RecoverPoint for VMs

Installation and Deployment Guide

5.3

D≪LLEMC

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

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://support.dell.com) to ensure that you are using the latest version of this document.

Purpose

This document describes how to install and configure a Recoverpoint for Virtual Machines system.

Audience

This document is intended for Virtualization Administrators who manage, maintain and scale their virtual environments, and Application Administrators who monitor application performance.

Related documentation

The following publications provide additional information:

- *RecoverPoint for VMs Release Notes*
- *RecoverPoint for VMs Quick Start Installation Poster*
- *RecoverPoint for VMs Installation and Deployment Guide*
- *RecoverPoint for VMs Product Guide*
- *RecoverPoint for VMs HTML5 Plugin Administrator's Guide*
- *RecoverPoint for VMs Flex Plugin Administrator's Guide*
- *RecoverPoint for VMs CLI Reference Guide*
- *RecoverPoint for VMs Deployment REST API Programming Guide*
- *RecoverPoint for VMs Functional REST API Programmer's Guide*
- *RecoverPoint for VMs Security Configuration Guide*
- *RecoverPoint for VMs Scale and Performance Guide*
- *RecoverPoint for VMs Cloud Solutions 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.

Internationalization and localization

The RecoverPoint for VMs 5.3 HTML5 plugin is available only in English. The vSphere Web Client Flex plugin is also available only in English.

The following documents are available in English, Simplified Chinese, Japanese, Korean, German, French, Latin-American Spanish, and Brazilian Portuguese:

- *RecoverPoint for VMs 5.3 Installation and Deployment Guide*
- *RecoverPoint for VMs 5.3 Installation Poster*
- *RecoverPoint for VMs 5.3 Release Notes* **New features** section.

Introduction to RecoverPoint for VMs

RecoverPoint for VMs is a virtualized solution that provides data replication, protection, and recovery within the VMware vSphere environment.

Definition of key terms and a system diagram help you to understand the system operation.

Topics:

*   RecoverPoint for VMs system

RecoverPoint for VMs system

Key components of the RecoverPoint for VMs system are defined and illustrated.

Key system components that are involved in this installation include:

vRPA	The virtual RecoverPoint Appliance is a data appliance that manages data replication. Create the vRPAs you need by using the vSphere Client or vSphere Web Client from the vCenter Server.
vRPA cluster	A group of up to 8 vRPAs that work together to replicate and protect data. You will create the vRPA clusters and connect them to the system by using the RecoverPoint for VMs Deployer wizards.
RecoverPoint for VMs plugin	The vSphere Client (HTML5) or vSphere Web Client (Flex) user interface for managing VM replication. The Flex plugin is automatically installed after you create the vRPA cluster. The HTML5 plugin is installed after the plugin server is deployed and configured.
RecoverPoint for VMs plugin server	A dedicated VM deployed from an OVA template that provides replication management over one or more RecoverPoint for VMs systems that are hosted on the same vCenter Server. The plugin server provides a single endpoint for the vSphere HTML5 plugin (when running vSphere 6.7 U1 or later) and the REAPER API.
RecoverPoint for VMs splitter	Proprietary software installed on every ESXi host in an ESX cluster that is involved in RecoverPoint replication or running virtual RPAs. It intercepts every write for a protected VM and sends a copy of the write to the assigned vRPA and then on to the designated storage volume(s). The splitter is automatically installed on every ESXi hosts in the ESX cluster where the vRPA cluster resides as part of the vRPA cluster deployment, or is automatically installed on ESXi hosts as part of any additional ESX clusters that are registered to the vRPA cluster.
RecoverPoint for VMs system	One or more connected vRPA clusters.

