RecoverPoint for VMs

vSphere HTML5 Plugin Administrator's Guide

5.3

DELLEMC

Notes, cautions, and warnings

(i) **NOTE:** A NOTE indicates important information that helps you make better use of your product.

⚠ **CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.**

⚠ **WARNING: A WARNING indicates a potential for property damage, personal injury, or death.**

Contents

Figures

Tables

Preface

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

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

(i) **NOTE:** This document was accurate at publication time. Go to Online Support (https://www.dell.com/support) to ensure that you are using the latest version of this document.

Purpose

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

Audience

This document is intended for use by vSphere administrators who are responsible for managing the RecoverPoint for VMs system.

Related documentation

The following publications provide additional information:

- *RecoverPoint for VMs Release Notes*
- *RecoverPoint for VMs Installation and Deployment Guide*
- *RecoverPoint for VMs Flex Plugin Administrator's Guide*
- *RecoverPoint for VMs Deployment REST API Programming Guide*
- *RecoverPoint for VMs REST API Programmer's Guide*
- *RecoverPoint for VMs Security Configuration Guide*
- *RecoverPoint for VMs Scale and Performance Guide*
- *RecoverPoint for VMs Cloud Solutions Guide*
- *RecoverPoint for VMs CLI Reference Guide*
- **New!** *RecoverPoint for VMs RESTful API* at https://developer.dell.com/apis.

In addition to the core documents, we also provide white papers, technical notes, and demos.

Typographical conventions

This document uses the following style conventions:

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
Italic	Used for full titles of publications referenced in text
`Monospace`	Used for: ● System code ● System output, such as an error message or script ● Pathnames, filenames, prompts, and syntax

	• Commands and options
Monospace italic	Used for variables
`Monospace bold`	Used for user input
[ ]	Square brackets enclose optional values
|	Vertical bar indicates alternate selections - the bar means "or"
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Product documentation

- For release notes and user guides, go to **Online Support** at https://www.dell.com/support.
- For API documentation, see https://developer.dell.com/apis.

Product information

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

Where to get help

Go to **Online Support** at https://www.dell.com/support and click **Contact Support**. To open a service request, you must have a valid support agreement. Contact your sales representative for details about obtaining a valid support agreement or with questions about your account.

Where to find the support matrix

Consult the **Simple Support Matrix** for RecoverPoint for VMs at https://elabnavigator.emc.com.

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 https://contentfeedback.dell.com/s.

Before you begin

Before you start protecting your data in RecoverPoint for VMs, perform the tasks in this section in the provided sequence.

This guide provides the procedures for protecting, recovering and managing VMs:

- **Using the RecoverPoint for VMs Plugin for vSphere Client (HTML5) version 6.7 U1 or later**. When using the RecoverPoint for VMs Flex Plugin, see the *RecoverPoint for VMs Flex Plugin Administrator's Guide*.
- **With on-premises local and/or remote RecoverPoint for VMs copies**. You can also protect your VMs by creating a copy of them in the Amazon Cloud. To deploy the RecoverPoint for VMs Cloud Solution and protect your virtual machines with a cloud copy, see the *RecoverPoint for VMs Cloud Solutions Guide*.

Before you begin:

- See the *RecoverPoint for VMs Product Guide* for a detailed description of the **vSphere HTML5 plugin** and the **vSphere HTML5 plugin server**.
- See the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.
- System installation must be complete. See the *RecoverPoint for VMs Installation and Deployment Guide* for more information.
- The plugin server must be registered with the vCenter Server that you are connected to, or one linked to the vCenter you are connected to. See Managing the plugin server on page 68 for more information.

Topics:

- Create your license files
- Open RecoverPoint for VMs
- Add license
- Register for Customer Support

Create your license files

When a RecoverPoint for VMs sales order is approved, an order confirmation email is automatically sent to the email addresses provided during order entry. The email provides the information you need to begin license activation.

About this task

For more information about software licensing, see these resources:

- Software licensing documentation
- Software Licensing Central Activation video

Steps

1. If you are starting from the Dell Digital Locker, log in or create an account, search for your order, select your product, and click **Activate now**.
   This action takes you to the first step of the license activation wizard in the Software Licensing Center (described in step 3 on page 10).
2. If you do not have the order confirmation email, you can activate your software directly from the Software Licensing Center:
   a. Click **ACTIVATE MY SOFTWARE**.
   b. Provide information such as Licensing Authorization Code (LAC) or Sales Order #.
3. In the **SELECT PRODUCTS** step of the activation wizard, select the product you want to activate, and click **START THE ACTIVATION PROCESS**.
4. In the **SELECT A MACHINE** step:
   a. If you are activating a new instance or rehosting to a new target machine, add a new machine name, and click **SAVE MACHINE & CONTINUE TO NEXT STEP**.

b. If you are increasing capacity on an existing machine, use the **SEARCH MACHINES** widget, select the wanted machine, and click **NEXT: ENTER DETAILS**.
5. In the **ENTER DETAILS** step, enter the quantity of entitlements that you want to activate and the vCenter ID. To learn how to obtain the vCenter ID, click **Machine Details FAQs** and select **RecoverPoint for Virtual Machines**.
6. In the **REVIEW** step, review your selections, and click **ACTIVATE**. This action generates the license keys.
7. In the **COMPLETE** step, view and download the license keys that you generated. To return to the Software Licensing Center home page, click **HOME**.

Results

The entitlements are converted to license files.

Next steps

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

Open RecoverPoint for VMs

Display the RecoverPoint for VMs plugin in your vSphere Client.

Steps

1. Connect to a vCenter Server hosting RecoverPoint for VMs components.
2. Click **LAUNCH VSPHERE CLIENT (HTML5)**. You can also launch the **vSphere Client (HTML5)** directly by entering `https://vCenter-IP or FQDN:/ui/` into your address bar. The HTML5 plugin supports vSphere 6.7 U1 and later versions. In vSphere 7.0 U2 or later, it is recommended to use RecoverPoint for VMs 5.3 SP2 or later.

vmware

Getting Started

LAUNCH VSPHERE CLIENT (HTML5)

Documentation

VMware vSphere Documentation Center

Best practice is to **LAUNCH VSPHERE CLIENT (HTML5)** as your primary client, and only use the **vSphere FLEX plugin** if you are using a vSphere version prior to 6.5 U1, or if you need a feature that is not currently supported through the **vSphere HTML5 plugin** of your RecoverPoint for VMs version. See the *RecoverPoint for VMs Flex Plugin Administrator's Guide* for more information.

In RecoverPoint for VMs 5.3 SP2 and later versions, when running a vSphere version that supports FLEX, only use the **vSphere FLEX plugin** to:
- protect VMs with a copy in the Amazon cloud.
- enable pre-emptive support services.
- display recovery activity reports.
- display performance statistics for consistency groups and vRPA clusters.
- protect VMs with more than one local copy and two remote copies.
- select an existing copy VM when adding a VM to an existing consistency group.

In RecoverPoint for VMs versions prior to 5.3 SP2, when running a vSphere version that supports FLEX, also use the **vSphere FLEX plugin** to monitor system health, components, limits, events, usage, and capacity.

> (i) **NOTE:** On vSphere versions that do not support FLEX, you can only perform these tasks through the *RecoverPoint for VMs RESTful API* at https://developer.dell.com/apis.

In all other cases, click **LAUNCH VSPHERE CLIENT (HTML5)** to display the RecoverPoint for VMs **vSphere HTML5 plugin**.

3. Display the RecoverPoint for VMs plugin in the vSphere Client:
   - Click the **RecoverPoint for VMs** menu item in your main vSphere Client **Menu**.

-OR-

   - Click the **RecoverPoint for VMs** menu item in your vSphere Client **Navigator**.

Results

The RecoverPoint for VMs Dashboard on page 26 is displayed.

Add license

Add a RecoverPoint for VMs license for each vCenter Server that is hosting a protected VM. Adding a license automatically registers the product and enables support.

Prerequisites

Create your license files on page 10. To learn more about the types of licenses that are available, see RecoverPoint for VMs licensing on page 107.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, click **System** > **Licenses** > **Add**.

Add a RecoverPoint for VMs License   ✕

Access your entitlements in the 'Software Licenses' section of support.dell.com or by clicking the link in the LAC email sent to the email address provided during order entry. Activate your entitlements and download your license (*.lic) files from support.dell.com. Add a license for each vCenter Server that is protecting a VM or hosting a vRPA cluster. To learn more about the types of licenses that are available, see the RecoverPoint for VMs Administrator's Guide.

   ⓘ   When you add a socket-based license to a system with VM-based licenses, the system converts VM licenses to socket licenses at a ratio of 15 VMs per socket. When the ratio is not an even conversion, the value is rounded up.

Select a license (*.lic) file   ↲

                                    CANCEL   ADD LICENSE

The **Add a RecoverPoint for VMs License** screen is displayed.

2. Click the browse icon, select a license file (*.lic), and click **Add License**.

Results

Your license is added to the RecoverPoint for VMs system, and its usage statistics and expiration are displayed in the **RecoverPoint for VMs Licenses** table. When both VM-based and socket-based licenses are added, the VM licenses are converted to socket licenses at a rate of 15 VMs to one socket.

To register your system, an email with your license details is automatically sent to emailalerts@dell.com.

Register for Customer Support

If the automatic registration email was not sent during license addition, or a change was made to your system follow this procedure. By ensuring that Customer Support has up-to-date information on the configuration of your system, you enable Dell to provide you with the most effective support possible.

About this task

Refer to the *RecoverPoint for VMs CLI Guide* for more detailed information.

Steps

1. Create an SSH connection to a vRPA cluster management IP address, and use your RecoverPoint for VMs admin username and password to log into the Boxmgmt CLI and, from there, open the Sysmgmt CLI.

2. In the Sysmgmt CLI, run the **`set_registration_params`** command, and provide all of the requested information.

Results

You should receive notification that your registration parameters have been successfully configured. Run the **`get_registration_params`** command if you want to confirm that the information in your system registration is correct.

Protecting VMs

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

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

For additional information about VM management, including how to unprotect VMs, see Managing virtual machines.

(i) **NOTE:** Protecting a virtual machine with fault tolerance enabled is not supported.

Before protecting your VMs, ensure you have:

- Completed the tasks described in Before you begin on page 10.
- Powered on the virtual machines that you want to protect.
- Registered all linked vCenter Servers hosting production VMs and copy VMs as described in Managing the plugin server on page 68.

Topics:

Protect a VM in a new group

Protect a virtual machine in a new consistency group.

Steps

1. Connect to the vSphere Client of your production site.
2. Select **VMs and Templates** view.
3. Right-click a powered-on VM, and select **RecoverPoint for VMs** > **Protect VM...**.
   The **Protect VM** dialog is displayed.

Protect VM ✕

NOTE: In the **Protect VM** dialog, all of the fields are pre-populated with sensible values, so you can safely click **PROTECT** now, and manage the protection policies later, if necessary.

4. (Optional) Click **Edit Settings** to change the default VM protection policy.

- **Disk Provisioning**: Default is `Same as source`. Defines how the copy VMDKs are provisioned; `Same as source`, `Thick provision lazy zeroed`, `Thick provision eager zeroed` or `Thin provision`.
- **Replicate Hardware Changes**: Default is `Enabled`. Automatically replicates the hardware settings of all production virtual machines to their copy VMs whenever an image is accessed on the copy VMs. When enabled, RecoverPoint for VMs replicates the virtual machine version, CPU, memory, resource reservations, and network adapter status and type, and MAC addresses (only to remote copy VMs).

  NOTE: Replication of SR-IOV Passthrough Adapter is not supported. If the ESXi at a copy does not support the production VM version, no hardware changes are replicated.

- **VMDKs**: Displays the number of VMDKs that will be replicated, and their total size. Clear a VMDK checkbox to exclude the VMDK from replication.

5. (Optional) Change the default consistency group protection policies.

- **Consistency Group**: Default is `cg_<vmname>`. Defines the consistency group name.
- **vRPA Cluster**: Defines the vRPA cluster used to replicate and manage the production data to the copies.

- **Journal Datastore**: Defines the datastore that RecoverPoint for VMs will automatically provision a **3GB** VMDK on for the production journal. By default, RecoverPoint automatically registers up to 15 datastores for the production and copy journals and automatically selects the datastore with the most free space.

  (i) **NOTE:** RecoverPoint for VMs will attempt to create the production journal on the selected datastore. If it cannot, the system will attempt to create the production journal on another registered datastore. If you have more than 15 datastores and would like to register an additional datastore that is not in the list, register the other datastores according to Managing journal datastore registration on page 83.

6. (Optional) Change the default copy protection policy.

   Update the copy policies:

   - **vRPA Cluster**: Defines the vRPA cluster used to replicate and manage the production VM data to the storage at this copy.
   - **Sync/Async**: Default is **Asynchronous** with **RPO** (Recovery Point Objective) of **25 Seconds**. The RPO is the point in time to which you are required to recover data, for a specific application, as defined by the organization. RPO defines the maximum lag that is allowed on a link.
   - **vCenter Server** and **ESX Cluster**: Defines the vCenter Server and ESX cluster hosting the copy VMs.
   - **Copy Datastore**: Defines the datastore to use for the copy VM data.

7. (Optional) Click the copy's **Advanced Configuration** icon to update the advanced copy policies:

   - **Journal Datastore**: Defines the datastore that RecoverPoint for VMs will automatically provision a 10GB VMDK on for the copy journal, unless **Manually select copy VM** is selected as the **Copy VM Creation** method. By default, RecoverPoint automatically registers up to 15 datastores for the production and copy journals and automatically selects the datastore with the most free space.

     (i) **NOTE:** RecoverPoint for VMs will attempt to create the copy journal on the selected datastore. If it cannot, the system will attempt to create the copy journal on another registered datastore. If you have more than 15 datastores and would like to register an additional datastore that is not in the list, register the other datastores according to Managing journal datastore registration on page 83.

- **Journal Size**: Default is `10GB`. The larger the copy journal, the more history can be saved.
- **Copy VM Creation**: Default is `Automatically create copy VM`. You can:
  - Click **Select a Resource...** and select the ESXi host to host the copy VM.
  - Select **Manually select copy VM** > click **Select a VM...** , and select a VM from the list.

8. (Optional) For added protection, click **ADD A COPY** to protect the VM with an additional copy:
   - Up to two copies can be created during VM protection. For additional protection, Add a copy to an existing group on page 24 to create more copies.
   - After adding a copy, you can click **Delete Copy** to delete a copy. The last copy cannot be deleted.

     🗑

9. Click **PROTECT**.

Results

The specified virtual machine is protected, and the group production data starts being replicated to the copy VMs according to the specified policies. If an unregistered ESX cluster, or the VMware Resource Pool of an unregistered ESX cluster was selected to host a copy VM, the unregistered ESX cluster is automatically registered with the specified vRPA cluster, a splitter is installed on every ESXi host in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

Next steps

- See Managing VM protection policies on page 84 for additional VM protection policies.
- See Managing group protection policies on page 86 for additional group and copy protection policies.
- For additional protection, Add a copy to an existing group on page 24 to create more copies.

Protect a VM in an existing group

Protect a virtual machine in an existing consistency group.

About this task

The VM that you select to protect will be added to an existing consistency group that is already protecting a VM. For best performance, you should only protect one VM per consistency group.

⚠ **CAUTION: If the image of the VM that you want to protect is larger than the journal size of the copy, the system automatically enters *one-phase distribution mode* upon protection.**

Steps

1. Connect to the vSphere Client of your production site.
2. Select **VMs and Templates** view.
3. Right-click a powered-on VM, and select **RecoverPoint for VMs** > **Protect VM in Existing Group...**

   The **Protect this VM in an Existing Group** dialog is displayed.

Protect this VM in an Existing Group ✕

🔲 other_os	Edit Settings

Consistency Groups   HIDE

'cg_vm2' is selected.

🔍 Search

	Consistency Group ↑	Production vCenter Server	Production vRPA Cluster	Status
⦿	cg_vm2	VM-RP-LAB-VC-288.rp.ilcoe.lab.e…	Site1	🟢 Enabled

Items per page   20 ⌄

1 Consistency group

Copies

vCenter Server	Target ESXi Cluster	Copy Datastore
VM-RP-LAB-VC-288…	Site 1 ⌄	DEV_RPVE34_SIte ⌄

<div align="right">

CANCEL   **PROTECT**

</div>

> ⓘ **NOTE:** All of the fields are pre-populated with sensible values, so you can safely select the consistency group, and click **PROTECT** now. You can manage the VM protection settings later, if necessary, as described in Managing VM protection policies on page 84.

4. Select the consistency group to protect your production VMs.

   When a consistency group is selected, the group copies are displayed and you can change the **Target ESXi Cluster** and **Copy Datastore** of the group copy VMs.

5. (Optional) Click **Edit Settings** to change the default VM protection policies.

   ### Hardware

   Disk Provisioning   Same as source ⌄

   🔵 Replicate Hardware Changes

   VMDKs (4 /4)

☑ Hard disk 1 - SCSI (0:0)	5 GB
☑ Hard disk 2 - SCSI (0:1)	1 GB
☑ Hard disk 3 - SCSI (0:2)	2 GB
☑ Hard disk 4 - SCSI (0:3)	3 GB

   See Managing VM protection policies on page 84 for a detailed description of these VM protection policies, and others that cannot be defined during VM protection.

6. Click **PROTECT**.

Results

The specified virtual machine is protected in the specified consistency group.
- A volume sweep occurs on the newly added VM and a short initialization occurs on all other VMs in the consistency group.
- If an unregistered ESX cluster, or the VMware Resource Pool of an unregistered ESX cluster, was selected to host a copy VM, the unregistered ESX cluster is automatically registered with the specified vRPA cluster, a splitter is installed on every ESXi host in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.
- RecoverPoint for VMs will attempt to create the journals on the selected datastores. If for any reason journal creation fails, the system will attempt to create the journal on another registered datastore.

- If the image of the VMs that you want to protect is larger than the journal size of the copy, the system automatically enters *one-phase distribution mode* upon protection.

