |
| **IP Address** | IPv4 address for this virtual machine. | Customizable.<br>Can contain either a static IPv4 address or it's DHCP string.<br>Virtual machines can have multiple virtual network adapters. Configure each virtual network adapter with one static IPv4 address.<br>Can be left blank when using IPv6. |
| **Subnet Mask** | IPv4 subnet mask for this virtual machine. | Customizable.<br>Can be left blank when using IPv6. |
| **Gateway(s)** | One or more IPv4 gateways for this virtual machine. | Customizable.<br>Separate multiple values with a semicolon (;).<br>Can be left blank when using IPv6. |
| **IPv6 Address** | IPv6 address for this virtual machine. | Customizable.<br>Can contain either a static IPv6 address or it's DHCP string.<br>Virtual machines can have multiple virtual network adapters. Configure each virtual network adapter with one static IPv6 address.<br>Can be left blank when using IPv4. |
| **IPv6 Subnet Prefix Length** | IPv6 subnet mask for this virtual machine. | Customizable.<br>Can be left blank when using IPv4. |
| **IPv6 Gateway(s)** | One or more IPv6 gateways for this virtual machine. | Customizable.<br>Separate multiple values with a semicolon (;).<br>Can be left blank when using IPv4. |
| **DNS Server(s)** | Address of the DNS server(s). | Customizable.<br>Can be left blank.<br>• Can contain one or more IPv4 DNS servers for each NIC.<br>• For Windows virtual machines, this setting applies only to the configured adapter |

**Table 6** Network configuration file (continued)

| Setting | Description | Guidelines |
|---|---|---|
| | | when a value other than **Adapter ID** 0 is defined.<br>• For Linux virtual machines, this setting applies to all adapters.<br>• Separate multiple values with a semicolon (;). |
| **DNS Suffix(es)** | Suffix(es) for DNS servers.<br>The global settings for all adapters on both Windows and Linux virtual machines. | Customizable.<br>Can be left blank. |