RecoverPoint for VMs system on page 11 provides a reference diagram that shows the vRPAs and vRPA clusters within the RecoverPoint for VMs system, including the RecoverPoint for VMs Plugin Servers. The diagram shows how these components interconnect within the VMware vSphere environment.
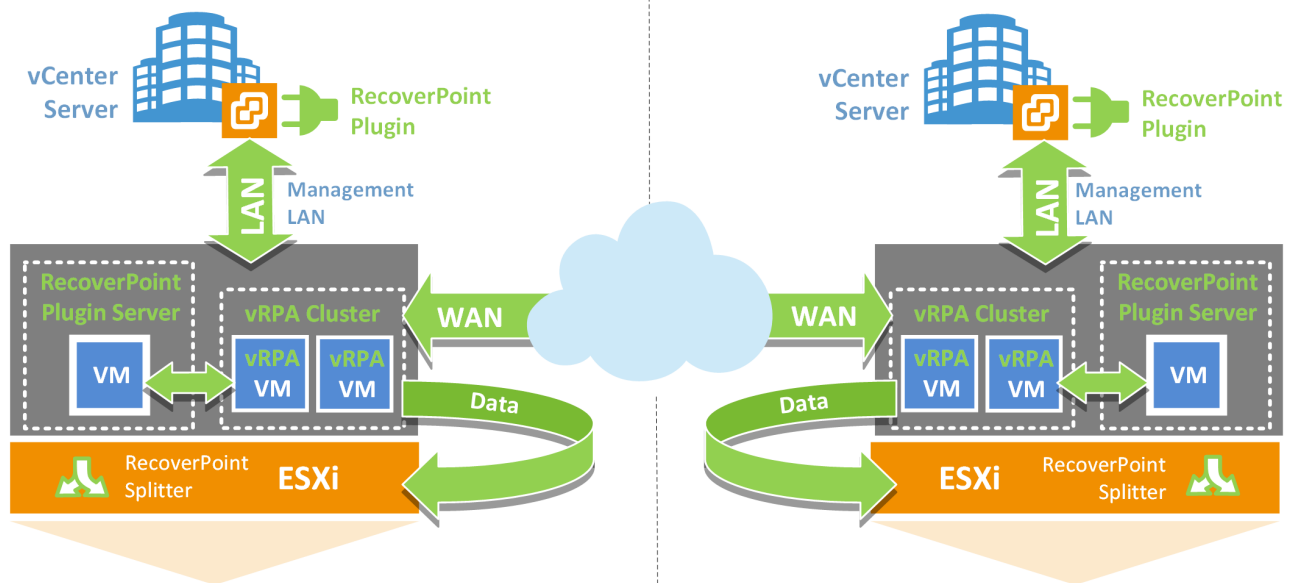
Figure 1. RecoverPoint for VMs system

Preparing to install RecoverPoint for VMs

Guidelines help you choose the number of vRPAs and vRPA clusters, vRPA performance profile, and network adapter topology.

Preparing the VMware network and determining storage capacity sets the stage for a successful installation.

Topics:

RecoverPoint for VMs networking example

A reference diagram is a valuable tool for planning your RecoverPoint for VMs system. The diagram shows an example of the network that interconnects key system components.

For clarity, Networking example on page 13 shows the components and interconnections of only one site in a small system. The IP addresses are for illustration purposes only.