Protect multiple VMs in a new group

Protect multiple virtual machines hosted on an ESX cluster in a new consistency group.
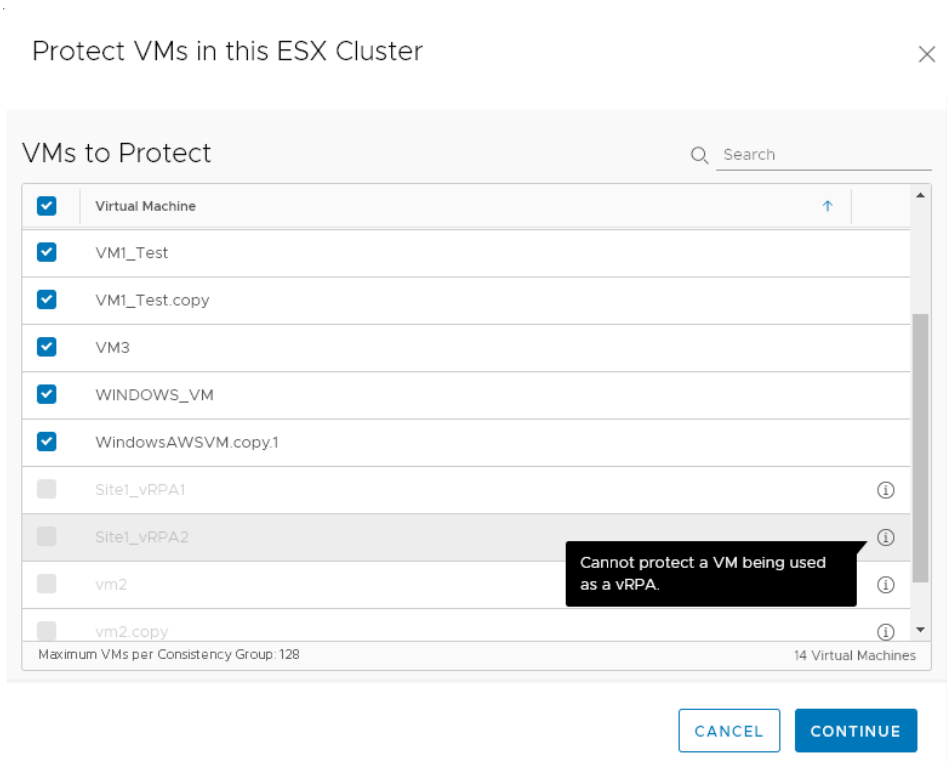
About this task

All of the VMs that you select in the following procedure will be added to a single consistency group. For best performance, you should only protect one VM per consistency group.

Steps

1. Connect to the vSphere Client of your production site.
2. Select **Hosts and Clusters** view.
3. Right-click on an ESX cluster, and select **RecoverPoint for VMs** > **Protect VMs...**.

    The **Protect VMs in this ESX Cluster** dialog is displayed.

    Protect VMs in this ESX Cluster   ✕

    VMs to Protect   🔍 Search

☑	Virtual Machine	↑
☑	VM1_Test	
☑	VM1_Test.copy	
☑	VM3	
☑	WINDOWS_VM	
☑	WindowsAWSVM.copy.1	
☐	Site1_vRPA1	ⓘ
☐	Site1_vRPA2	ⓘ
☐	vm2	ⓘ
☐	vm2.copy	ⓘ

    Cannot protect a VM being used as a vRPA.

    Maximum VMs per Consistency Group: 128   14 Virtual Machines

    CANCEL   CONTINUE

    RecoverPoint for VMs automatically detects if a VM is not able to be protected. Scroll to the bottom of the VM list and click the Info icon next to an excluded VM to display the reason for its exclusion.

4. Select the VMs that you want to protect, and click **CONTINUE**.

    The **Protect VMs** dialog is displayed.

Protect VMs   ✕

🗗 VM3	Edit Settings	🗗 KATE_WINDOWS	Edit Settings	🗗 other_os	Edit Settings
🗗 WINDOWS_VM	Edit Settings	🗗 VM1_Test.copy	Edit Settings	🗗 Deploy_RPCenter_mur...	Edit Settings
🗗 VM1_Test	Edit Settings	🗗 LINUX_VM	Edit Settings	🗗 WindowsAWSVM.copy.1	Edit Settings

Protected by

Consistency Group
cg_9-vms

Production

vRPA Cluster	Journal Datastore
Site1 ⌄	DEV_RPVE34_SIte ⌄

Copies   ＋ ADD A COPY

vRPA Cluster	Sync 🔵 Async	vCenter Server	Target ESXi Cluster	Copy Datastore		
Site1 (Local Copy) ⌄	RPO 25 seconds	VM-RP-LAB-VC-28 ⌄	Site 1 ⌄	DEV_RPVE34_SIte ⌄ ⓘ	⚙	🗑

> Space will soon be insufficient.
> 271.85 GB space is required.

                                                                     CANCEL   **PROTECT**

> ⓘ **NOTE:** In the **Protect VM** dialog, all of the fields are pre-populated with sensible values, so you can safely click **PROTECT** now, and manage the protection settings later, if necessary, as described in Managing VM protection policies on page 84 and Managing group protection policies on page 86.

5. (Optional) Click **Edit Settings** to change the default VM protection policies.

Hardware

Disk Provisioning   Same as source   ⌄

🔵 Replicate Hardware Changes

VMDKs (4 /4)

☑ Hard disk 1 - SCSI (0:0)	5 GB
☑ Hard disk 2 - SCSI (0:1)	1 GB
☑ Hard disk 3 - SCSI (0:2)	2 GB
☑ Hard disk 4 - SCSI (0:3)	3 GB

See Managing VM protection policies on page 84 for a detailed description of these VM protection policies, and others that cannot be defined during VM protection.

6. (Optional) Change the default consistency group and production protection policies.

Protected by

Consistency Group
cg_other_os

Production

vRPA Cluster	Journal Datastore
Site1 ⌄	DEV_RPVE34_SIte ⌄

See Managing group protection policies on page 86 for a detailed description of these group protection policies, and others that cannot be defined during VM protection.

7. (Optional) Change the default copy protection policies.

vRPA Cluster	Sync	Async	vCenter Server	Target ESXi Cluster	Copy Datastore		
Site2 (Remote Cop ⌄	RPO 25 seconds		VM-RP-LAB-VC-28 ⌄	Site 2 ⌄	DEV_RPVE34_Site ⌄	⚙	🗑

See Managing group protection policies on page 86 for a detailed description of these copy protection policies, and others that cannot be defined during VM protection.

8. (Optional) Update the advanced copy policies, see Protect a VM in a new group on page 15 for details.

9. (Optional) For added protection, click **ADD A COPY** to protect the VMs with an additional copy:

   - Up to two copies can be created during VM protection. For additional protection, use Add a copy to an existing group on page 24 to create additional copies.
   - After adding an additional copy, you can click the **Delete Copy** icon to delete a copy. The last copy cannot be deleted.

   🗑

10. Click **PROTECT**.

Results

The specified virtual machines are protected in a new consistency group.

- If an unregistered ESX cluster, or the VMware Resource Pool of an unregistered ESX cluster, was selected to host a copy VM, the unregistered ESX cluster is automatically registered with the specified vRPA cluster, a splitter is installed on every ESXi host in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.
- RecoverPoint for VMs will attempt to create the journals on the selected datastores. If for any reason journal creation fails, the system will attempt to create the journal on another registered datastore.

Protect multiple VMs in an existing group

Protect multiple virtual machines hosted on an ESX cluster in an existing consistency group.

About this task

All of the VMs that you select in the following procedure will be added to a single consistency group. For best performance, you should only protect one VM per consistency group.

> ⚠ **CAUTION:** If the image of the VMs that you want to protect is larger than the journal size of the copy, the system automatically enters *one-phase distribution mode* upon protection.

Steps

1. Connect to the vSphere Client of your production site.
2. Select **Hosts and Clusters** view.
3. Right-click on an ESX cluster, and select **RecoverPoint for VMs** > **Protect VMs in Existing Group…**.

   The **Protect VMs in an Existing Group** dialog is displayed.

Protect VMs in an Existing Group

VMs to Protect

	Virtual Machine	
☑	other_os	
☑	VM1_Test	
☑	VM1_Test.copy	
☑	VM3	
☑	WINDOWS_VM	
☑	WindowsAWSVM.copy.1	
☐	Site1_vRPA1	ⓘ
☐	Site1_vRPA2	ⓘ
☐	vm2	ⓘ

Cannot protect a VM being used as a vRPA.

Maximum VMs per Consistency Group: 128   14 Virtual Machines

CANCEL   CONTINUE

RecoverPoint for VMs automatically detects if a VM is not able to be protected. Scroll to the bottom of the VM list and click the Info icon next to an excluded VM to display the reason for its exclusion.

4. Select the VMs that you want to protect, and click **CONTINUE**.

The **Protect these VMs in an Existing Group** dialog is displayed.

Protect these VMs in an Existing Group

| 🗗 VM3 | Edit Settings | 🗗 other_os | Edit Settings | 🗗 LINUX_VM | Edit Settings |

Consistency Groups HIDE

'cg_vm2' is selected.

	Consistency Group	Production vCenter Server	Production vRPA Cluster	Status
⦿	cg_vm2	VM-RP-LAB-VC-288.rp.ilcoe.lab.e...	Site1	● Enabled
○	cg_newAWS	VM-RP-Lab-H-134.rp.ilcoe.lab.em...	Site1	● Enabled

Items per page 20 ⌄   2 Consistency groups

Copies

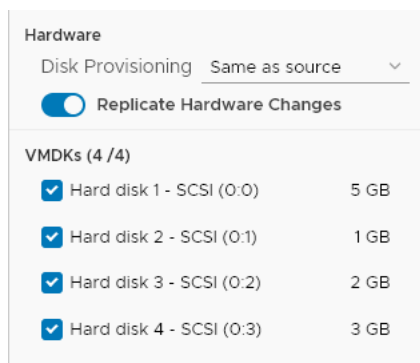| vCenter Server	Target ESXi Cluster	Copy Datastore
VM-RP-LAB-VC-288...	Site 1 ⌄	DEV_RPVE34_Site ⌄

CANCEL   PROTECT

(i) **NOTE:** All of the fields are pre-populated with sensible values, so you can safely select the consistency group, and click **PROTECT** now. You can manage the VM protection settings later, if necessary, as described in Managing VM protection policies on page 84.

5. Select the consistency group to protect your production VMs.

   When a consistency group is selected, the group copies are displayed and you can change the **Target ESXi Cluster** and **Copy Datastore** of the group copy VMs. .

6. (Optional) Click **Edit Settings** to change the default VM protection policies.

Hardware		
Disk Provisioning	Same as source	∨

   (●) Replicate Hardware Changes

   VMDKs (4 /4)
   - ☑ Hard disk 1 - SCSI (0:0)   5 GB
   - ☑ Hard disk 2 - SCSI (0:1)   1 GB
   - ☑ Hard disk 3 - SCSI (0:2)   2 GB
   - ☑ Hard disk 4 - SCSI (0:3)   3 GB

   See Managing VM protection policies on page 84 for a detailed description of these VM protection policies, and others that cannot be defined during VM protection.

7. Click **PROTECT**.

Results

The specified virtual machines are protected in the specified consistency group.
- A volume sweep occurs on the newly added VMs and a short initialization occurs on all other VMs in the consistency group.
- If an unregistered ESX cluster, or the VMware Resource Pool of an unregistered ESX cluster, was selected to host a copy VM, the unregistered ESX cluster is automatically registered with the specified vRPA cluster, a splitter is installed on every ESXi host in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.
- RecoverPoint for VMs will attempt to create the journals on the selected datastores. If for any reason journal creation fails, the system will attempt to create the journal on another registered datastore.
- If the image of the VMs that you want to protect is larger than the journal size of the copy, the system automatically enters *one-phase distribution mode* upon protection.

Add a copy to an existing group

For added protection, add another copy to an existing consistency group.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client, select **Protection** > **Consistency Groups**.
2. Select a group, and click the more group actions button **(...)** > **Add a copy**.
3. In the **Add a Copy** dialog:

Add a Copy to 'MyGroup'   ✕

| vRPA Cluster | Sync 🔵 Async | vCenter Server | Target ESXi Cluster | Copy Datastore | ⚙ |
| WELL (Remote Co| ⌄ | RPO 25 seconds | VM-RP-Lab-VC-154 ⌄ | RPVENV8_Site3 ⌄ | vsanDatastore (1.2 ⌄ | |

[CANCEL] [**ADD COPY**]

ⓘ | **NOTE:** All of the fields are pre-populated with sensible values, so you can safely click **ADD COPY** now, and manage the protection policies later, if necessary.

4. (Optional) Change the default copy protection policy.

   Update the copy policies:

   ● **vRPA Cluster**: Defines the vRPA cluster used to replicate and manage the production VM data to the storage at this copy.
   ● **Sync/Async**: Default is **Asynchronous** with **RPO** (Recovery Point Objective) of **25 Seconds**. The RPO is the point in time to which you are required to recover data, for a specific application, as defined by the organization. RPO defines the maximum lag that is allowed on a link.
   ● **vCenter Server** and **ESX Cluster**: Defines the vCenter Server and ESX cluster hosting the copy VMs.
   ● **Copy Datastore**: Defines the datastore to use for the copy VM data.

5. (Optional) Update the advanced copy policies, see Protect a VM in a new group on page 15 for details.
6. Click **ADD COPY**.

Results

A copy is added to the consistency group, and the group production data starts being replicated to the copy VMs according to the specified policies. If an unregistered ESX cluster, or the VMware Resource Pool of an unregistered ESX cluster, was selected to host a copy VM, the unregistered ESX cluster is automatically registered with the specified vRPA cluster, a splitter is installed on every ESXi host in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

Next steps

See Managing group protection policies on page 86 for additional copy protection policies.

Monitoring protection

After protecting your VMs, use the **Dashboard** to monitor the number and status of protected VMs, consistency groups, group sets, vRPA clusters, system alerts and limits, as well as your RecoverPoint for VMs license usage.

Topics:

- The RecoverPoint for VMs Dashboard
- Monitor system alerts
- Monitor system events
- Monitor system limits
- Monitor system components
- Monitor group and copy protection

The RecoverPoint for VMs Dashboard

The **RecoverPoint for VMs Dashboard** presents a high-level overview of the RecoverPoint for VMs system to help you analyze and monitor your system health.

Figure 1. The RecoverPoint for VMs Dashboard

The **Dashboard** is displayed each time you log into RecoverPoint for VMs. Use the **Dashboard** to monitor the status of your system licenses, limits, alerts, protected VMs, consistency groups, group sets, and recovery activities.

In every **Dashboard widget**, you can:

- **Click the help icon** to display more detailed information on the widget system component or activity.

- **Click a status in a legend** to hide or show the system components or activities in the clicked status from the chart.
- **Click a color of a status in a chart** to go to the relevant system component or activity screen, and display only the system component or activity in the clicked status.

The **Dashboard** in RecoverPoint for VMs 5.3 SP2 and later versions contains the following **widgets** that do not exist in previous RecoverPoint versions.

- **Limits**

  - Displays the `number` and `status` of the limits imposed on system components like consistency groups, splitters, and vRPA clusters.
  - A system component's limit status can be '**Critical**', '**Warning**', or '**OK**'.
  - Click a status in the chart to display all system components with that status in the **Monitoring** > **System Limits** screen.
  - For more information, see how to Monitor system limits on page 34.
- **Protected VMs Size**

  - Displays the `total size` (in GB) of all protected VMs on this vCenter Server
  - For more information, see Protecting VMs on page 15.
- **vRPA Clusters**

- ○ Displays the `number` and `status` of all registered vRPA clusters on all registered and linked vCenter Servers.
- ○ The status of a vRPA cluster can be '**`Error`**', '**`Warning`**', or '**`OK`**'.
- ○ Click a status in the chart to display all vRPA clusters with that status in the **System** > **Administration** > **vRPA Clusters** screen.
- ○ For more information, see Managing vRPA clusters on page 78 and Collecting logs from vRPA clusters on page 103.
- **License Usage**

- ○ Displays the `number of protected sockets` out of the `total number of licensed sockets`.
- ○ A system license's usage status can be '**`Trial`**', '**`OK`**', or '**`Violated`**'.
- ○ Click a status in the chart to display all licenses with that status in the **System** > **Licenses** screen.
- ○ For more information, see Managing your RecoverPoint for VMs licenses on page 68.

The **Dashboard** in RecoverPoint for VMs versions prior to 5.3 SP2 also contains these **widgets**:

ⓘ **NOTE:** RecoverPoint for VMs versions prior to 5.3 SP2 do not contain the **Limits**, **vRPA Clusters**, **License Usage**, and **Protected VMs Size** widgets.

- **Alerts**

- ○ Displays the `number` and `types` of alerts in the system.

- ○ An alert's type can be '**Error**' or '**Warning**'.
- ○ Click an alert type in the chart to display all alerts of that type in the **Monitoring** > **Alerts** screen.
- ○ For more information, see Monitor system alerts on page 30.
- **VM Protection**

- ○ Displays the `number` and `status` of protected VMs on the vCenter Server that you are connected to, or a registered vCenter Server linked to the vCenter that you're connected to.
- ○ A protected VM's status can be '**Active**', '**Error**' or '**Paused**'.
- ○ Click a status in the chart to display all protected VMs with that status in the **Protection** > **Protected VMs** screen.
- ○ For more information, see Protecting VMs on page 15.
- **Group Protection**

- ○ Displays the `number` and `status` of all consistency groups in the system.
- ○ A consistency group's status can be '**Active**', '**Inactive**', '**Warning**' or '**Error**'.
- ○ Click a status in the chart to display all groups with that status in the **Protection** > **Consistency Groups** screen.
- ○ For more information, see Protecting VMs on page 15.
- **Group Recovery Activities**

- ○ Displays the `number` and `status` of all consistency group recovery activities.

- ○ The status of an activity can be '**Error**', '**Action Needed**', or '**In Progress**'.
  - ○ Click a status in the chart to display all groups with that status in the **Recovery Activities** > **Consistency Groups** screen.
  - ○ For more information, see
- ● **Group Set Recovery Activities**

- ○ Displays the `number` and `status` of all group set recovery activities.
  - ○ The status of an activity can be '**Error**', '**Action Needed**', or '**In Progress**'.
  - ○ Click a status in the chart to display all group sets with that status in the **Recovery Activities** > **Group Sets** screen.
  - ○ For more information, see

Monitor system alerts

System alerts is a mechanism that enables vRPA clusters to send events about system components in real time. Monitor system alerts to troubleshoot your RecoverPoint for VMs environment.

Prerequisites

Ensure you have performed the procedure for and have added valid licenses as described in

About this task

To monitor your system alerts, click **Monitoring** > **Alerts**. A system alert's type can be `Warning` or `Error`.

> ⓘ **NOTE:** You can also monitor your system alerts in the

Monitor system events

Monitor system events to troubleshoot your RecoverPoint for VMs environment.

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

By default, the following information is displayed for every event in the **Event Logs**:

- **Level**, which can be: `Info`, `Warning`, or `Error`.
- **Scope**, which can be: `Normal`, `Detailed`, or `Advanced`.
- **Time** and date that the event log was generated.
- **vRPA Cluster** reporting the event.
- **Event ID** that allows the event to be excluded from the events log using the event logs filter.
- **Topic**, which can be: `Splitter`, `Consistency Group`, `Management`, `Cluster`, `RPA`, or `Array`.
- **Summary** of the event.

To monitor your system events, click **Monitoring** > **Event Logs**.

Figure 2. RecoverPoint for VMs event logs

1. Note the total number of events in the event logs.
2. Use the table controls to move to the next page, a previous page, or control the number of events displayed per page.
3. Click the **Event Filter** to control which events are displayed in the **Event Logs** and which are hidden.

ⓘ **NOTE:** Click **APPLY** after changing the event filter settings.

- Click **vRPA Cluster** to select the events for a specific vRPA cluster to display. By default, the events of all vRPA clusters are displayed.
- Click **Time Range** to select the events of a specific time period to display. Select **Unbound** to display all events. Display events based on your local time (the default) or GMT.

- Click **Topics** to hide or display events for specific system components. The event topic can be: `Splitter`, `Consistency Group`, `Management`, `Cluster`, `RPA`, or `Array`.
- Click **Scope** to hide or display logs of specific event scope. The event scope can be: `Normal`, `Detailed`, or `Advanced`
- Click **Level** to hide or display events of a specific level. The event level can be `Info`, `Warning`, or `Error`.
- Click **Event IDs to exclude** to select the events to exclude from display in the events log.

4. Note the date and time that the **Event Logs** were **Last Updated** and use the **Refresh** icon to update the **Event Logs**.
5. Hover over the **Summary** of an event with an elipses (**...**) after it, to display hidden text.

While troubleshooting:

● Use the search bar to display only events that include specific text.

● Click the **Clear Filters** button to easily clear all event filters.

● Click an arrow to expand an event and display the event **Description** and **Details**.

Monitor system limits

Monitor the limits imposed on your RecoverPoint for VMs system and system components to troubleshoot your system.

ⓘ **NOTE:** System limits are imposed by your RecoverPoint for VMs licenses, see Managing your RecoverPoint for VMs licenses on page 68 for more information.

To monitor the state of your system limits, click **Monitoring** > **System Limits**. The **System Limits** screen displays the limits imposed on a system, or on consistency groups, vRPA clusters, or splitters in a system.

A system component's limit status can be '`Critical`', '`Warning`', or '`OK`'. Ensure an OK status is displayed for all of your system limits.

Figure 3. System limits

Figure 4. Consistency group limits

Figure 5. vRPA cluster limits

Figure 6. Splitter limits

> (i) **NOTE:** You can also monitor your system **Limits** in the The RecoverPoint for VMs Dashboard on page 26.

Monitor system components

Monitor RecoverPoint for VMs system components to better understand and troubleshoot your RecoverPoint for VMs environment.

To monitor the state of your system components, click **Monitoring** > **Components**. In the **System Components** screen, ensure an `OK` status is displayed next to each of your RecoverPoint for VMs system components.

> ⓘ **NOTE:** You can also monitor the state of **vRPA clusters** in the

Monitor group and copy protection

Monitor the status of replication for consistency groups and copies, when managing or troubleshooting your system.

Steps

1. Select **Protection** > **Consistency Groups**.
2. Expand a group.
3. Note the **Transfer Status (1)** and **State (2)** of each consistency group and the **Status (3)** of each copy.

Results

- The **Transfer Status (1)** of a consistency group can be:
  - **Active**: Data is being transferred to a copy.
  - **Initializing**: A copy is being initialized: volume sweep, short init, or full sweep.
  - **High Load**: The system enters a temporary high-load state while data is being transferred to a copy, when the journal is full and cannot accept new writes. The system attempts to resolve the high-load state without user action.
  - **Paused by System**: System paused replication so data is not being transferred. If this state occurs for long periods of time, check the system alerts and events in the **Dashboard** for more information.
  - **Error**: An error has occurred.
  - **Permanent High Load**: The system enters a permanent high-load state while data is being transferred to a copy. A permanent high-load can occur after a temporary high-load. The system pauses replication and waits for user action.
  - **Paused**: User paused replication so data is not being transferred to a copy.
  - **Disabled**: User disabled a copy so data is not being transferred.
- The **State (2)** of a consistency group can be:
  - **Enabled**: A group is enabled for replication.
  - **Failed over**: A multi-copy group has completed temporary failover.
  - **Being recovered**: A group is in the process of recovering production.
  - **Partially suspended**: Some of a group's copies have been momentarily suspended while being upgraded.
  - **Suspended**: All of a group's copies have been momentarily suspended while being upgraded.
  - **Disabled**: A group is disabled for replication.
- The **Status (3)** of a copy can be:
  - **OK**: Data can be transferred to the copy.
  - **Initializing**: A copy is being initialized: volume sweep, short init, or full sweep.
  - **High Load**: A copy enters a temporary high-load state while data is being transferred to the copy, when the journal is full and cannot accept new writes. The system attempts to resolve the high-load state for the copy without user action.
  - **Paused by System**: System paused replication so data is not being transferred to a copy. If this state occurs for long periods of time, check the system alerts and events in the **Dashboard** for more information.
  - **Error**: An error has occurred on the copy.
  - **Permanent High Load**: A copy enters a permanent high-load state while data is being transferred to the copy. A permanent high-load can occur after a temporary high-load. The system pauses replication to the copy and waits for user action.
  - **Paused**: User paused replication so data is not being transferred to a copy.
  - **Disabled**: User disabled a copy so data is not being transferred.

VM automation and orchestration

RecoverPoint for VMs provides the following features that automate and orchestrate the recovery of your copy VMs.

Topics:

* Create a bookmark
* Automation
* Orchestration

Create a bookmark

Label a snapshot of a virtual machine, a consistency group, or a group set, for identification during testing and recovery. RecoverPoint for VMs creates crash-consistent snapshots.

About this task

Creating a bookmark on a protected VM creates a bookmark on all copy VMs of the group containing the protected VM. Creating a bookmark on a consistency group that is part of a group set creates a bookmark on all copy VMs of all groups in the group set.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client:
   * To bookmark a snapshot of a consistency group, click **Protection** > **Consistency Groups**.
   * To bookmark a snapshot of a protected virtual machine, click **Protection** > **Protected VMs**.
   * To bookmark a snapshot of a group set, click **Protection** > **Group Sets**.
2. Select the consistency group, protected VM, or group set that you want to bookmark.
3. Click **BOOKMARK**.
4. In the **Bookmark** dialog:

Bookmark VM 'vm2'   ✕

Bookmark

MyBookmark

Snapshot Type

Crash-consistent   ⌄

Snapshot Consolidation Policy

◯ Consolidate snapshot

This snapshot must survive   daily   ⌄ consolidations

CANCEL   CREATE BOOKMARK

* **Bookmark**: Enter a name for the snapshot. The bookmark is the name that is used to identify the snapshot during testing and recovery.

- **Snapshot Type**: Default is `crash-consistent`. Change this value to application-consistent only if you know this snapshot to be application consistent. Selecting `application-consistent` does not create an application-consistent snapshot, it only labels the snapshot as known to be application-consistent.
- **Snapshot Consolidation Policy**
  - **Consolidate snapshot**: Default is `disabled`.
  - **This snapshot must survive `daily`, `weekly`, or `monthly` consolidations**: Default is `daily`.
    - **Daily**: Snapshot survives daily consolidations but is consolidated weekly and monthly.
    - **Weekly**: Snapshot survives daily and weekly consolidations but is consolidated monthly.
    - **Monthly**: Snapshot survives daily, weekly, and monthly consolidations.

5. Click **CREATE BOOKMARK**.

Results

A crash-consistent snapshot is created for the specified VM, group, or group set, with the specified label and consolidation policy.

If the bookmark was created on a:
- Protected VM, the system creates a bookmark on all copy VMs of the group containing the protected VM.
- Consistency group that is part of a group set, the system creates a bookmark on all copy VMs of all groups in the group set.

Next steps

To display bookmarks, go to **Protection** > **Consistency Groups**, expand a group, and select **Snapshots** from the copy commands.

Automation

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

VM automation can be defined when protecting VMs, or later through the VM protection policy.

To configure VM automation after protection, select **Protection** > **Protected VMs**, select a VM, and click **PROTECTION POLICY**.

Protection Policy of VM 'Windows'

Disk Provisioning Same as source

Replicate VM hardware changes

Replicate MAC addresses to the local copy

Automatically protect newly added VMDKs

Protected VMDKs

	Protected VMDK	↑	Path	Size
☑	Hard disk 1		SCSI (0:0)	20 GB
☑	Hard disk 2		SCSI (0:1)	200 MB

CANCEL   UPDATE POLICY

Automatic protection of newly added VMDKs

Define whether or not VMDKs that are added to a protected VM should automatically be protected.

About this task

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

Steps

1. Click **Protection** > **Protected VMs**.
2. Select the virtual machine for which you want to disable the automatic protection of any newly added VMDKs, in the future.
3. Click **PROTECTION POLICY**.
4. Disable or enable **Automatically protect newly added VMDKs**.
5. Click **UPDATE POLICY**.

Results

The protection policy is updated and RecoverPoint for VMs will use the new policy the next time a VMDK is added to this VM.

Provisioning copy VMDKs

Define the way copy VMs are provisioned, per consistency group.

About this task

By default, copy VMDKs are provisioned `Same as source`. Use this procedure to change the default behavior.

Steps

1. Click **Protection** > **Protected VMs**.
2. Select a protected VM.
3. Click **PROTECTION POLICY**.
4. In the **Disk Provisioning** drop-down, select `Same as source`, `Thick provision lazy zeroed`, `Thick provision eager`, or `Thin provision`.
5. Click **UPDATE POLICY**.

Results

Newly added copy VMDKs are provisioned according to the specified settings. Copy VMDKs that were already provisioned will not be re-provisioned, and will keep the provisioning method defined during initial protection.

Excluding a VMDK from replication

Include or exclude protected VMDKs from replication.

About this task

Protected VMs containing non-persistent VMDKs cannot be replicated. They should be excluded from replication or their VMDK type should be changed through vSphere Client.

- Excluded VMDKs are not replicated, but their corresponding copy VMDKs are not deleted. Excluded copy VMDKs are created at the copy, but writes going to excluded VMDKs are not replicated to their copy VMDKs.
- Copy VMDKs are created for any production VMDKs, both included and excluded. For most efficient use of storage resources, ensure that disk provisioning is configured as `Thin provision` (or `Same as source`, if production is thin provisioned) before adding the VMDK or upon VM protection.
- Changing the disk type of a non-persistent VMDK to a persistent VMDK does not automatically include the VMDK in replication, even if **Automatically protect newly added VMDKs** is enabled.

Steps

1. Click **Protection** > **Protected VMs**.
2. Select the VM whose VMDKs you want to exclude from replication.
3. Click **PROTECTION POLICY**.
4. Clear the checkbox next to each **Protected VMDK** that you want to exclude from replication.
5. Click **UPDATE POLICY**.

Results

In the future, the excluded VMDKs are not displayed in the list of snapshots that you can select when Recovering VMs on page 52 from a previous point in time, even when recovering snapshots from a time previous to the VMDK removal.

Automatic replication of VM hardware changes

Enables or disables the automatic replication of hardware changes made to a protected VM.

About this task

By default any hardware changes made to a protected VM through the vSphere Client VM Properties are replicated to all copy VMs. Use this procedure to change the default system behavior.

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

(i) **NOTE:** Replication of the SR-IOV NIC type is not supported. If the ESXi at a copy does not support the production VM version, no hardware resources are replicated.

Steps

1. Click **Protection** > **Protected VMs**.
2. Select a protected VM.
3. Click **PROTECTION POLICY**.
4. Switch **Replicate VM hardware changes** off.
5. Click **UPDATE POLICY**.

Results

Replication of the protected VM hardware changes is enabled or disabled, as defined.

Orchestration

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

VM start-up sequence

Define the order in which VMs in a consistency group will power on during testing and recovery.

Prerequisites

Install VMware Tools on every production VM. When VMware Tools is installed on a production VM, the VM is considered *powered on* only after its operating system loads. When VMware Tools is not installed on a production VM, the VM is considered *powered on* when it is powered on.

About this task

The VM start-up sequence:

- Is initiated when a copy snapshot is accessed during testing or recovery.
- Moves to the next VM in the start-up sequence only when a VM is *powered on*.

- Enables you to define a VM as **Critical** to ensure that no other VMs power on if the critical VM fails to power on first.
- Can contain an **Operation**, with one **user script** and one **user prompt** that will run before VM power-on and after VM power-on. Operations run in a strict sequence: **script** > **prompt** > **power-up** > **script** > **prompt**.

  See Defining user prompts on page 44 and Defining user scripts on page 44 for more information.

The following diagram illustrates the order of sequences:

Steps

1. Click **Protection** > **Consistency Groups**.
2. Select a consistency group, and click **PROTECTION POLICY**.
   The **Group Protection Policy** dialog is displayed.
3. In the **Group Protection Policy** dialog, click **General** > **VM STARTUP SEQUENCE**.

   If the power-on sequence also performs an operation (contains scripts and prompts), a checkmark is displayed in the **Operation** column. See Defining user scripts on page 44 and Defining user prompts on page 44 for more information.

4. Enable **Set same priority for all VMs in group**, or define a **Priority** for each VM in the group.
   By default, all priorities are set to **3**.

Priority	Description
1	The first VMs to power on in the group

Priority	Description
2	The second VMs to power on in the group
3	The third VMs to power on in the group
4	The fourth VMs to power on in the group
5	The last VMs to power on in the group

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

Results

During testing and recovery, the VMs in the group are powered on in the defined order of priority. All of the VMs with the same priority power-on simultaneously.

Defining user prompts

Define a message to display to the administrator that will perform VM copy testing and recovery. User prompts remind the administrator to perform specific tasks before continuing the start-up sequence.

About this task

When defining a VM start-up sequence on page 42, you can add a user prompt before power-on and a user prompt after power-on. Administrator's will have to dismiss the prompt before the start-up sequence continues. If a timeout is defined, the prompt is automatically dismissed when the time-out period passes. If no time-out is defined and a prompt is not dismissed, the start-up sequence does not continue until the prompt is dismissed.

Steps

1. Select **Protection** > **Consistency Groups**, select a consistency group, and click **PROTECTION POLICY**.
2. In the **Protection Policy** dialog, select **General** > **VM STARTUP SEQUENCE**.
3. Expand a VM, and enable **Prompt user** before or after VM power-on.
4. Type a descriptive name for the prompt.
5. Type the prompt message.
6. Optionally, set the time-out period.

Results

During testing and recovery, administrator's will have to dismiss the prompt before the start-up sequence continues, unless a timeout is defined.

Defining user scripts

Run commands immediately before or after VMs are powered-on during testing and recovery.

Prerequisites

- An external host must be configured. One external host can be defined per vRPA cluster. See Managing external host registration on page 84 for more information.
- An SSH server must be installed on each external host.

About this task

When defining a VM start-up sequence on page 42, you can also define the scripts that will be run before or after VMs are powered-on.

- The scripts are run with `ssh` on the external host that you designate.
- Each script has a mandatory time-out period. The recovery flow is blocked until the script runs successfully. A prompt indicates if the script failed.
- You can define one user script before power-on, and one user script after power-on per VM.

- The maximum size of the script name and parameters is 1024 bytes.

Steps

1. Select **Protection** > **Consistency Groups**, select a consistency group, and click **PROTECTION POLICY**.
2. In the **Protection Policy** dialog, select **General** > **VM STARTUP SEQUENCE**.
3. Expand a VM, and enable **Run script** before or after VM power-on.
4. Type a descriptive name for the script.
5. Type the script command, including parameters (separated by a space).
6. Set the time-out period (mandatory).
7. Specify the number of retries. If the script does not run within the set time or the script fails, the system retries the script this number of times.

Results

During testing and recovery, these scripts will run before the start-up sequence continues.

Group start-up sequence

Define the order in which VMs of each consistency group in a group set power-on during testing and recovery.

Prerequisites

Install VMware Tools on every production VM. When VMware Tools is installed on a production VM, the VM is considered *powered on* only after its operating system loads. When VMware Tools is not installed on a production VM, the VM is considered *powered on* when it is powered on.

About this task

The group start-up sequence:

- Is initiated when a copy snapshot is accessed during testing or recovery.
- Moves to the next group of VMs in the start-up sequence only when the last group of VMs are all *powered on*.

Steps

1. Click **Protection** > **Group Sets**.
2. Select a group set.
3. Click **[...]** > **Group priority**, and define a **Priority** for each consistency group in the group set.

Group Priority in 'MyGroupSet'   ✕

Q Search

| cg_Centos-KATE | Priority |
| | 1 - Highest ⌄ |

| cg_rpc_5.3_214 | Priority |
| | 5 - Lowest ⌄ |

CANCEL   **SAVE**

Priority	Description
1	The first group (VMs) to power on in the group set

Priority	Description
2	The second (VMs) to power on in the group set
3	The third group (VMs) to power on in the group set
4	The fourth group (VMs) to power on in the group set
5	The last group (VMs) to power on in the group set

4. Click **SAVE**.

Results

During testing and recovery, group VMs are powered on in the defined order of priority. All group VMs with the same priority power-on simultaneously.

Create a group set

Add a group set to RecoverPoint for VMs.

About this task

A group set is a collection of consistency groups that you can bookmark, enable, disable, pause and resume replication for, and test and recover as a group. You can also create parallel bookmarks on all groups in the group set, at a frequency that you define. Group sets are useful for consistency groups that are dependent on one another or that must work together as a single unit.

> (i) **NOTE:** You cannot enable parallel bookmarking for a group set containing a group that is part of another group set that has parallel bookmarking enabled.

Steps

1. Select **Protection** > **Group Sets**.

   The **Group Sets** screen is displayed.

2. Click **ADD**.
   The **Add Group Set** dialog is displayed.

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

a. Choose the vRPA cluster that the consistency groups you want to add to the group set are replicating from. All groups in a group set must be replicating from the same vRPA cluster.

b. Enter a descriptive name for the group set.

c. (Optional) To create bookmarks for all consistency groups in the group set at the same interval, enable **Parallel Bookmarks** and set the desired bookmark interval frequency.

d. Select the consistency groups to add to the group set.

4. Click **ADD GROUP SET**.

Results

A new group set is created with the specified settings.

Next steps

See for additional group set capabilities.

Re-IP rules

Create Re-IP rules to update the network configuration of one or more copy VMs when *testing a copy*, *failing over*, or *recovering production*.

Prerequisites

- This feature is supported for VMs running Microsoft Windows server versions 8, 10, 2008 R2, 2012, and 2016, Red Hat Linux server versions 6.5 and 7.2, Red Hat Enterprise Linux (RHEL) server version 7.1, and Ubuntu Studio 15.10.
- VMware Tools should be installed on each relevant production VM.
  - For Linux SLES12, automatic network configuration is not supported unless *Open VM Tools* version 9.4.0.25793 and `deployPkg` has been manually installed. See *VMWare KB article 2075048* for detailed information about how to install `deployPkg`.
  - For VMs running *Open VM Tools* versions lower than 9.10, automatic network configuration is not supported unless `deployPkg` has been manually installed. See *VMWare KB article 2075048* for detailed information about how to install `deployPkg`.
- Read the Copy VM network configuration guidelines on page 111.
- Ensure you do not lose your production VM network configuration during failback by also creating re-ip rules for your production VMs.

About this task

You can specify the testing network of one or more VMs at a copy, or of all copy VMs in the system.
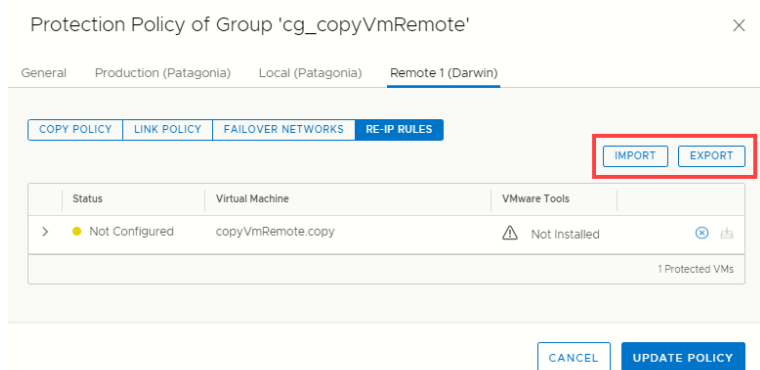
ⓘ **NOTE:** Re-IP configuration using glue scripts is not available in the vSphere HTML5 plugin. You can configure glue scripts using the Flex plugin; for instructions, see the *RecoverPoint for VMs Flex Plugin Administrator's Guide*.