The remaining sections of this chapter will help you to plan a system that meets your specific requirements.
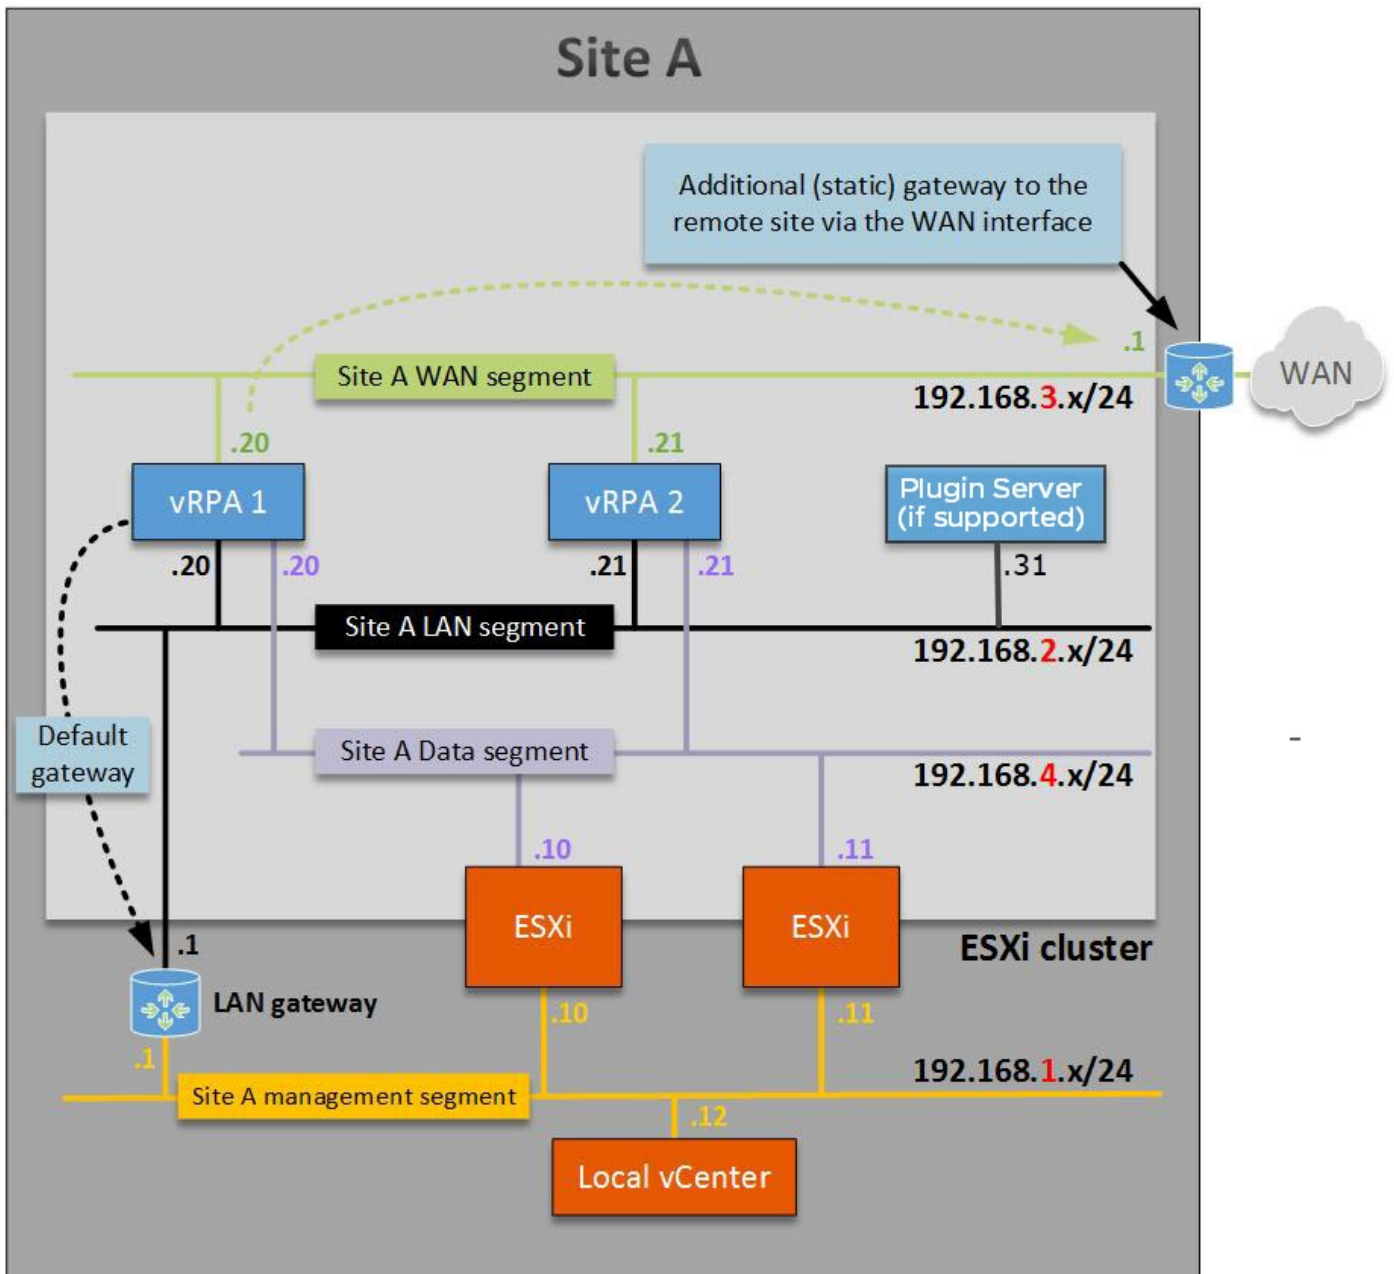
Figure 2. Networking example