Re-IP a few copy VMs

Update the network configuration of a small number of VMs at a copy.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client, select **Protection** > **Consistency Groups**.
2. Select a consistency group, and click **PROTECTION POLICY**.
   The group **Protection Policy** dialog is displayed.
3. Select the tab of the **Production** or a **Copy** and click the **RE-IP RULES** tab.

4. To retrieve the network configuration of all copy VMs at all vRPA clusters in the system, click **Get values from production**.
5. Update the network values of each copy according to the Copy VM network configuration guidelines on page 111.
6. Click **UPDATE POLICY**.

Results

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

Re-IP many copy VMs

Simultaneously update the network configuration of multiple VMs at a copy.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client, select **Protection** > **Consistency Groups**.
2. Select a consistency group, and click **PROTECTION POLICY**.
   The group **Protection Policy** dialog is displayed.
3. Select the tab of the **Production** or a **Copy** and click the **RE-IP RULES** tab.
4. To retrieve the network configuration of all copy VMs at all vRPA clusters in the system, click **Get values from production**.

Protection Policy of Group 'cg_newVm'

5. To facilitate populating JSON with the production values, click **UPDATE POLICY**.

6. Repeat steps 1-3.

7. To save the current network configuration of all virtual machines at the selected copy to a local JSON file, click **Export**.

8. Open the JSON file, and modify the network configuration of relevant virtual machines according to the Copy VM network configuration guidelines on page 111.

9. To apply the new network configuration, click **Import** and select the modified JSON file .

10. Click **UPDATE POLICY**.

Results

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

Re-IP all copy VMs in the system

Simultaneously update the network configuration of all VMs of all copies in the system.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client, select **System** > **Orchestration** and use the buttons in the **RE-IP Copy VMs** section to update your copy network settings.

2. To save the current network configuration to a local JSON file, click **Export**.

3. Open the JSON file, and modify the network configuration of relevant copy VMs according to the Copy VM network configuration guidelines on page 111.

4. To apply the new network configuration to the system, click **Import** and select the modified JSON file.

Results

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

Failover networks

Automatically associate the VM network adapters (vNICs) of copy VMs with specific port groups upon failover, or during copy testing.

About this task

Failover networks can be configured during or after VM protection. Configured failover networks are made available for selection during testing and failover.

Steps

1. In the RecoverPoint for VMs vSphere Client plugin, click **Protection** > **Consistency Groups**, select a group, and click **PROTECTION POLICY**.
   The group **Protection Policy** dialog is displayed.

2. Select a copy and click **FAILOVER NETWORKS**.

3. Expand a VM to display its network adapters, and for each adapter, select the network to be used after failover.
   Use the search filter to easily identify the required network.

4. Click **UPDATE POLICY**.

Results

The failover networks are configured.

Next steps

Select `Preconfigured failover networks` when defining the **Testing Network** for copy testing, and when defining the **Target Network** before failing over.

Recovering VMs

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

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

Topics:

- Test a copy
- Failover to a copy
- Recover production from a copy

Test a copy

Test a copy of a consistency group or a group set.

Prerequisites

You may want to add journal volumes to a copy journal to ensure that you have ample space for copy testing. For detailed instructions on how to add journal volumes to a copy, see Managing group protection policies on page 86.

About this task

From time to time, and especially before you begin recovery, test your copy snapshots to ensure they are suitable for recovery. Then, Create a bookmark on page 39 so that suitable snapshots are easily identifiable during recovery.

Steps

1. Depending on whether you want to test a copy of a consistency group or a group set, select **Protection** > **Consistency Groups** or **Protection** > **Group Sets**.
2. Select the group or group set whose copy you want to test, and click **TEST A COPY**.
3. In the **Test a Copy** dialog:
   - If you selected a consistency group, select the **Copy to Test**. The vRPA cluster of the copy is displayed, for easy identification.

Test a Copy

×

Copy to Test

Remote Copy 1 (NASA_Site2) ⌄

Snapshot to Test   CHANGE

Latest

Testing Networks   CHANGE

Isolated per group

Power on copy VMs during testing

CANCEL   **START**

- If you selected a group set, select the **vRPA Cluster** containing a copy that you want to test. If there are multiple copies at the specified vRPA cluster, RecoverPoint for VMs automatically selects the copy to test. Consistency groups in the group set that do not have a copy at the specified vRPA cluster will be excluded from the activity.

  (i) **NOTE:** To finish this activity, navigate to the **Recovery Activities** screen, **Consistency Groups** tab.

Test a Copy

×

vRPA Cluster

NASA_Site1 ⌄

Snapshot to Test   CHANGE

Latest

Testing Networks   CHANGE

Isolated per group

Power on copy VMs during testing

CANCEL   **START**

---

4. Select the copy **Snapshot to Test**.

   Default is `latest`. When selecting a copy snapshot:
   - You may want to start with the **Latest** (default) snapshot that is known to be valid.
   - You can search for snapshots by name using the search field.
   - You can select snapshots by **Bookmark** or **Point in Time**.

- Use the **zoom in** icon to display all snapshots between the snapshot whose **zoom in** icon you clicked, and the one before it. You can zoom into a snapshot up to 4 times. After zooming into a snapshot you can zoom out by clicking the originating snapshot timestamp.

- After a snapshot is selected, you can click **REVERT TO LATEST** to revert to the latest snapshot that includes all writes made to it during testing.

5. (Optional) Select the copy **Testing Networks**.

   To avoid IP conflicts between the production and copy VMs, best practice is to use a dedicated testing network. Therefore, by default, RecoverPoint for VMs auto-provisions an isolated network for all VMs in the group or group set .

You can also:

- Create an isolated network for each ESX.
- Use pre-configured Failover networks on page 50.
- Use a dedicated network.

6. Specify whether or not you want RecoverPoint for VMs to **Power on copy VMs during testing**. Default is `enabled`.
7. Click **START** to access the copy snapshot.

Results

The specified snapshot is accessed and a success message is displayed. You can now start testing the copy image.
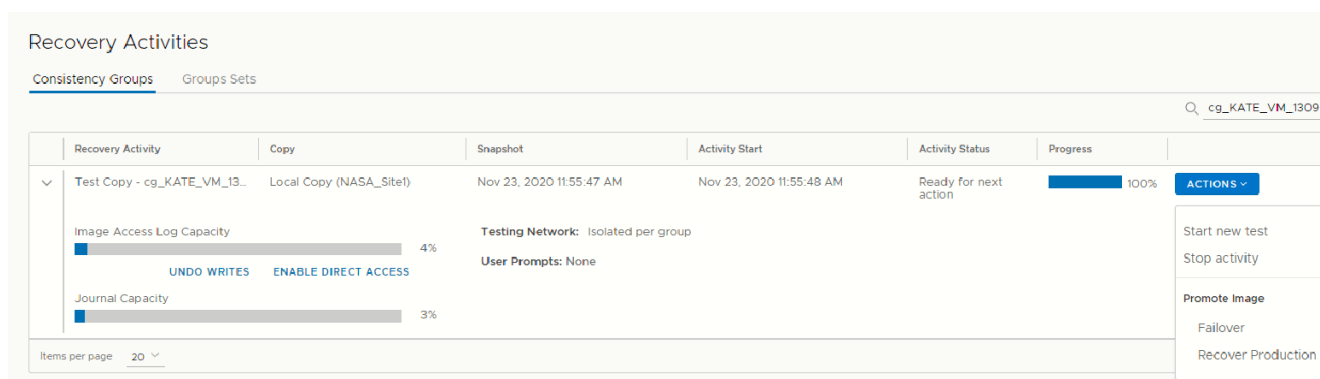

> ⓘ **NOTE:** If you selected to test a copy of a group set, the success message identifies the copy that the system selected, at the vRPA cluster that you selected.

Next steps

Click the **MORE INFO** link in the success message to go to the **Recovery Activities** screen, and display the activity progress and options. In the **Recovery Activities** screen:

- If you tested a copy of a consistency group:

The **Consistency Group Recovery Activities** screen is displayed. The **Activity Status** and **Progress** columns indicate the progress of image access. After access is enabled to the copy snapshot, the **Activity Status** column displays **Ready for next action**, and you can:

> ⓘ **NOTE:** 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. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies on page 86.

- ○ Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- ○ Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

  > ⚠ **CAUTION: When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.**

- ○ Click **ACTIONS** > **Start new test** to select another snapshot to test, or to redefine the testing network.
- ○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- ○ Click **ACTIONS** > **Promote image: Failover** (5.3.1 or later) to Failover to the copy image that you testedin step 8 on page 61 without needing to roll back the writes that were made to the copy snapshot while write access to the copy volumes was enabled.

- ○ Click **ACTIONS** > **Promote image: Recover Production** (5.3.1 or later) to recover production from the copy image that you tested in step 8 on page 66 without needing to roll back the writes that were made to the copy snapshot while write access to the copy volumes was enabled.

- If you tested a copy of a group set:

- ○ Click **OPEN** to display a **Detailed Status** for each group in the group set. The **Detailed Status** screen is displayed.
  (i) **NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

- ○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

(i) **NOTE:** After finding a suitable snapshot, Create a bookmark on page 39 to label the snapshot so it is easily identifiable for recovery.

Failover to a copy

Failover to a copy of a consistency group or a group set, and (optionally) failback to the production. You can failover (and failback) consistency groups or group sets.

About this task

Failover consists of two stages:
- Testing the copy image
- Failover

Failback consists of the same stages.

Before failover, you will have an opportunity to test your copy snapshots to ensure they are suitable for failover.

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

Steps

1. Depending on whether you want to failover a group or a group set, select **Protection** > **Consistency Groups** or **Protection** > **Group Sets**.
2. Select the group or group set that you want to failover, and click **FAILOVER**.
3. In the **Test a Copy for Failover** dialog:
   - If you selected a consistency group, select the copy and vRPA cluster that you want to failover to in the **Failover to Copy** field.

Test a Copy for Failover   ✕

Failover to Copy

Remote Copy 1 (NASA_Site2) ⌄

Failover to Snapshot   CHANGE

Latest

Testing Networks   CHANGE

Isolated per group

CANCEL   **START**

   - If you selected a group set, select the **vRPA Cluster** containing a copy that you want to failover to. If there are multiple copies at the specified vRPA cluster, RecoverPoint for VMs automatically selects the copy. Consistency groups in the group set that do not have a copy at the specified vRPA cluster will be excluded from the activity.

Test a Copy for Failover   ✕

vRPA Cluster

NASA_Site2 ⌄

Failover to Snapshot   CHANGE

Latest

Testing Networks   CHANGE

Isolated per group

CANCEL   **START**

4. Select the snapshot that you want to failover to by clicking **CHANGE** next to **Failover to Snapshot**. Default is the `Latest` snapshot.

When selecting a copy snapshot:

- You may want to start with the **Latest** (default) snapshot that is known to be valid.
- You can search for snapshots by name using the search field.
- You can select snapshots by **Bookmark** or **Point in Time**.

- Use the **zoom in** icon to display all snapshots between the snapshot whose **zoom in** icon you clicked, and the one before it. You can zoom into a snapshot up to 4 times. After zooming into a snapshot you can zoom out by clicking the originating snapshot timestamp.

- After a snapshot is selected, you can click **REVERT TO LATEST** to revert to the latest snapshot that includes all writes made to it during testing.

5. (Optional) Select the copy **Testing Networks**.

To avoid IP conflicts between the production and copy VMs, best practice is to use a dedicated testing network. Therefore, by default, RecoverPoint for VMs auto-provisions an isolated network for all VMs in the group or group set .

You can also:

- Create an isolated network for each ESX.
- Use pre-configured
- Use a dedicated network.

6. Click **START** to access the copy snapshot.

   The specified snapshot is accessed and a success message is displayed.

   Click **MORE INFO** in the success message to go to the **Recovery Activities** screen.

   > **NOTE:** If you selected to test a copy of a group set, the success message identifies the copy that the system selected, at the vRPA cluster that you selected.

7. Test the copy image for failover:

   In the **Recovery Activities** screen, wait for the **Activity Status** to show **Ready for next action** and the **Progress** status bar, indicating the state of image access to reach 100%.

   Then:

   - To select a consistency group for failover, ensure the **Consistency Groups** tab is selected.

   > **NOTE:** By default, replication starts immediately after failover. In RecoverPoint for VMs 5.3.1 and later versions, disable **Start transfer** before failing over to pause replication after failover.

   - Click **ACTIONS** > **Start new test** to select another snapshot to test, or to redefine the testing network.
   - Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
   - (Optional) Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.

- ○ (Optional) Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

  ⚠ **CAUTION: When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.**

- ○ In **Failover Networks**, you can use the default pre-configured failover networks, by keeping **Use or edit pre-configured failover networks** selected. You can also edit a pre-configured network, or choose to **Use current testing networks**.

  ⓘ **NOTE:** 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. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies on page 86.

- ● To select a group set for failover, click the **Group Sets** tab.

ⓘ **NOTE:** By default, replication starts immediately after failover. In RecoverPoint for VMs 5.3.2 and later versions, disable **Start transfer** before failing over to pause replication after failover.

- ○ Click **OPEN** to display the **Detailed Status** of all consistency groups in the group set. After access is enabled to the copy snapshot, the **Status** column of all groups displays **Ready for next action**.

  ⓘ **NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

(i) **NOTE:** After finding a suitable snapshot, you may want to Create a bookmark on page 39 to label the snapshot so it is easily identifiable during failover.

8. Failover to the copy.

   Click **ACTIONS** > **Failover**.

Results

- If the selected consistency group or group set has only one copy, failover starts.
  ○ The role of the **Production** becomes **Remote/Local Copy**.
  ○ The role of the **Remote/Local Copy** becomes **Production**.
  ○ The production VM and copy VM change roles, but their names do not change. Therefore, after failover, new production VMs will still be named *YourVMName*.copy and the new copy VMs are still named *YourVMName*.
  ○ The production journal becomes the copy journal and the copy journal becomes the production journal. You may want to add journal volumes as described in Managing group protection policies on page 86.
  ○ The marking information in the production journal is deleted, the copy journal is deleted, and the consistency group undergoes a full sweep.
    ⚠ **CAUTION: During the full sweep, data is not transferred synchronously.**

- If the consistency group or group set has copies other than the copy to which you are failing over (even if they are disabled or replication to them is paused), a temporary failover begins:
  ○ The role of the **Production** changes to **Temporary Production**.
  ○ The role of the **Remote/Local Copy** changes to **Temporary Remote/Local Copy**.
  ○ The roles of any other (unlinked) copies become **Standalone**.
  ○ Replication pauses for the other copies and the direction of replication between the production and the failed-over copy changes.

Next steps

After temporary failover, if your consistency group or group set had more than one copy (even if they are disabled or replication to them is paused), in the **Recovery Activities** screen:

- **Failback to the original production**. Select the recovery activity, click **ACTIONS** > **Test for failback** and run the above procedure beginning with step 3, substituting "failback" for "failover" throughout.

  After failing back to the production, if you added volumes to the production journal after failover, to reset the production journal to its original size (by default, 3 GB) without triggering a full sweep click **Protection** > **Consistency Groups** > **PROTECTION POLICY**, select the group's **Production** copy, and click **RESET SIZE** in the **Journal Volumes** section.

- **Set the copy as the new production**. Select the recovery activity and click **ACTIONS** > **Set as production**. If there are standalone (unlinked) copies, the **Set this Copy as the New Production** dialog is displayed.

  In the **Set this Copy as the New Production** dialog for consistency groups:

1. Configure each standalone copy for consistency groups (or all standalone copies for group sets).

   Standalone copies are not linked to the production, and you must decide how to handle them before failover. By default, RecoverPoint for VMs does not delete copy VMs but it does disable them. You can **Enable** any required standalone copies and select a replication mode (sync or async), or **Delete** them from the consistency group. Deleting a copy does not delete the VMs from storage.

   ⚠️ **CAUTION: Disabled copy VMs require a full sweep when they are re-enabled.**

2. Click **SET AS PRODUCTION** to permanently failover.
   - The role of the **Production** becomes **Remote/Local Copy**.
   - The role of the **Remote/Local Copy** becomes **Production**.
   - The standalone copies are handled as specified.
   - The production VM and copy VM change roles, but their names do not change. Therefore, after failover, new production VMs will still be named *YourVMName*`.copy` and the new copy VMs are still named *YourVMName*.
   - The production journal becomes the copy journal and the copy journal becomes the production journal. The production journal does not contain the copy history, so it is by default, a much smaller journal. Therefore, after failover, when the production becomes the copy, you may want to add journal volumes to the new copy journal to ensure that you have ample space for copy testing. For detailed instructions on how to add journal volumes to a copy journal, see Managing group protection policies on page 86
   - The marking information in the production journal is deleted, the journal of the copy to which you failed over is deleted, and the consistency group undergoes a full sweep.

     ⚠️ **CAUTION: During the full sweep, data is not transferred synchronously.**

Recover production from a copy

Production recovery corrects file or logical corruption by rolling the production back to a previous point-in-time. You can recover production of consistency groups or group sets.

About this task

Before you begin recovery, you should test your copy snapshots to ensure they are suitable for recovery.

Production recovery consists of two stages:

- Testing the copy image
- Recovering the production from the copy image

Steps

1. Depending on whether you want to recover the production VMs of a group or a group set, select **Protection** > **Consistency Groups** or **Protection** > **Group Sets**.
2. Select the group or group set whose production VMs you want to recover and click **RECOVER PRODUCTION**.
3. In the **Test a Copy for Production Recovery** dialog:
   - If you selected a consistency group, select the copy and vRPA cluster from which you want to recover production in the **Recover Production From Copy** field.

- If you selected a group set, select the **vRPA Cluster** containing a copy from which you want to recover production. If there are multiple copies at the specified vRPA cluster, RecoverPoint for VMs automatically selects the copy. Consistency groups in the group set that do not have a copy at the specified vRPA cluster will be excluded from the activity.

4. Select the snapshot from which you want to recover production by clicking **CHANGE** next to **Recover Production From Snapshot**. Default is the `Latest` snapshot.

   When selecting a copy snapshot:
   - You may want to start with the **Latest** (default) snapshot that is known to be valid.
   - You can search for snapshots by name using the search field.

- You can select snapshots by **Bookmark** or **Point in Time**.

- Use the **zoom in** icon to display all snapshots between the snapshot whose **zoom in** icon you clicked, and the one before it. You can zoom into a snapshot up to 4 times. After zooming into a snapshot you can zoom out by clicking the originating snapshot timestamp.

- After a snapshot is selected, you can click **REVERT TO LATEST** to revert to the latest snapshot that includes all writes made to it during testing.

5. (Optional) Select the copy **Testing Networks**.

   To avoid IP conflicts between the production and copy VMs, best practice is to use a dedicated testing network. Therefore, by default, RecoverPoint for VMs auto-provisions an isolated network for all VMs in the group or group set .

You can also:

- Create an isolated network for each ESX.
- Use pre-configured Failover networks on page 50.
- Use a dedicated network.

6. Click **START** to access the copy snapshot.

The specified snapshot is accessed and a success message is displayed.

> ✓ Accessed snapshot **'Latest'** of **Remote Copy 1** at vRPA cluster **'NASA_Site2'**.   [MORE INFO]   ✕

Click **MORE INFO** in the success message to go to the **Recovery Activities** screen.

ⓘ **NOTE:** If you selected to test a copy of a group set, the success message identifies the copy that the system selected, at the vRPA cluster that you selected.

7. Test the copy image for production recovery:

In the **Recovery Activities** screen, wait for the **Activity Status** to show **Ready for next action** and the **Progress** status bar, indicating the state of image access to reach 100%.

Then:

- To select a consistency group for production recovery, ensure the **Consistency Groups** tab is selected.

- ○ Click **ACTIONS > Start new test** to select another snapshot to test, or to redefine the testing network.
- ○ Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- ○ (Optional) Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- ○ (Optional) Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

  ⚠ **CAUTION: When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.**

  ⓘ **NOTE:** 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. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies on page 86.

- To select a group set for production recovery, click the **Group Sets** tab.

- ○ Click **OPEN** to display the **Detailed Status** of all consistency groups in the group set. After access is enabled to the copy snapshot, the **Status** column of all groups displays **Ready for next action**.

    (i) **NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

- ○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

8. Recover production from the copy.

   Click **ACTIONS** > **Recover production**.

Results

- Data transfer from the production to all copies is paused, and will resume only after production recovery is complete.
- Host access to the recovered production volumes, and the recovering copy volumes is blocked.
- Recovered production volumes are overwritten. Any writes made to the copy during testing are transferred to the production, unless you clicked **UNDO WRITES** in step 7.
- The group undergoes a short initialization process to synchronize the new production data at the copy.

Monitoring recovery activities

Monitor ongoing testing, failover, failback and production recovery activities of consistency groups and group sets, using the **Dashboard**.

Use The RecoverPoint for VMs Dashboard on page 26 to monitor your recovery activities. The **Dashboard** provides an overview of all ongoing recovery activities in the system. Clicking the status of a recovery activity in a dashboard widget automatically displays the relevant system screen, displaying only the system components in the clicked status. To manage recovery activities, see Managing recovery activities on page 91.

Figure 7. Monitoring recovery activities

Managing RecoverPoint for VMs

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

Topics:

Managing your RecoverPoint for VMs licenses

Manage RecoverPoint for VMs licenses.

To monitor your RecoverPoint for VMs license usage, use the describes the licensing process.

To license RecoverPoint for VMs:

1.
2.

Managing the plugin server

RecoverPoint for VMs **plugin server** is supported in vSphere 6.7 U1 and later versions. One plugin server is supported per vCenter Server (or multiple linked vCenter Servers).

Prerequisites

- Ensure you have a plugin server installed and registered with the vCenter Server that you are connected to. See the *RecoverPoint for VMs Installation and Deployment Guide* for more information on installing plugin servers.
- Ensure you have consulted the *RecoverPoint for VMs Product Guide* for a more detailed description of the plugin server, its architecture, installation, and functionality when vCenter Servers are linked.

● See Managing linked vCenter Server registration on page 71.

In the **System** > **Administration** screen:

● The **vCenter Servers** tab displays all vCenter Servers registered with the plugin server of the vCenter Server that you are connected to.
● The **vRPA Clusters** tab displays all vRPA clusters that are hosted on all linked vCenter Servers registered with the plugin server of the vCenter Server that you are connected to.

Changing the plugin server certificate

Use this procedure to change the plugin server certificate before the plugin server has been configured using **Deployment Manager**.

About this task

Use this procedure, for instance, if you want to use a certificate that has been signed by your organization's internal certificate authority.

Steps

1. Connect to the plugin server with root permissions.
2. Create a backup of the existing certificate and key files:

   `/etc/nginx/ssl/rpcenter.cert`

   `/etc/nginx/ssl/rpcenter.key`
3. Disable the firewall on the plugin server.

   Run the command **`/sbin/SuSEfirewall2 off`**
4. Upload the new certificate and key files to `/etc/nginx/ssl`.
5. Rename the new certificate file to **`rpcenter.cert`** and the new key file to **`rpcenter.key`**.
6. Reboot the plugin server VM.
7. In the **RecoverPoint for VMs Deployer**, click **Configure plugin server** home screen.

   Enter the **`plugin server IP address`** in IPv4 format, confirm the new certificate, and click **Configure**.

   For more information, see the "Configure the plugin server" in the *RecoverPoint for VMs Installation and Deployment Guide*.

Results

RecoverPoint for VMs is configured to use the new plugin server certificate.

Next steps

ⓘ **NOTE:**

Check that the certificate is the same across all vRPAs of the same cluster before adding the vRPA to the cluster.

Log into vSphere Client from the relevant vCenter Server and check that the RecoverPoint for VMs HTML5 plugin is displayed.

Changing a registered plugin server certificate

Use this procedure to change the plugin server certificate after the plugin server has already been configured using **Deployment Manager**.

About this task

Use this procedure, for instance, if you want to use a certificate that has been signed by your organization's internal certificate authority.

Steps

1. Connect to the plugin server with root permissions.
2. Create a backup of the existing certificate and key files:

```
/etc/nginx/ssl/rpcenter.cert
/etc/nginx/ssl/rpcenter.key
```

3. Disable the firewall on the plugin server.

   Run the command **/sbin/SuSEfirewall2 off**

4. Upload the new certificate and key files to `/etc/nginx/ssl`.

5. Rename the new certificate file to **rpcenter.cert** and the new key file to **rpcenter.key**.

6. Power off the plugin server VM.

7. Unregister the RecoverPoint for VMs HTML5 plugin from the relevant vCenter Server.

   See "Unregistering the plugin from the Managed Object Browser" in the *RecoverPoint for VMs Installation and Deployment Guide*.

8. Power on the plugin server VM.

9. Navigate to `https://RPCIP/ui`.

10. Click **Authorize** and enter the vCenter Server Credentials.

11. Navigate to **DELETE /vcs/{vc-id}** near the bottom of the Swagger page.

12. Select **Try it Out**, enter the vCenter Server serial number, and select **Execute**.
    A 204 response is returned.

13. In the **RecoverPoint for VMs Deployer**, click **Configure plugin server** home screen.

    Enter the **plugin server IP address** in IPv4 format, confirm the new certificate, and click **Configure**.

    For more information, see the "Configure the plugin server" in the *RecoverPoint for VMs Installation and Deployment Guide*.

Results

RecoverPoint for VMs is configured to use the new plugin server certificate.

Next steps

ⓘ **NOTE:**

Ensure the certificate is the same across all vRPAs of the same cluster before adding the vRPA to the cluster.

Log into vSphere Client from the relevant vCenter Server and check that the RecoverPoint for VMs HTML5 plugin is displayed.

Managing vCenter Server registration with plugin server

Manage the registration of linked vCenter Servers with the plugin server.

About this task

Use the **System** > **Administration** > **vCenter Servers** tab to register linked vCenter Servers with the plugin server, update the registration of a vCenter Server with the plugin server, or unregister a vCenter Server from the plugin server.

ⓘ **NOTE:** The first vCenter Server is registered with the plugin server upon system deployment. See the *RecoverPoint for VMs Installation and Deployment Guide* for more information.

Figure 8. Plugin server vCenter Servers screen

Steps

1. To register a linked vCenter Server with the plugin server, click **Register Linked vCenter**.
2. To update the registration information of a vCenter Server that is already registered with the RecoverPoint for VMs plugin server (for instance, if the vCenter Server password expired, or if you changed the vCenter Server certificate or credentials), and the plugin server session is still active and you can still use the UI:
   a. Click the **Edit** icon.
   b. Re-enter the vCenter Server credentials for the plugin server in the **System** > **Administration** > **vCenter Servers** tab, or for each vRPA cluster under **System** > **Administration** > **vRPA Clusters** tab.
   c. Use the Sysmgmt CLI command `update_vcenter_server_registration` to enter the new vCenter Server credentials for each relevant vRPA cluster. See the *RecoverPoint for VMs CLI Reference Guide* for more details.
3. To unregister a vCenter Server from the plugin server, click the **Delete** icon.
4. Alternatively, if the plugin server session has expired and the UI is disconnected:
   a. Use the Sysmgmt CLI command `update_vcenter_server_registration` to enter the new vCenter Server credentials for each relevant vRPA cluster, see the *RecoverPoint for VMs CLI Reference Guide* for more details.
   b. Use the **Configure plugin server** option in the RecoverPoint for VMs Deployer for a vRPA cluster to update the plugin server with the new certificate, see the Configure plugin server procedure in the *RecoverPoint for VMs Installation and Deployment Guide* for more details.

Next steps

After unregistering a vCenter Server from the plugin server and after updating a plugin server with a new certificate, log out and log back into vSphere, or wait for the session to be re-established.

Managing linked vCenter Server registration

If you have vCenter Servers that are linked together using **vCenter Embedded Linked Mode**, you can see all vRPA clusters, protected VMs, copy VMs, and the plugin servers that reside on multiple vCenter Servers in a single vSphere **Inventory** view, and protect and recover VMs that reside on multiple vCenter Servers.

(i) **NOTE:** The first vCenter Server per site is registered during system deployment through the **RecoverPoint for VMs Deployer**. Additional linked and non-linked vCenter Servers are registered through the **RecoverPoint for VMs vSphere plugin**.

RecoverPoint for VMs:
- Automatically registers the first vCenter Server when you install the first vRPA cluster.
- Supports up to 7 vCenter Servers linked together using **vCenter Embedded Linked Mode**.
- Uses one user authentication method for all linked vCenter Servers.

To register a linked system, after deploying the vRPA clusters:

1. Ensure you have used the **RecoverPoint for VMs Deployer** > **Install Plugin Server** button to install and register a plugin server with every vCenter Server hosting a protected VM, a copy VM, a plugin server or a vRPA cluster. The plugin server installs the **vSphere HTML5 plugin** on every registered vCenter Server. When a vCenter Server is not registered with any plugin server, the **vSphere HTML5 plugin** interface is not available through the vSphere Client, and you cannot operate

your RecoverPoint for VMs system. See the *RecoverPoint for VMs Installation and Deployment Guide* for more information on installing the plugin server.

2. Register all linked vCenter Servers and the plugin server through the **RecoverPoint for VMs vSphere plugin**:
   ● Ensure you have registered all linked vCenter Servers with each vRPA cluster hosted on the linked vCenter Servers.
   ● RecoverPoint for VMs automatically registers remote vRPA clusters when a local vRPA cluster is registered, but does not automatically register the vCenter Server on which the remote vRPA clusters are hosted.
   ● When a vCenter Server is not registered with at least one vRPA cluster, the **RecoverPoint for VMs** > **Protect VMs...** menu is disabled, and you cannot protect VMs from the vSphere **Inventory**.
   ● Click **System** > **Administration** > **vCenter Servers** and ensure you have registered all linked vCenter Servers with a plugin server.
   ● When a vCenter Server is not registered with any plugin server, the **RecoverPoint for VMs** > **Protect VMs...** menu is disabled, and you cannot protect VMs from the vSphere **Inventory**.
3. Connect to a vCenter Server directly to:
   ● Manage the plugin server of the linked vCenter Server.
   ● Add a RecoverPoint for VMs license to the linked vCenter Server.

   Licence usage information is combined for all linked vCenter Servers.

To display all linked vCenter Servers in a linked RecoverPoint for VMs system, click the plugin server **INSTANCE** at the top of the RecoverPoint for VMs vSphere plugin.

Local linked vCenter Servers

If you have local vCenter Servers that are linked together using **vCenter Embedded Linked Mode** and remote non-linked vCenter Servers, you can connect to the local vCenter Servers to display all of the production VMs and vRPA clusters at the production site in one vSphere Client inventory. When connecting to the remote vCenter Server, only the remote vRPA clusters and copy VMs will be displayed in the vSphere Client inventory.

Figure 9. Local linked vCenter example

As illustrated in the Local linked vCenter example on page 73, after installing your vRPA clusters:

1. Launch the **RecoverPoint for VMs Deployer** by clicking **System** > **Administration** > **vRPA Cluster** and clicking the **Display more vRPA cluster options** icon for a vRPA cluster that is registered to a linked vCenter:

2. Connect to **Local vCenter Server1** or **Local vCenter Server2** and register the vCenter with the **Local Plugin Server**.
3. Connect to the **Remote vCenter Server** and register it with the **Remote Plugin Server**.
4. In the **RecoverPoint for VMs vSphere plugin**:
   - When Managing the plugin server on page 68, ensure the **System** > **Administration** > **Registered vCenter Servers** table displays linked **Local vCenter Server 2** is registered with the **Local Plugin Server**.
   - When Managing vCenter Server registration with plugin server on page 70 of **Local vCenter Server1** and **Local vCenter Server2**:
     - Register **Local vCenter Server 1** with **Local vRPA Cluster 1**.
     - Register **Local vCenter Server 2** with **Local vRPA Cluster 2**.
   - Click **INSTANCE** to display both the local and remote plugin servers and all linked and registered vCenter Servers at the production site.

- When connected to the **Remote vCenter Server**:
  - When Managing vCenter Server registration with plugin server on page 70, register **Remote vCenter Server** with **Remote vRPA Cluster 1** and **Remote vRPA Cluster 2**.
  - When Managing the plugin server on page 68, register **Remote vCenter Server** with the **Remote Plugin Server**.
  - Click **INSTANCE** to display the remote plugin server and the remote vCenter Server.

(i) **NOTE:**

In VMware Cloud Foundation (VCF) multi-tenancy environments (such as VxRail):

- It is best practice to register vRPA clusters with different vCenter Servers (see *Workload Domain* or *WLD* in VCF documentation).
- All VMs of all WLDs are displayed in the RecoverPoint for VMs plugin, no matter which user is logged in. For example, a user from a tenant on WLD 1 will see and can perform recovery on consistency groups from WLD 2. The RecoverPoint for VMs vSphere plugin does not currently support role-based authentication.
- You can see, manage, protect and recover only the VMs of vRPA clusters that have been directly registered with a vCenter Server.

Remote linked vCenter Servers

If you have vCenter Servers that are linked together using **vCenter Embedded Linked Mode**, you can display, manage, protect and recover VMs hosted on multiple linked vCenter Servers, across remote sites.
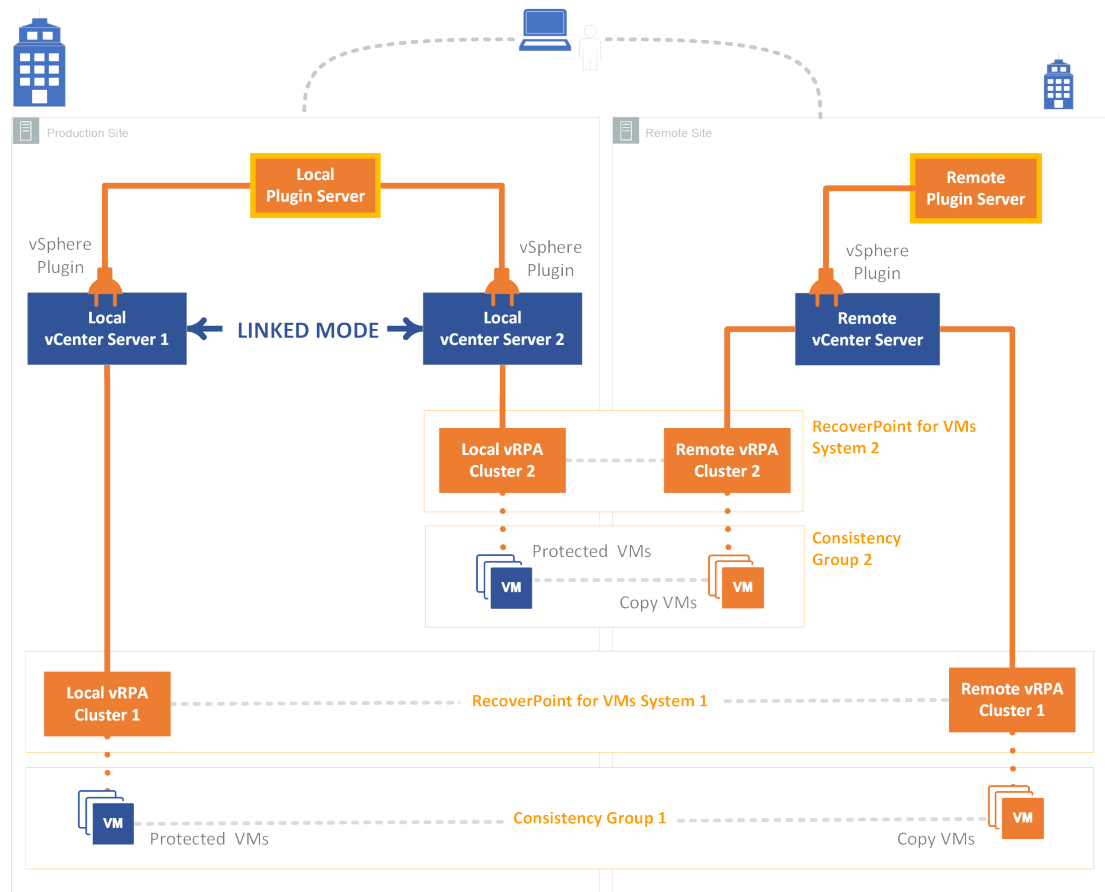

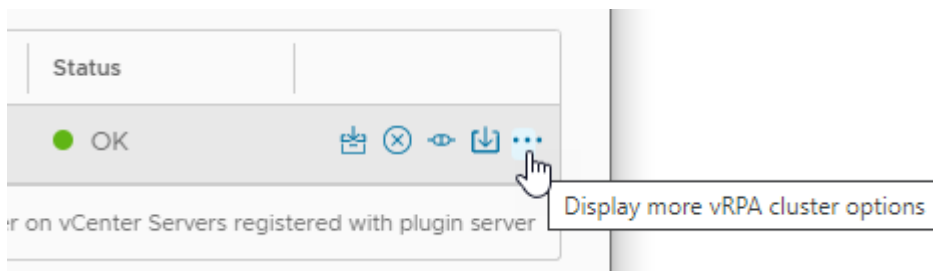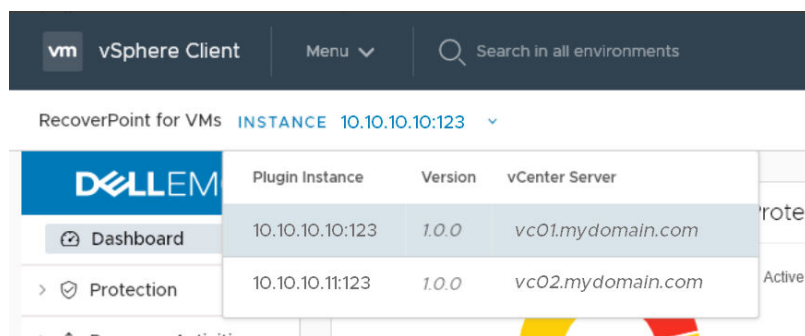**Figure 10. Remote linked vCenter example**

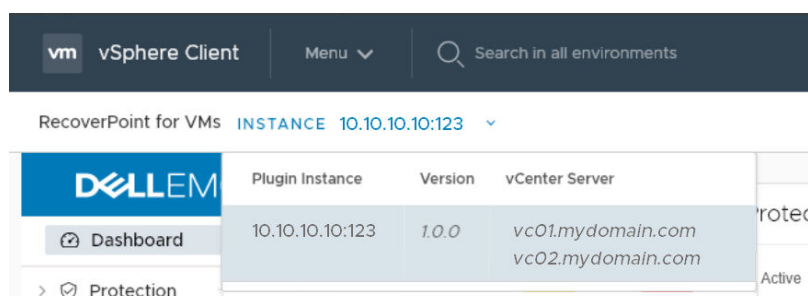As illustrated in the Remote linked vCenter example on page 75, after installing your vRPA clusters:

Register the vCenter Servers with the **Local Plugin Server** and **Remote Plugin Server** during system deployment using the **RecoverPoint for VMs Deployer**.