Planning your system

System limitations

Understanding system limitations facilitates the installation of the RecoverPoint for VMs system.

Successful operation of RecoverPoint for VMs depends on a persistent vSphere deployment.

For a comprehensive and up-to-date list of system limitations, see the *RecoverPoint for Virtual Machines Release Notes*.

Allocating IP addresses

Knowing how many IP addresses you need for the RecoverPoint for VMs system helps you to allocate the IP addresses before the installation is scheduled.

The RecoverPoint for VMs system requires these IP addresses:

- Cluster management IP address for each vRPA cluster
- An IP address for each vRPA network adapter (see Choosing a network adapter topology on page 15)
- An IP address for each VMkernel port
- An IP address for each plugin server
- An IP address for the NTP server (recommended)
- An IP address for the primary and secondary DNS servers (optional)

To allocate the necessary IP addresses for the RecoverPoint for VMs system, consult with the network administrator.

Document these addresses in an installation data form or spreadsheet before you begin the installation.

Documenting the installation settings

Creating an inventory of the RecoverPoint for VMs system ensures that you have all the required settings before the installation begins.

As you perform the required planning, create an installation data form or spreadsheet to record the values that you type during the installation. See Installation data forms on page 56 for examples.

Adhere to a consistent naming and numbering convention for the components of the RecoverPoint for VMs system. For example, use the following format for each vRPA:

*<vRPA_cluster_name>*_vRPA-*<n>*, where n=1, 2, ... 8, and where there are between 1 and 5 uniquely named clusters.

Choosing a vRPA topology

The first step in planning the RecoverPoint for VMs system is to determine how many vRPAs you need in each vRPA cluster and how many vRPA clusters you need in the system.

How many vRPAs?

Determining the number of vRPAs in the system is based on protection requirements, existing storage capacity, VMware infrastructure, and replication requirements such as high availability or product evaluation.

For typical installations, two vRPAs per vRPA cluster is sufficient. Two vRPAs per vRPA cluster provide the high availability that most production environments require.

For production environments that do not require high availability or for product evaluation in non-production environments, a single vRPA per cluster is also possible.

To scale up and support higher throughput, you may non-disruptively add vRPAs (up to 8) to each vRPA cluster.

All vRPA clusters in a system must have the same number of vRPAs.

The actual number of vRPAs that you need for each vRPA cluster depends on the capabilities of your storage, network, ESXi hosts, and the scale and performance requirements of your system.

For specific details and examples, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide*.

How many vRPA clusters?

The number of vRPA clusters you need is based on protection requirements, or on whether you require local or remote replication, or both.

For most installations, you will install two vRPA clusters in your RecoverPoint for VMs system.

For local replication, you need only one vRPA cluster. To support remote replication, two vRPA clusters are required. The maximum number of vRPA clusters in a system is five.

A vRPA cluster is confined to a single ESXi cluster. All vRPAs in a vRPA cluster must be in the same ESXi cluster, although an ESXi cluster can support multiple vRPA clusters.

A vRPA cluster protects VMs on the same or a different ESXi cluster. This capability requires connections between the vRPA cluster and the ESXi hosts (see Preparing the network on page 16).

For specific details and examples, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide.*

Choosing a vRPA performance profile

The vRPA performance profile defines the number of virtual CPUs, RAM, and VMDK capacity allocated to each vRPA. You choose a performance profile depending on the number of protected VMs and expected throughput.

For most installations, 2 CPUs and 8 GB RAM is sufficient.

The actual vRPA performance profile that you need depends on these factors:

- IOPS and throughput of protected VMs
- The number of VMs protected by the vRPA cluster

You can change the resource allocation later by using the vSphere Client or vSphere Web Client.

Decide which of these vRPA performance profiles you need:

- 2x virtual CPUs / 8 GB RAM
- 4x virtual CPUs / 8 GB RAM
- 8x virtual CPUs / 8 GB RAM

The vRPAs for all of these performance profiles support at least 256 VMs. For all profiles, each vRPA VM has a 35 GB VMDK capacity.

This selection is made when you create vRPAs from the OVF wizard in the vSphere Client or Web Client.

ⓘ **NOTE:** By default, all RAM is reserved and vCPU reservation is set to 3400MHz.

If required, you can add memory and CPU resources after initial OVA deployment. For each vRPA, power off the vRPA, select Edit Settings for the vRPA VM, and add the needed resources.

For details and examples, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide*.

I/O throttling

I/O throttling is used to slow down storage reads that are part of any initialization process including the initial full synchronization (full-sweep).

I/O throttling mitigates the negative impact of initialization on production performance. As a result, however, the initialization process may take longer than expected.

I/O throttling is initially enabled at 400 MBps.

You can use the `config_io_throttling` Sysmgmt CLI command to set the I/O throttling setting.

For more information on I/O throttling, see the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines CLI Reference Guide*.

Choosing a network adapter topology

RecoverPoint for VMs supports LAN, WAN, and data interfaces distributed across multiple network adapters or combined into one. The choice depends on the requirements for high availability and performance.

Combining multiple interfaces on one network adapter is recommended for small environments. The advantage is a smaller network footprint and ease of installation and management.

Where high availability and performance are desired, you should separate the LAN and WAN interfaces from the data interfaces (recommended for most installations). For even better performance, place each network on a separate virtual switch.

When using multiple network adapters, each adapter should be assigned an IP on a different subnet. Assigning adapter IPs on the same subnet can cause connectivity issues that result in packets being sent through the wrong interface. If necessary, combine multiple adapters according to your preferred network topology.

Decide which of these network adapter topologies you need:

One network adapter	WAN + LAN + Data combined Fewer IP addresses to create and manage Not for high availability solutions
Two network adapters (the default and recommended configuration)	WAN + LAN combined, Data separated Better performance, high availability
Two network adapters	LAN + Data combined, WAN separated Better performance, high availability DHCP for LAN is not supported
Three network adapters	WAN, LAN, and Data separated Better performance, high availability DHCP for LAN is not supported
Four network adapters	WAN and LAN separated, Data separated on two dedicated network adapters Compatible with previous releases Best performance and high availability. Use different subnet masks for the two Data IP addresses. DHCP for LAN is not supported

ⓘ **NOTE:** IPv6 is supported on vRPA LAN and WAN interfaces, but not on vRPA Data interfaces.

This selection is made when you run the Install a vRPA cluster wizard in the RecoverPoint for VMs Deployer.