● Register the **Local vCenter Server** with the **Local Plugin Server**.
● Register the **Remote vCenter Server** with the **Remote Plugin Server**.

In the **vSphere plugin**, register linked vCenter (see Managing vCenter Server registration with plugin server on page 70).

● Register the **Local vCenter Server** with **Local vRPA Cluster 1** and **Local vRPA Cluster 2**.
● Register the **Remote vCenter Server** with **Remote vRPA Cluster 1** and **Remote vRPA Cluster 2**.

> ⓘ **NOTE:** To open the **Deployer** from the **vSphere plugin**, you can click **System** > **Administration** > **vRPA Clusters**, select a vRPA cluster and click **Display more vRPA cluster options**. See the *RecoverPoint for VMs Installation and Deployment Guide* for more information.

When Managing the plugin server on page 68, register the **Remote vCenter Server** with the **Remote Plugin Server**.

After registering your vCenter Servers with the plugin server and the vRPA clusters, you can see and protect VMs of both **System 1** and **System 2** from either vCenter Server.

Click **INSTANCE** to display both plugin server instances and both vCenter Servers are displayed when you are connected to the **RecoverPoint for VMs vSphere plugin** of either vCenter Server.

Displaying the plugin server info

To display the plugin server, click the plugin server **INSTANCE** at the top of the **vSphere plugin**.

About this task

When vCenter Servers are linked, all Managing linked vCenter Server registration on page 71 that have been registered with the plugin server are displayed. All plugin servers in the linked system are also displayed.

Collecting plugin server logs

Collect plugin server logs for support purposes.

About this task

This procedure collects logs from all vRPAs on all vCenter Servers registered with (or linked to one that is registered with) the plugin server, see Collecting logs from vRPA clusters on page 103 to collect logs from specific vRPA clusters.

Steps

1. Click **System** > **Administration**
   The **RecoverPoint for VMs Plugin Server Administration** screen is displayed.

2. Click **Actions** > **Collect plugin server logs**.

The collection of plugin server logs begins and a status indicator is displayed.

◯ Collecting plugin server logs...

Results

When the collection process is complete, a success message is displayed with the location of the plugin server logs.

⊘ Get the plugin server logs:
https://▮▮▮▮▮▮/logs/recover_point_logs_2020-04-
07T17_15_12Z.tar.gz   **COPY**   ✕

Next steps

Click **Copy**, open a browser window, and paste the copied URL into the browser address bar to retrieve the files.

Upgrading the plugin server

Upgrading the plugin server upgrades the HTML5 plugin and the API.

Prerequisites

Download a plugin server upgrade file from the RecoverPoint for VMs product support section of https://www.dell.com/support.

About this task

Plugin server releases are not tied to RecoverPoint for VMs releases. Upgrade packages can upgrade all services running on the plugin server. When a plugin server is being upgraded, the vSphere HTML5 plugin is not functional until upgrade is complete.

Steps

1. Click **System** > **Administration**
   The **RecoverPoint for VMs Plugin Server Administration** screen is displayed.

2. Click **Actions** > **Upgrade plugin server**
3. Select the plugin server upgrade file that you downloaded from https://www.dell.com/support, and click **OK**.

Results

During upgrade the RecoverPoint for VMs HTML5 plugin cannot be operated.

Next steps

Wait for upgrade to complete to operate your RecoverPoint for VMs system.

Managing vRPA clusters

Manage the vRPA clusters on all vCenter Servers that are registered with (or Linked to a vCenter Server registered with) the plugin server.

Monitor the state of your **vRPA clusters** in the system in The RecoverPoint for VMs Dashboard on page 26.

Use the **System** > **Administration** > **vRPA Clusters** screen to collect logs from a vRPA cluster, exclude a vRPA cluster from plugin server management, connect a vRPA cluster to another vRPA cluster, upgrade a vRPA cluster, or display more vRPA cluster options.

- To collect logs from one or more (but not all) vRPA clusters, see Collecting logs from vRPA clusters on page 103.
- To exclude a vRPA cluster from plugin server management, click the **exclude icon**.

To include all excluded vRPA clusters, click **ACTIONS** > **Clear excluded vRPA clusters**.

- To connect a vRPA cluster to another vRPA cluster, click the **Connect this vRPA cluster** icon.

The following system message is displayed:

- To upgrade a vRPA cluster, click the **Upgrade this vRPA cluster** icon.

The following system message is displayed:

- To display more vRPA cluster options click the **ellipses icon [...]** .

The following system message is displayed:

In the system message that is displayed, click **Copy**, open a browser window, and paste the copied URL into the browser address bar.

The **RecoverPoint for VMs Deployer** is displayed.

In the **RecoverPoint for VMs Deployer** home screen, to complete:

- Click **Collect Logs** and see Collecting logs from vRPA clusters on page 103.
- Click **Connect vRPA clusters** and follow the onscreen instructions.
- Click **Upgrade a vRPA cluster** and follow the onscreen instructions.
- Other vRPA cluster options, select the required option under **More actions**.

For detailed information on how to perform all vRPA cluster actions, see the *RecoverPoint for VMs Installation and Deployment Guide*.

Managing system component registration

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

About this task

After initial system configuration, manage system component registration using the options in the **System** menu.

For a detailed description of how to initially configure the RecoverPoint for VMs system, see

Register vCenter Server to vRPA cluster

Use this procedure to register a vCenter Server to a vRPA cluster.

Prerequisites

- All vCenter Servers that manage production VMs and copy VMs must be registered at the relevant vRPA cluster before you protect VMs.
- It is recommended 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.

  For more information about the location of the security certificate, refer to VMware documentation at www.vmware.com.

Steps

1. Click **System** > **vCenter Server**, and select the vRPA cluster to which you want to register a vCenter Server.

2. Click **ADD**.

Register vCenter Server   ✕

vCenter Server
Enter IP address or host name

Port
443

Username
Enter username

Password
Enter username password

CANCEL   REGISTER

3. Enter the IP and credentials of the vCenter Server to be registered.
4. Click **REGISTER**.

Results

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

Next steps

Ensure that a plugin server is installed on the newly registered vCenter Server. For a linked vCenter Server, see Managing vCenter Server registration with plugin server on page 70.

Managing ESX cluster registration

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

About this task

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

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **System** > **ESX Clusters**.

2. If you are replicating remotely, select the vRPA cluster at which you want to register ESX clusters.
3. Click **Add**.

4. In the **Register ESX Cluster** dialog box:
   a. Select the ESX cluster that you want to register.
   b. Click **REGISTER**.

Results

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

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

Next steps

Registered ESX clusters can be (1) updated or (2) deleted. To validate ESX cluster registration, see the *RecoverPoint for VMs Flex Plugin Administrator's Guide*.

Managing journal datastore registration

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

About this task

(i) **NOTE:** When you protect a consistency group, the **Protect VMs Wizard** will automatically select a datastore from the list of registered datastores, unless you specify a specific registered datastore to use. RecoverPoint for VMs will attempt to create the journal on the selected datastore. If for any reason journal creation fails, the system will attempt to create the journal on a different registered datastore.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **System** > **Datastores**.

2. If you are replicating remotely, select the vRPA cluster at which you want to register datastores, and click **Add...** .
   The **Register Datastore** dialog box is displayed.
3. In the **Register Datastore** dialog box:
   a. Select one or more datastores to register.
   b. Click **REGISTER**.

Results

The datastore is registered at the specified vRPA cluster.

Next steps

Registered journal datastores can be deleted.

(i) **NOTE:** A datastore with a scratch partition on its root path must not be used to host journals.

Managing external host registration

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

Prerequisites

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

Steps

1. In the vSphere HTML5 plugin, select **System** > **Orchestration**, and select the vRPA cluster for which you want to define an external host.
2. Click **ADD** under the **External Host** widget.
3. In the **Register External Host** dialog box, type the **Name**, **IP**, **User**, and **Password** of the external host for the selected vRPA cluster.
4. Optionally:
   - To verify connectivity with the external host, click **Check Connectivity** .
   - To unregister the external host from the specified vRPA cluster, click **Remove**.

Managing protected VMs

This section describes how to manage the protection of protected VMs.

About this task

After initial protection, use the RecoverPoint for VMs vSphere HTML5 plugin to manage VMs either through **Protection** > **Protected VMs** or **Protection** > **Consistency Groups**. For a detailed description of how to protect VMs, see Protecting VMs.

Managing VM protection policies

Update the protection policies of protected virtual machines.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client , select **Protection** > **Protected VMs**.
2. Select a virtual machine, and click **PROTECTION POLICY**.
   The VM **Protection Policy** dialog is displayed.

Protection Policy of VM 'Windows'   ✕

Disk Provisioning | Same as source   ▼

🔵 Replicate VM hardware changes

⚪ Replicate MAC addresses to the local copy

🔵 Automatically protect newly added VMDKs

Protected VMDKs

☑	Protected VMDK	↑	Path	Size
☑	Hard disk 1		SCSI (0:0)	20 GB
☑	Hard disk 2		SCSI (0:1)	200 MB

CANCEL   **UPDATE POLICY**

3. Update the VM protection policies:
   - **Disk provisioning**: Default is `Same as source`. Defines the way in which the copy VMDKs are to be provisioned; `Same as source`, `Thick provision lazy zeroed`, `Thick provision eager zeroed` or `Thin provision`.
   - **Replicate VM hardware changes**: Default is `Enabled`. Automatically replicates the hardware settings of all production virtual machines to their copy VMs whenever an image is accessed on the copy VMs. When enabled, RecoverPoint for VMs replicates the virtual machine version, MAC address, CPU, memory, resource reservations, and network adapter status and type. Replication of SR-IOV Passthrough Adapter is not supported. If the ESX at a copy does not support the production VM version, no hardware resources are replicated.
   - **Replicate MAC addresses to local copy**: Default is `Disabled`. If two *remote copies* of the same production VM are on the same vCenter and in the same network, you cannot power on both copy VMs simultaneously, as they will both have the same MAC address. Therefore, by default, the MAC address of remote copy VMs network adapters (NICs) on a different vCenter than their production VMs are replicated to the copy. However:
     - When **Replicate VM hardware changes** is disabled, MAC address replication to the remote copies is also disabled.
     - To avoid IP conflicts, by default, the MAC addresses are not replicated to the local copy VMs on the same vCenter as their production VMs. If a local copy VM is not on the same network and ESX as its production VM, enable **Replicate MAC addresses to local copy** to replicate the MAC addresses.
   - **Automatically protect newly added VMDKs**: Default is `Enabled`. Automatically includes any VMDKs that are added to a VM, after it is already protected.
   - **Protected VMDKs:** Displays the number of VMDKs that will be replicated, their **Path**, and their total size. Clear a VMDK checkbox to exclude the VMDK from replication.
4. Click **UPDATE POLICY**.

Results

The VM protection policies are updated.

Stop protecting a VM

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

Steps

1. In the RecoverPoint for VMs vSphere HTML5 plugin to manage VMs either through **Protection** > **Protected VMs**
2. Select the production VM that you want to stop protecting.
3. Click **UNPROTECT**

Alternatively, to unprotect all VMs in a consistency group, select **Protection** > **Consistency Groups**, select the consistency group, and from the **[...]** menu, select **Unprotect**.

Results

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

Managing consistency groups

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

About this task

After initial protection, use the RecoverPoint for VMs vSphere HTML5 plugin to manage consistency groups through the **Protection** > **Consistency Groups** screen.

Managing group protection policies

Update the protection policies of consistency groups and their copies.

Steps

1. In the RecoverPoint for VMs plugin for **vSphere Client (HTML5)**, select **Protection** > **Consistency Groups**.
2. Select a consistency group, and click **PROTECTION POLICY**.
   The group **Protection Policy** dialog is displayed.

Protection Policy of Group 'MyGroup'   ×

General   Production (Site1)   Remote 1 (Site2)

┌───┐
│ **GROUP POLICY**   VM STARTUP SEQUENCE   │
│   │
│ Consistency Group   │
│ MyGroup   │
│   │
│ Primary vRPA   │
│ vRPA 2 ⌄   │
│   │
│ Bandwidth Priority   │
│ Normal ⌄   │
└───┘

                                          CANCEL   **UPDATE POLICY**

---

3. Update the group policies:
   a. Select **General** > **GROUP POLICY**.
      ● **Consistency Group**: The name of the consistency group in the RecoverPoint for VMs system. Default is `cg_<vmname>`.
      ● **Primary vRPA**: The vRPA that you prefer to replicate the consistency group. When the primary vRPA is not available, the consistency group will switch to another vRPA in the vRPA cluster. When the primary vRPA becomes available, the consistency group will switch back to it.

> (i) **NOTE:** If your vRPA cluster is the only vRPA cluster in the system, it is a single point of failure in cases of disaster. Consider adding additional vRPAs to the vRPA cluster to ensure high availability.

- **Bandwidth Priority**: Only relevant for remote replication when two or more consistency groups are using the same **Primary vRPA**. Default is `Normal`. Select the bandwidth priority to assign to this consistency group. The priority determines the amount of bandwidth allocated to this group in relation to all other groups using the same primary vRPA.

b. To define the order in which VMs in a consistency group will power on during testing and recovery. select **General** > **VM STARTUP SEQUENCE**, and refer to VM start-up sequence on page 42.

4. Update the group production policies:

> (i) **NOTE:** The production policies are only applied after failover.

Select the **Production** tab and click **PRODUCTION POLICY** and **RE-IP RULES** to define the production policies to be applied after failover (when the production becomes a copy). Refer to the next step for a detailed description of the production policies.

5. Update the group copy policies:

a. Select a copy and click the **COPY POLICY** tab to update the copy protection settings:

- **Required retention policy**: Default is `disabled`. Defines how far in time the copy image can be rolled back. Enable to define a retention policy in `minutes`, `hours`, `days`, `weeks`, or `months`. An alert is displayed if the copy image cannot be rolled back according to the required retention policy.
- **Consolidate RecoverPoint for VMs snapshots**: Default is `disabled`. Automatic snapshot consolidation cannot be enabled for a group that is part of a group set.
  - **Do not consolidate any snapshots for at least**: Default is `2 days`. Can be defined in `hours`, `days`, `weeks`, or `months`. Defines 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 is `5 days`. Snapshots are consolidated every 24 hours. Select Indefinitely to consolidate all subsequent snapshots in 24-hour intervals.
    - If **Indefinitely** is not selected, and no weekly consolidations are specified, the remaining snapshots are consolidated monthly.
    - If **Indefinitely** is selected, weekly and monthly consolidations are disabled, and the remaining snapshots are consolidated daily.
  - **Consolidate snapshots that are older than x to one snapshot per week for y weeks**: Default is `4 weeks`. Snapshots are consolidated every 7 days. Select Indefinitely to consolidate all subsequent snapshots in seven-day intervals.
    - If **Indefinitely** is not selected, the remaining snapshots are consolidated monthly.
    - If **Indefinitely** is selected, monthly consolidations are disabled, and the remaining snapshots are consolidated weekly.
- **Journal Volumes**: Displays the size and datastore of each journal volume, and allows you to add or remove journal volumes.

○ (Production journal only) This option is only displayed if more than one journal volume was added to the production journal. Click **RESET SIZE** if journal volumes were added to the production journal after a temporary failover. After failing back to production, use this button to reset the production journal to its original size (by default, **3GB**) without triggering a full sweep.

○ Click the **delete icon** to remove a journal volume. The last journal volume at a copy cannot be deleted.

⚠ **CAUTION: Removing a journal volume causes a full sweep on all VMs in the consistency group. The full sweep duration depends on the size of the data being replicated, network resources, and storage performance.**

Click **ADD** to add volumes to the copy journal. The **Add Journal Volume** dialog is displayed.

ⓘ **NOTE:** You can safely click **ADD** now, as the default settings provide a sensible configuration for most systems.

- **Journal Size**: Default is **10GB** for the copy journals and **3GB** for the production journal. The default size of the production journal is smaller, and in the vast majority of cases, will not require any additional volumes. The larger a copy journal, the more history can be saved.

- **Select Registered Datastore**: Default is `Automatically`. By default, RecoverPoint for VMs automatically registers up to 15 datastores for the production and copy journals and automatically selects the datastore with the most free space. When set to `Manually`:

- **Registered Datastores**: Select a registered datastore or register a new datastore for the journal. By default, RecoverPoint for VMs automatically registers up to 15 datastores for the group journals and automatically selects the datastore with the most free space. RecoverPoint for VMs will attempt to create

a journal volume on the selected datastore. If it cannot, the system will attempt to create the journal volume on another registered datastore.

- If you have more than 15 datastores and would like to register a datastore that is not in the list, click **REGISTER DATASTORE** to register an additional datastore.

b. Select a copy and click the **RE-IP RULES** tab to create Re-IP rules on page 47 to update the network configuration of copy VMs during testing and recovery.

c. (Not relevant for the production) Select a copy and click the **LINK POLICY** tab to update the copy link protection settings:

- **Async** or **Sync**: Default is `Async`. Defines the way in which data is replicated from the production to the copy. Data can be replicated synchronously (`sync`) or asynchronously (`async`). When `sync` is selected, you can also define the following policies:
  - **Dynamic by latency**: Default is `disabled`. When enabled, RecoverPoint for VMs alternates between synchronous and asynchronous replication modes, as necessary, according to latency conditions.
    - **Start async replication above**: When the specified limit (in `milliseconds`) is reached, RecoverPoint for VMs automatically starts replicating asynchronously.
    - **Resume sync replication below**: When the specified limit (in `milliseconds`) is reached, RecoverPoint goes back to replicating synchronously.
  - **Dynamic by throughput** Default is `disabled`. When enabled, RecoverPoint for VMs alternates between synchronous and asynchronous replication modes, as necessary, according to throughput conditions.
    - **Start async replication above**: When the specified limit (in `MB`) is reached, RecoverPoint for VMs automatically starts replicating asynchronously.
    - **Resume sync replication below** : When the specified limit (in `MB`) is reached, RecoverPoint goes back to replicating synchronously.
- **RPO**: Defines the maximum lag allowed on a link.
- **WAN Compression**: Default is `enabled`. Only relevant for asynchronous remote replication. To compress data before transferring it to a remote vRPA cluster, select a level of compression. Enabling and disabling compression causes a short pause in transfer and a short initialization. Compression can reduce transfer time significantly, but increases the source vRPA's CPU utilization.
- **Deduplication**: Default is `disabled`, but deduplication can be enabled whenever compression is enabled. Eliminates repetitive data before transferring the data to a remote vRPA cluster. Enabling and disabling deduplication causes a short pause in transfer and a short initialization. Deduplication can reduce transfer time significantly, but increases the source vRPA's CPU utilization.

d. (Not relevant for the production) Select a copy and click the **FAILOVER NETWORKS** tab to configure Failover networks on page 50 to automatically associate the VM network adapters (vNICs) of a copy VM with specific port groups upon failover or during copy testing.