For high-availability deployments in which clients have redundant physical switches, route each data card to a different virtual switch with a separate network adapter.

For each network adapter, you have the option to assign a dynamic or static IP address.

When using Dynamic Host Configuration Protocol (DHCP):

- Separating WAN and LAN interfaces on different network adapters is supported only when using static IP addresses for the LAN interface
- Redundant, highly available DHCP servers in the network ensure that when a vRPA restarts, it acquires an IP address

Preparing the VMware environment

Supported vSphere versions

For the most up-to-date information on supported VMware vCenter and vSphere versions, refer to the *Simple Support Matrix* available online at https://www.dell.com/support.

Preparing the network

The RecoverPoint for VMs splitter communicates with the vRPAs through a VMKernel port. Setting up separate VMkernel ports is the best practice for isolating splitter traffic from other network traffic. You isolate the traffic by placing the vRPA data interface and a dedicated VMkernel port on a private (separate) subnet. Avoid using the same subnet that is used also for high availability (vMotion) and hosts (applications).

Depending on your existing network, you may not need to configure any additional VMkernel ports, or if so, you can do so later from the RecoverPoint for VMs plug-in, even after you have protected VMs and are ready to begin replication. RecoverPoint for VMs informs you after protection if there are possible communication issues between vRPAs and splitters. When needed, RecoverPoint for VMs assists in automatically creating VMkernel ports for all the ESXi hosts in the ESX cluster. The procedure is described in the *RecoverPoint for Virtual Machines HTML Plugin Administrator's Guide* and the *RecoverPoint for Virtual Machines Flex Plugin Administrator's Guide.*

Alternatively, you can configure VMkernel ports manually by following the procedure in

The number of VMkernel ports you need is based on the network adapter topology you previously selected. If you decided to use four network adapters for the topology, create two VMkernel ports. Otherwise, one VMkernel port is required.

Establishing vCenter-to-vRPA communication

During installation, the vCenter server communicates with the vRPAs over port 443 to acquire the RecoverPoint for VMs plug-in. The ESXi clusters communicate over the network with the vRPA targets.

Steps

- Ensure that you open port 443 between the vCenter Server and the vRPAs in both directions, and port 7225 in the vCenter Server-to-vRPAs direction only.
- Ensure that ESXi clusters can communicate with their vRPA targets. Configure the ESXi firewall profile to allow communication through the network.
- See the *RecoverPoint for Virtual Machines Security Configuration Guide* for more information.

Plugin server communication

- Ensure that port 443 is open for bi-directional communication between the plugin server and the vCenter Server.
- Ensure you open port 9443 for plugin server-to-vRPA communication.

See the *RecoverPoint for Virtual Machines Security Configuration Guide* for more information.

Preparing the storage

Determining the amount and types of storage you need requires careful planning, guidelines, and sizing tools.

RecoverPoint for VMs replicates VMs on any type of storage that VMware supports including VMFS, NFS, vSAN, and vVols.

Ensure that all ESXi hosts in the cluster where the vRPAs reside share the datastore for the repository VMDK.

RecoverPoint for VMs requires additional storage for journal VMDKs to store point-in-time history. This storage is needed at local and remote sites. The amount of journal storage you need depends on site-specific installation and replication requirements and requires careful planning. A general guideline is to begin with a number that is 15–25% of the total protected VM capacity. If required, you may add additional storage later. To size the system according to estimated workloads, use the RecoverPoint Sizer tool. See https://psapps.emc.com/recoverpoint-sizer-basic.

The total storage capacity that is required includes:

- Storage for production VMs at the production site
- Storage for replica VMs at the replica site
- Storage for journal VMDKs
- 35 GB for each vRPA in the RecoverPoint for VMs system

A persistent scratch location on the ESXi host is required for storing splitter configuration information. The scratch location (`/scratch/log`) requires at least 500 MB of free storage space on a permanently available persistent storage device.