6. Click **UPDATE POLICY**.

Results

The group protection policies are updated.

Managing group sets

Manage group sets in RecoverPoint for VMs.

About this task

A group set is a collection of consistency groups that you can bookmark, enable, disable, pause and resume replication for, and test and recover as a group. You can also create parallel bookmarks on all groups in the group set, at a frequency that you define. Group sets are useful for consistency groups that are dependent on one another or that must work together as a single unit.

Steps

1. Select **Protection** > **Group Sets**.

   The **Group Sets** screen is displayed.

2. In the **Group Sets** screen:
   ● Note the number of consistency groups in the group set, the parallel bookmarking status and the vRPA cluster from which all groups in the group set are replicating.
   ● Expand a group set to display the names of the consistency groups that are currently in the group set and their properties.
   ● Select a group set, and:
     ○ Click **ADD** to create another group set. See Create a group set on page 46 for more information.
     ○ Click **UPDATE** to modify the group set configuration. See Create a group set on page 46 for more information.
     ○ Click **BOOKMARK** to apply a label and/or consolidation policy to all copy VMs of all consistency groups in the group set. See Create a bookmark on page 39 for more information.
     ○ Click **TEST A COPY**, **RECOVER PRODUCTION**, or **FAILOVER** to test a copy, failover, or recover production of all consistency groups in the group set. See Recovering VMs on page 52 for more information.
     ○ Click the more commands **[...]** button to display additional group set commands:

- **Group priority**: Set the power-on priority of all consistency groups in the group set. See Group start-up sequence on page 45 for more information.
- **Disable groups** or **Enable groups**: Disables or enables all copies of all consistency groups in the group set.

  ⚠ **CAUTION: Disabling a group stops all copy activities, and deletes all copy journals. A full sweep is required when transfer resumes.**
- **Pause transfer** or **Resume transfer**: Pauses or resumes transfer of all copies of all consistency groups in the group set.

Managing recovery activities

Manage ongoing testing, failover, failback and production recovery activities using the **Recovery Activities** screen.

Use the **Recovery Activities** screen to:

- Manage testing on page 92.
- Manage failover and failback on page 96.
- Manage production recovery on page 99.
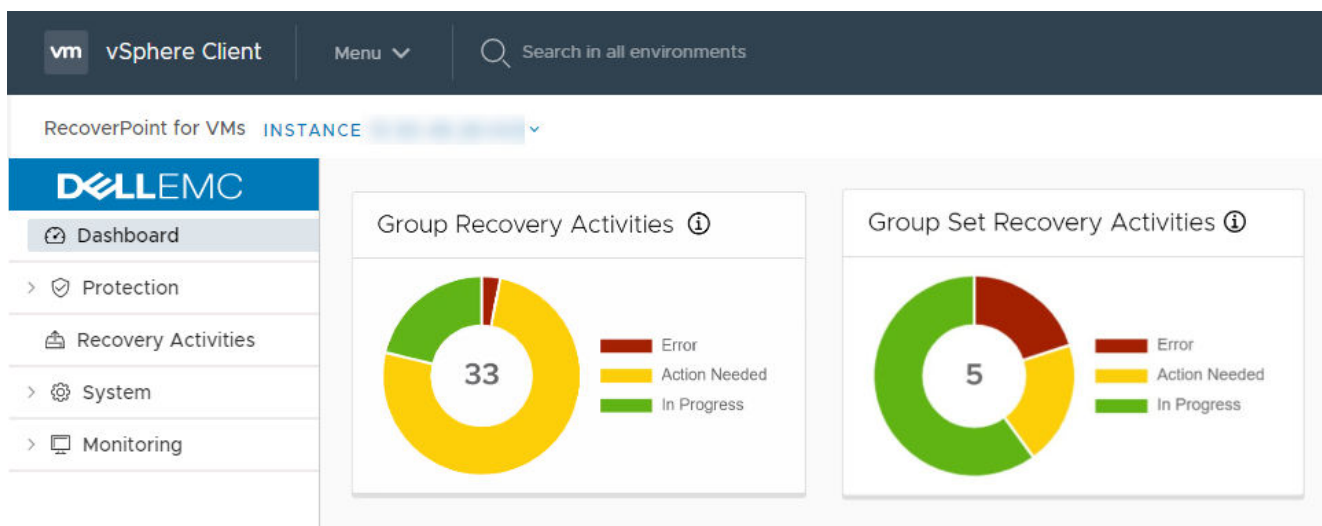
To manage an ongoing recovery activity of a consistency group, select **Recovery ActivitiesConsistency Groups**.

To manage an ongoing recovery activity of a group set, select **Recovery ActivitiesGroup Sets**.

Use the The RecoverPoint for VMs Dashboard on page 26 to monitor the state of ongoing recovery activities, for both consistency groups and group sets.

Manage testing

Manage the ongoing testing of a copy of a consistency group or group set using the **Recovery Activities** screen.

(i)**NOTE:**

When testing a copy, wait for the **Activity Status** to show **Ready for next action** and the **Progress status bar** to reach 100%, indicating the specified snapshot image has been accessed.

● If you tested a copy of a consistency group:

The **Consistency Group Recovery Activities** screen is displayed. The **Activity Status** and **Progress** columns indicate the progress of image access. After access is enabled to the copy snapshot, the **Activity Status** column displays **Ready for next action**, and you can:

(i)**NOTE:** 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. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies on page 86.

○ Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.

○ Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

⚠ **CAUTION: When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.**

○ Click **ACTIONS** > **Start new test** to select another snapshot to test, or to redefine the testing network.

○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.

- Click **ACTIONS** > **Promote image: Failover** (5.3.1 or later) to jump directly to the Failover stage of failover to the copy image that you just tested (step 8 on page 61), without needing to roll back the writes that were made to the copy snapshot while write access was enabled.
  - Click **ACTIONS** > **Promote image: Recover Production** (5.3.1 or later) to jump directly to the production recovery stage of recovering production from the copy image that you just tested (step 8 on page 66), without needing to roll back the writes that were made to the copy snapshot while write access to the copy volumes was enabled.
- If you tested a copy of a group set:

  - Click **OPEN** to display a **Detailed Status** for each group in the group set. The **Detailed Status** screen is displayed.
    ⓘ **NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

  - Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.
- If you tested a copy of a consistency group:

The **Consistency Group Recovery Activities** screen is displayed. The **Activity Status** and **Progress** columns indicate the progress of image access. After access is enabled to the copy snapshot, the **Activity Status** column displays **Ready for next action**, and you can:

> (i) **NOTE:** 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. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies on page 86.

- ○ Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- ○ Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

  > ⚠ **CAUTION: When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.**

- ○ Click **ACTIONS** > **Start new test** to select another snapshot to test, or to redefine the testing network.
- ○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- ○ Click **ACTIONS** > **Promote image: Failover** (5.3.1 or later) to jump directly to the Failover stage of failover to the copy image that you just tested (step 8 on page 61), without needing to roll back the writes that were made to the copy snapshot while write access was enabled.
- ○ Click **ACTIONS** > **Promote image: Recover Production** (5.3.1 or later) to jump directly to the production recovery stage of recovering production from the copy image that you just tested (step 8 on page 66), without needing to roll back the writes that were made to the copy snapshot while write access to the copy volumes was enabled.

● If you tested a copy of a group set:

- ○ Click **OPEN** to display a **Detailed Status** for each group in the group set. The **Detailed Status** screen is displayed.

  > (i) **NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

Detailed Status of 'group-set'

- ○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.
- ● If you tested a copy of a consistency group:

The **Consistency Group Recovery Activities** screen is displayed. The **Activity Status** and **Progress** columns indicate the progress of image access. After access is enabled to the copy snapshot, the **Activity Status** column displays **Ready for next action**, and you can:

> ⓘ **NOTE:** 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. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies on page 86.

- ○ Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- ○ Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

  > ⚠ **CAUTION: When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.**

- ○ Click **ACTIONS** > **Start new test** to select another snapshot to test, or to redefine the testing network.
- ○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- ○ Click **ACTIONS** > **Promote image: Failover** (5.3.1 or later) to jump directly to the Failover stage of failover to the copy image that you just tested (step 8 on page 61), without needing to roll back the writes that were made to the copy snapshot while write access was enabled.
- ○ Click **ACTIONS** > **Promote image: Recover Production** (5.3.1 or later) to jump directly to the production recovery stage of recovering production from the copy image that you just tested (step 8 on page 66), without needing to roll back the writes that were made to the copy snapshot while write access to the copy volumes was enabled.

- If you tested a copy of a group set:

- ○ Click **OPEN** to display a **Detailed Status** for each group in the group set. The **Detailed Status** screen is displayed.
  - (i) **NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

- ○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

After finding a suitable snapshot, Create a bookmark on page 39 to label the snapshot so it is easily identifiable for recovery.

Manage failover and failback

Manage the ongoing failover or failback of a copy of a consistency group or group set using the **Recovery Activities** screen.
(i) **NOTE:**

When failing over (or failing back) to a copy, wait for the **Activity Status** to show **Ready for next action** and the **Progress status bar** to reach 100%, indicating the specified snapshot image has been accessed.

When the **Activity Status** is **Ready for next action**:
- To select a consistency group for failover, ensure the **Consistency Groups** tab is selected.

NOTE: By default, replication starts immediately after failover. In RecoverPoint for VMs 5.3.1 and later versions, disable **Start transfer** before failing over to pause replication after failover.

○ Click **ACTIONS** > **Start new test** to select another snapshot to test, or to redefine the testing network.

○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.

○ (Optional) Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.

○ (Optional) Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

⚠ **CAUTION: When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.**

○ In **Failover Networks**, you can use the default pre-configured failover networks, by keeping **Use or edit pre-configured failover networks** selected. You can also edit a pre-configured network, or choose to **Use current testing networks**.

NOTE: 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. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies on page 86.

● To select a group set for failover, click the **Group Sets** tab.

○ Click **OPEN** to display the **Detailed Status** of all consistency groups in the group set. After access is enabled to the copy snapshot, the **Status** column of all groups displays **Ready for next action**.

NOTE: Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

Detailed Status of 'group-set'   ✕

Consistency Groups	Copy	Snapshot	Status	Progress
cg_Win_286	Local Copy	May 26, 2020 2:04:2...	Ready for next action	▬▬▬ 100%
cg_Win	Standalone	May 26, 2020 2:06:1...	Ready for next action	▬▬▬ 100%

Items per page 10 ⌄   2 Detailed Statuses

CLOSE

- ○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

ⓘ **NOTE:** After finding a suitable snapshot, you may want to Create a bookmark on page 39 to label the snapshot so it is easily identifiable during failover.

Click **ACTIONS** > **Failover** to failover to the copy

- If the selected consistency group or group set has only one copy, failover starts.
  - ○ The role of the **Production** becomes **Remote/Local Copy**.
  - ○ The role of the **Remote/Local Copy** becomes **Production**.
  - ○ The production VM and copy VM change roles, but their names do not change. Therefore, after failover, new production VMs will still be named `YourVMName.copy` and the new copy VMs are still named `YourVMName`.
  - ○ The production journal becomes the copy journal and the copy journal becomes the production journal. You may want to add journal volumes as described in Managing group protection policies on page 86.
  - ○ The marking information in the production journal is deleted, the copy journal is deleted, and the consistency group undergoes a full sweep.

    ⚠ **CAUTION: During the full sweep, data is not transferred synchronously.**

- If the consistency group or group set has copies other than the copy to which you are failing over (even if they are disabled or replication to them is paused), a temporary failover begins:
  - ○ The role of the **Production** changes to **Temporary Production**.
  - ○ The role of the **Remote/Local Copy** changes to **Temporary Remote/Local Copy**.
  - ○ The roles of any other (unlinked) copies become **Standalone**.
  - ○ Replication pauses for the other copies and the direction of replication between the production and the failed-over copy changes.

After temporary failover, if your consistency group or group set had more than one copy (even if they are disabled or replication to them is paused), in the **Recovery Activities** screen:

- **Failback to the original production**. Select the recovery activity, click **ACTIONS** > **Test for failback** and run the above procedure beginning with step 3, substituting "failback" for "failover" throughout.

  After failing back to the production, if you added volumes to the production journal after failover, to reset the production journal to its original size (by default, 3 GB) without triggering a full sweep click **Protection** > **Consistency Groups** > **PROTECTION POLICY**, select the group's **Production** copy, and click **RESET SIZE** in the **Journal Volumes** section.

- **Set the copy as the new production**. Select the recovery activity and click **ACTIONS** > **Set as production**. If there are standalone (unlinked) copies, the **Set this Copy as the New Production** dialog is displayed.

  In the **Set this Copy as the New Production** dialog for consistency groups:

Set this Copy as the New Production ✕

Designate this copy at vRPA cluster 'Darwin' as the new
production of consistency group 'cg_newVm'.

Before you can permanently failover to this copy, decide what you want to do
with the other copies in this group.

For each other copy in the group, define a new copy protection policy, and
whether you want to disable or delete the copy after permanent failover.

| Standalone 2 (Patagonia) | Async ⬤ Sync | Disable ⬤ Enable | 🗑 |
| Standalone 2 (Darwin) | Async ⬤ Sync | Disable ⬤ Enable | 🗑 |

CANCEL **SET AS PRODUCTION**

1. Configure each standalone copy for consistency groups (or all standalone copies for group sets).

   Standalone copies are not linked to the production, and you must decide how to handle them before failover. By default, RecoverPoint for VMs does not delete copy VMs but it does disable them. You can **Enable** any required standalone copies and select a replication mode (sync or async), or **Delete** them from the consistency group. Deleting a copy does not delete the VMs from storage.

   ⚠ CAUTION: **Disabled copy VMs require a full sweep when they are re-enabled.**

2. Click **SET AS PRODUCTION** to permanently failover.
   ○ The role of the **Production** becomes **Remote/Local Copy**.
   ○ The role of the **Remote/Local Copy** becomes **Production**.
   ○ The standalone copies are handled as specified.
   ○ The production VM and copy VM change roles, but their names do not change. Therefore, after failover, new production VMs will still be named *YourVMName*.copy and the new copy VMs are still named *YourVMName*.
   ○ The production journal becomes the copy journal and the copy journal becomes the production journal. The production journal does not contain the copy history, so it is by default, a much smaller journal. Therefore, after failover, when the production becomes the copy, you may want to add journal volumes to the new copy journal to ensure that you have ample space for copy testing. For detailed instructions on how to add journal volumes to a copy journal, see Managing group protection policies on page 86
   ○ The marking information in the production journal is deleted, the journal of the copy to which you failed over is deleted, and the consistency group undergoes a full sweep.
     ⚠ CAUTION: **During the full sweep, data is not transferred synchronously.**

Manage production recovery

Manage the ongoing recovery of a copy of a consistency group or group set from the production using the **Recovery Activities** screen.

ⓘ NOTE: When recovering of a copy from the production, wait for the **Activity Status** to show **Ready for next action** and the **Progress status bar** to reach 100%, indicating the specified snapshot image has been accessed.

When the **Activity Status** is **Ready for next action**:

● To select a consistency group for production recovery, ensure the **Consistency Groups** tab is selected.

○ Click **ACTIONS** > **Start new test** to select another snapshot to test, or to redefine the testing network.
○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
○ (Optional) Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
○ (Optional) Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

⚠ **CAUTION:** When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.

ⓘ **NOTE:** 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. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies on page 86.

● To select a group set for production recovery, click the **Group Sets** tab.

○ Click **OPEN** to display the **Detailed Status** of all consistency groups in the group set. After access is enabled to the copy snapshot, the **Status** column of all groups displays **Ready for next action**.

ⓘ **NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

Detailed Status of 'group-set'

Consistency Groups	Copy	Snapshot	Status	Progress
cg_Win_286	Local Copy	May 26, 2020 2:04:2...	Ready for next action	100%
cg_Win	Standalone	May 26, 2020 2:06:1...	Ready for next action	100%

Items per page 10 ⌄   2 Detailed Statuses

CLOSE

○ Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

To recover production from the copy click **ACTIONS** > **Recover production**.

ⓘ **NOTE:**
- Data transfer from the production to all copies is paused, and will resume only after production recovery is complete.
- Host access to the recovered production volumes, and the recovering copy volumes is blocked.
- Recovered production volumes are overwritten. Any writes made to the copy during testing are transferred to the production, unless you clicked **UNDO WRITES**.
- The group undergoes a short initialization process to synchronize the new production data at the copy.

Configuring email alerts and reports

Run a series of **Sysmgmt CLI** commands to configure your system to send system alerts and reports to a specified email, in real time.

Prerequisites

Create an SSH connection to a vRPA cluster management IP address, and use your RecoverPoint for VMs **admin** username and password to log into the **Boxmgmt CLI** > **Sysmgmt CLI**.

About this task

ⓘ **NOTE:** For more information, refer to the *RecoverPoint for VMs CLI Reference Guide*.

Steps

1. Set up email access using the Sysmgmt CLI:
   a. Run the **set_smtp_server** command, and enter the IP address or DNS name for sending email notifications.
      You should receive confirmation that the SMTP server has been configured successfully.
   b. Run the **config_email** command, and provide the requested information.
      You should receive confirmation that your email mechanism has been successfully configured.
   c. Run the **enable_email** command, and choose the **enable_email** option.
      You should receive confirmation that email alerts have been enabled successfully.
2. (Optional) To add email users, run the **add_email_users** command, and provide the requested information.
   You should receive confirmation that email users have been added successfully.
3. (Optional) To configure system reports, in the Sysmgmt CLI, run the **config_system_reports** command, and provide the requested information.
   You should receive confirmation that system notifications have been successfully configured.

Results

Your email alerts and reports are configured.

Next steps

See the Monitor system alerts on page 30 for more information.

Troubleshooting

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

Topics:

Finding the vRPA cluster management IP

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

Steps

1. In the vSphere HTML5 plugin, select **System** > **Administration**, and the **vRPA Clusters** tab.
2. Select the vRPA cluster.
3. Note the **Management IP Address** for the selected vRPA cluster.

Collecting logs

Collecting logs is relevant only in support cases, and should be performed only when instructed to do so by Customer Support.

Collecting logs from vRPA clusters

In RecoverPoint for VMs 5.3 SP2 and later versions, you can initiate the log collection process from within the RecoverPoint for VMs vSphere plugin.

About this task

Logs can be collected from one or more vRPAs in multiple vRPA clusters, as long as they all reside on a vCenter Server registered with the plugin server, or a vCenter Server linked to a vCenter Server that is registered with the plugin server.

ⓘ **NOTE:** In RecoverPoint for VMs versions prior to 5.3 SP2, see the *RecoverPoint for VMs Installation and Deployment Guide* "Collect logs" section for information on how to collect logs from vRPA clusters.

Steps

1. Click **System** > **Administration** > **vRPA Clusters**.
2. Select a vRPA cluster and click the **Collect logs** icon.

3. In the system notification that is displayed, click **Copy**, open a browser window and paste the copied URL into the browser address bar.

4. If prompted, type the login credentials for the **admin** user and click **Sign in**.
5. In the **RecoverPoint for VMs Deployer**, click **Collect Logs**.

The **Collect Cluster Logs** dialog is displayed.

In the **Collect Cluster Logs** dialog:

   a. Enter a time frame for log collection.
   b. (Optional) Enter the IP addresses of other vRPA clusters from which to collect logs in the **Advanced** section.
   c. Click **Collect Logs**.

Results

Depending on the size of the environment, log collection may take several minutes to complete. When the collection process is complete, a success message is displayed with the location (i.e. vRPA cluster) containing the logs.

Logs collected successfully for:

Location	Log file
Site1	sysInfo-incomplete-Site1_KBox-1-2-2021.07.12.15.29.35.tar

To download the logs, please use your 'admin' credentials

Next steps

1. In the success message, click the name of a vRPA cluster to open a browser window to the location of the collected logs.
2. If prompted to, log into the vRPA cluster with your `admin` user credentials.
3. Click a vRPA log name to download the vRPA log.

   The name of each vRPA log has a `*.tar` extension and it includes the `<clustername><vrpaname>` and `<vrpaip>` for easy identification. The log collection date is displayed under **Last Modified**.

Directory Listing For [/]

Filename	Size	Last Modified
ic_report	4.4 kb	Mon, 12 Jul 2021 15:32:29 GMT
sysInfo-incomplete-Site1_KBox-1-███████████.tar	198360.0 kb	Mon, 12 Jul 2021 15:32:29 GMT
long_term_stats/		Mon, 12 Jul 2021 15:30:49 GMT

Collect plugin server logs

Collect plugin server logs for support purposes.

About this task

The procedure for collecting plugin server logs is detailed in Collecting plugin server logs on page 76.

Collecting RecoverPoint for VMs splitter logs

About this task

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

Steps

1. In the vSphere Client, select **Menu** > **Home and Clusters**.
2. Right click on the desired vCenter, and select **Export System Logs...**.
3. In the **Select hosts** pane of the **Export System Logs** screen:
   a. Select the ESXi hosts for which you want to export the system logs.
   b. (Optional) Select **Include vCenter and vSphere UI Client logs**.
   c. Select the system logs to be exported
   d. (Optional) Select **Gather performance data**, and specify a `duration` and `interval`
   e. (Optional) Set a password with which to encrypt the collected log data.
4. In the **Select logs** pane of the **Export System Logs** screen:
   a. Select the system logs to be exported

b. (Optional) Select **Gather performance data**, and specify a `duration` and `interval`

c. (Optional) Set a password with which to encrypt the collected log data.

5. Click **EXPORT LOGS**.

Adding new VMDKs

Best practices and system behavior when adding new VMDKs to a protected VM.

About this task

RecoverPoint for VMs automatically detects when a new VMDK is added to a protected VM through the vSphere Client VM Properties, and by default, automatically starts protecting each added VMDK.

● To disable automatic protection for all VMDKs added to a protected VM in the future, see Automatic protection of newly added VMDKs on page 41.

● To exclude specific VMDKs of a protected VM from protection, see Excluding a VMDK from replication on page 41.

Results

When a new VMDK is added to a protected VM, a volume sweep occurs on the added VMDK and a short initialization occurs on all other VMDKs in the consistency group, but no history is lost.

Removing a VMDK

Defines how to handle a VMDK which was removed from a protected VM, at the copy.

About this task

RecoverPoint for VMs automatically detects when a VMDK is removed from a protected VM through the vSphere Client, and displays an alert when there is hardware mismatch between a protected VM and its copy VM.

● If the production VMDK removal was intentional, follow the instructions in Excluding a VMDK from replication on page 41 to stop protecting it.

● If the production VMDK removal was unintentional, run #unique_85 to recover the removed VMDK. For recover production, select a snapshot that pre-dates VMDK removal.

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

Automatically expanding copy VMDKs

Best practices, troubleshooting, system behavior, and limitations of automatically expanding copy VMs when a protected VM is expanded.

About this task

When a protected VMDK is expanded, RecoverPoint for VMs automatically expands all corresponding copy VMDKs, with the following limitations:

● VMDKs can be expanded, but they cannot be shrunk.

● When a production VMDK is expanded, the system pauses replication of the consistency group while the system is busy resizing the corresponding copy VMDK.

● Automatic VMDK expansion fails if:
   ○ Replicating to RDM. After expanding the production VMDK or RDM, you must manually expand the copy RDM.
   ○ The datastore does not contain enough free space, and you should free up space in the copy VM datastore.
   ○ A snapshot has been taken of the virtual machine containing the copy VMDK. Enable access to the copy containing the VMDK and use the vCenter snapshot manager to delete all snapshots before disabling image access.
   ○ The version of the file system that you are running does not support the VMDK size. In this case, consider upgrading the file system version.

- The size of a copy VMDK is larger than the size of its corresponding production VMDK. In this case, to begin the automatic VMDK expansion process, you must manually expand the production VMDK. This manual expansion might be required if you failed over while automatic expansion was in progress, or if the copy VMDK was manually expanded.
- Replicating to RDM. After expanding the production VMDK or RDM, you must manually expand the copy RDM for replication to resume.
- A snapshot of a copy containing a VMDK marked for automatic expansion is selected during testing or recovery. In this case, you should disable image access for replication to resume.
- A protected VMDK is smaller than the size registered in the system settings. In this case you should contact Customer Support. This can happen, for example, if a production VMDK has been removed and re-added with a smaller size.
- One or more copy VMDKs have been marked for automatic expansion, but the system cannot automatically resize a RAW device. In this case, enable access to the copy VM with the problematic VMDK and manually expand it before disabling image access. If problem persists, contact Customer Support.

Results

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

Recovering from a cluster disaster

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

RecoverPoint for VMs licensing

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

VM-based licensing

With VM-based licensing, licenses are based on the number of supported VMs per vCenter Server. Only production VMs are counted in the number of supported VMs per vCenter Server. Licensing is enforced using the vCenter Server ID.

All vCenter Servers must be registered in RecoverPoint for VMs before their licenses can be added. vCenter Server registration is performed in the RecoverPoint for VMs Deployer UI. Refer to the *RecoverPoint for VMs Installation and Deployment Guide* for more information.

You can use the HTML5 plugin to register additional vCenter Servers. For details, see Register vCenter Server to vRPA cluster on page 81.

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

Failover has no effect on the license.

Socket-based licensing

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

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

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

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

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

```
Capacity = 30 VMs
```

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

```
Capacity = 2 sockets
```

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

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

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

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

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

License subscriptions

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

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

Register RecoverPoint by email or phone

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

About this task

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

Steps

1. Gather the required information.
   - Download the post-deployment form:

     a. Access https://www.dell.com/support
     b. Search for the term *Post-Deployment Form*

   - If you have access to a Flex plugin, fill out the RecoverPoint and RecoverPoint for VMs Post-Deployment Form, for every vRPA cluster
   - Export the RecoverPoint registration information, for every vRPA cluster. Using the Flex plugin:

     a. Select **Administration** > **vRPA Clusters**.
     b. Select the vRPA cluster for which you want to export a post-deployment form, and then click **Support**.

  c. In the Registration pane, click the **Export to CSV** button and save the file to the computer.

2. Send the information to the Install Base group:
   - Customers and partners: Email the post-deployment form to the Install Base group at rp.registration@emc.com.
   - Employees:
     - (Preferred) Use the Install Base Group under Post Sales at http://emc.force.com/BusinessServices.
     - Call in the information to the Install Base Group at 1-866-436-2411 – Monday to Friday (normal Eastern Time Zone working hours).

Creating VMkernel ports

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

Steps

1. Select **System** > **ESX Clusters**, and click the **Configure VMkernal ports** icon of an ESXi cluster.

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

3. Click **DONE**

Load balancing

About this task

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

Perform load balancing:

- When a new consistency group is added to the system. Wait 1 week after the new group is added to accumulate enough traffic history before performing load balancing.
- When a new vRPA is added to a vRPA cluster. Perform load balancing immediately after the vRPA is added.
- If the system enters high load frequently. When load balancing is required, the event logs display a message indicating so. When you see this message, perform load balancing.
- Periodically, to ensure that the system is always handling distributing loads evenly. A script can be created to periodically perform load balancing.

Steps

1. To balance the load on the vRPAs, use an `ssh` client to connect to the vRPA management IP address, and type the RecoverPoint `username` and `password` to log in to the CLI.
2. Run the `balance_load` command to balance the load. To view command parameters that can refine the search, run: `balance_load ?`

Copy VM network configuration guidelines

Use the following guidelines for

Table 1. Virtual machine network settings available through the GUI

Setting	Description	Guidelines
(VM) Operating System	The guest operating system of the specified VM.	- Not customizable. - Automatically populated by the system. - Possible values are *Windows*, *Linux*, or *Unknown*.
(VM) Host Name	The hostname of the specified VM.	- Only mandatory for virtual machines with a Linux operating system. - Customizable. - Value can be retrieved from production VM using **Get Value from Production** or **Get All Values from Production**.
(VM) DNS Domain	The DNS domain for the specified VM.	- Only relevant (and mandatory) for virtual machines with a Linux operating system. - Value should be in the format `example.company.com`.
(VM) DNS Server(s)	The global IP address that identifies one or more DNS servers for all adapters of the specified VM.	- Only relevant for virtual machines with a Linux operating system. - Customizable. - Can be left blank. - This setting applies to all virtual network adapters of the specified VM. - Separate multiple values with a semicolon (;). - Value can be retrieved from production VM using **Get Value from Production** or **Get All Values from Production**.
(VM) DNS Suffix(s)	The global settings of the suffixes for the DNS servers of all adapters on both Windows and Linux virtual machines.	- Customizable. - Can be left blank. - Value can be retrieved from production VM using **Get Value from Production** or **Get All Values from Production**.
(Adapter) IP Address	IPv4 address for this virtual network adapter.	- Can contain either a static IPv4 address or DHCP string. - Can be left blank when using IPv6. - Define one IPv4 address, one IPv6 address, or one of each, for the same virtual network adapter. Entering multiple IPv4 or IPv6 addresses for the same virtual network adapter is not supported. - Value can be retrieved from production VM using **Get Value from Production** or **Get All Values from Production**.
(Adapter) Subnet	IPv4 subnet mask for this virtual network adapter.	- Mandatory when an **IP Address** is entered. - Can be left blank when using IPv6. - Value can be retrieved from production VM using **Get Value from Production** or **Get All Values from Production**.
(Adapter) Gateway(s)	One or more IPv4 gateways for this virtual network adapter.	- Mandatory when an **IP Address** is entered. - Can be left blank when using IPv6. - Separate multiple values with a semicolon (;). - Value can be retrieved from production VM using **Get Value from Production** or **Get All Values from Production**.

Table 1. Virtual machine network settings available through the GUI (continued)

Setting	Description	Guidelines
(Adapter) IPv6 Address	IPv6 address for this virtual network adapter.	• Can contain either a static IPv6 address or it's DHCP string. • Can be left blank when using IPv4. • Define one IPv4 address, one IPv6 address, or one of each, for the same virtual network adapter. Entering multiple IPv4 or IPv6 addresses for the same virtual network adapter is not supported. • Value can be retrieved from production VM using **Get Value from Production** or **Get All Values from Production**.
(Adapter) IPv6 Subnet Prefix Length	IPv6 subnet mask for this virtual network adapter.	• Customizable. • Can be left blank when using IPv4. • Value can be retrieved from production VM using **Get Value from Production** or **Get All Values from Production**.
(Adapter) IPv6 Gateway(s)	One or more IPv6 gateways for this virtual network adapter.	• Customizable. • Mandatory when an IPv6 format **IP Address** is entered. • Can be left blank when using IPv4. • Separate multiple values with a semicolon (;). • Value can be retrieved from production VM using **Get Value from Production** or **Get All Values from Production**.
(Adapter) DNS Server(s)	IP address of one or more DNS server(s) for this virtual network adapter.	• Can be left blank. • Can contain one or more IPv4 DNS servers for each virtual network adapter (NIC). • Applies only to the configured adapter when a value other than **Adapter ID 0** is defined. • Separate multiple values with a semicolon (;). • Value can be retrieved from production VM using **Get Value from Production** or **Get All Values from Production**.
(Adapter) NetBIOS	Whether or not to activate NetBIOS on this virtual network adapter.	• Cannot be left blank. • Only relevant for virtual machines running a Windows operating system. • Default is **Enabled**. • Net BIOS should be enabled. • Valid values are **DISABLED**, **ENABLED**, **ENABLED_VIA_DHCP**.
(Adapter) Primary WINS	Primary WINS server of this virtual network adapter.	• Relevant for windows virtual machines only. • Customizable. • Can be left blank.
(Adapter) Secondary WINS	Secondary WINS server of this virtual network adapter.	• Relevant for windows virtual machines only. • Customizable. • Can be left blank.

Table 2. Network settings only available through the JSON file

Setting	Description	Guidelines
CG ID	The consistency group ID in the RecoverPoint for VMs system.	• Do not modify this field. • Automatically populated by the system. • Not customizable. • Can be left blank.

Table 2. Network settings only available through the JSON file (continued)

Setting	Description	Guidelines
CG Name	Name of the consistency group in the RecoverPoint for VMs system.	● Automatically populated by the system. ● Must be the name associated with the specified consistency ID in RecoverPoint for VMs. ● Customizable. ● Can be left blank.
VC ID	The vCenter Server ID in VMware.	● Do not modify this field. ● Automatically populated by the system. ● Not customizable. ● Can be left blank.
VC Name	The name of the vCenter Server hosting the virtual machine.	● Customizable. ● Can be left blank.
VM ID	The virtual machine ID that vCenter Server uses.	● Do not modify this field. ● Automatically populated by the system. ● Not customizable. ● Cannot be left blank.
VM Name	The name of the virtual machine.	● Customizable. ● Automatically populated by the system. ● Can be left blank.
NIC Index in vCenter	The index of the adapter in the order of virtual network adapters (NICs) in the virtual machine settings of the vCenter web client.	● Customizable. ● Cannot be left blank. ● Enter a numeric value. ● Enter a value of **0** to define the first virtual network adapter in the vSphere Web Client. Enter a value of **1** to define the next network adapter.

Changing the network adapter configuration of a protected VM

About this task

When the virtual network adapter (NIC) configuration of a production VM changes, any pre-existing copy VM network configuration may be adversely affected and may require re-configuration before it works. After adding or removing NICs from a protected virtual machine, re-configure the copy VM network using Re-IP rules on page 47.

If the NIC configuration of a production VM changes and the change is not reflected in the copy VM, ensure **Hardware changes** is enabled and *enable image access* by #unique_92.

Shared disks

Shared disks are VMDKs or RDMs that are mapped to and accessible from multiple VMs. They are commonly used with database systems such as Oracle RAC and Microsoft SQL Server, where shared disks hold the database and all nodes in the cluster access it.

Use the procedures in this appendix to manage systems that use shared VMDKs or RDMs.

Topics:

Add VM to a host cluster

Use this procedure to add a VM to a Microsoft Failover Cluster or Oracle RAC whose VMs are already protected by RecoverPoint for VMs.

Steps

1. Add all shared disks to the new production VM.
2. Create a VM with the same configuration at the copy.
3. Add all shared disks at the copy to the new copy VM.
4. Using the Flex plugin, add the new VM to the existing consistency group. You must manually select the copy VM that you created in Step 2.

   (i) **NOTE:** You can manually select the copy VM only when using the Flex plugin.

5. Add the newly protected VM to the Microsoft Failover Cluster or Oracle RAC.

Results

The newly added VM is now protected.

Add shared disk to VMs that belong to a consistency group

Use this procedure to add a shared disk to all VMs that belong to an existing consistency group.

Steps

1. Add a shared disk (VMDK or RDM) to the production VMs that belong to the host Microsoft Failover Cluster or Oracle RAC.
2. Add a shared disk with the same configuration to the parallel VMs at the copy.
3. Use the RecoverPoint for VMs plugin to protect the new shared disk by including the shared disk on all of the production VMs.
4. Add the shared disk to the Microsoft Failover Cluster or Oracle RAC.

Results

The new shared disk is added to all of the VMs that belong to the consistency group.

Recover production after deletion of production VM or shared disk

Use this procedure to restore a shared disk (VMDK or RDM) that has been removed from production VMs.

About this task

RecoverPoint for VMs automatically recreates removed VMs, but does not recreate removed shared disks.

Steps

1. If the shared disk is removed from all of the VMs that were mapped to it, then recreate the shared disk.
   If one or more of the VMs remain, then the shared disk is not deleted. Do not create a new shared disk.
2. Recover production to a selected image.
   Recover production will not proceed if you have not recreated the removed shared disk.
3. After production recovery, to resume replication from production to copy, remap the shared disk to all of the VMs to which it was previously mapped.

Results

The shared disk is restored, and replication can resume.

Failover after deletion of production VM or shared disk

Use this procedure for failover of a consistency group to a copy following removal of a production VM or shared disk (VMDK or RDM).

About this task

RecoverPoint for VMs automatically recreates removed VMs, but does not recreate removed shared disks.

Steps

1. Failover the consistency group to a copy.
   Even when production VMs or shared disks are removed, it is still possible to successfully failover the consistency group.
2. After failover, to resume replication from the copy to production, restore the removed VMs and shared disks at production.
   If one or more of the VMs remain, then the shared disk is not deleted. Do not create a new shared disk. Ensure that the shared disk is added to all of the VMs that were previously mapped to it.
   If the shared disk at production is removed from all of the VMs that were mapped to it, then recreate the shared disk, and add it to all of the VMs that were previously mapped to it.

Results

Failover to the copy is successful, and data replication can resume from the copy to production.

Remove a VM from its host cluster

Use this procedure to unprotect a VM that belongs to a Microsoft Failover Cluster or Oracle RAC.

Steps

1. Remove a VM from its host cluster.
2. Unmap all shared disks from the VM.
3. Remove the VM from its consistency group.

Results

The VM is removed from its host cluster.

Unprotect a shared disk

Use this procedure to unprotect a shared disk (VMDK or RDM).

Steps

1. Use the RecoverPoint for VMs plugin to exclude from RecoverPoint for VMs protection all VMs that are mapped to the shared disk .
2. For any shared disk that is removed from the production VMs, remove the parallel shared disks from the copy VMs. Snapshots taken during the time when the shared volume is partially protected are inconsistent.

Results

The shared disk is removed from protection.