ⓘ **NOTE:** Each ESXi host should have its own dedicated datastore for the scratch directory.

For more details and examples, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide*.

For additional guidelines and sizing tools, contact Customer Support.

ⓘ **NOTE:** Shared virtual disks (VMDK/RDM) are not supported.

Secure boot considerations

RecoverPoint for VMs version 5.3 SP2 supports secure boot. If secure boot is enabled on your ESXi hosts ensure that you install RecoverPoint for VMs 5.3 SP2 or later.

Understanding the installation flow

Understanding the stages of the installation work flow helps you to successfully install the RecoverPoint for VMs system, so that your system is ready for protecting VMs.

Stages of the installation flow on page 18 shows the major stages of the installation flow. Procedures in the installation flow on page 18 provides details of the required procedures for each stage of the installation flow.
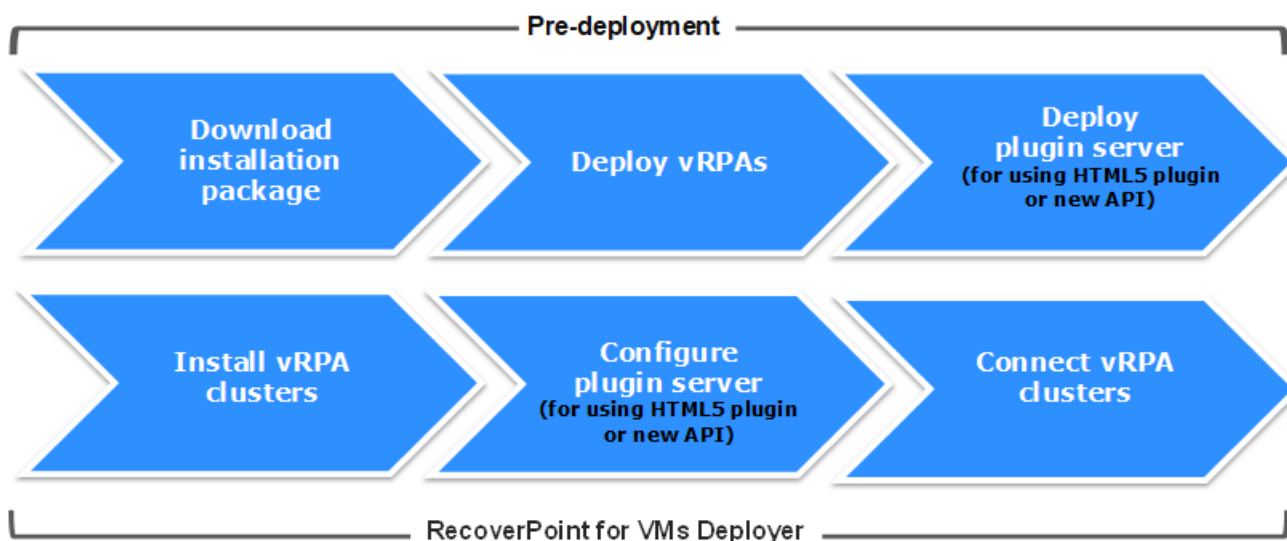
Figure 3. Stages of the installation flow

Table 1. Procedures in the installation flow

Stage of installation flow	Sequence of procedures in the installation flow	Interface
Download installation package	Download the installation package on page 19	Online support site
Deploy vRPAs	Deploy vRPAs on page 20	vSphere Client, vSphere Web Client
Deploy plugin server	Deploy the plugin server on page 21	vSphere Client
Install vRPA clusters	Install vRPA clusters on page 22	RecoverPoint for VMs Deployer
Configure plugin server	Configure the plugin server on page 23	RecoverPoint for VMs Deployer
Connect vRPA clusters	Connect vRPA clusters on page 24	RecoverPoint for VMs Deployer

Use the RecoverPoint for VMs plugin for the vSphere Client (HTML5) or the vSphere Web Client (Flex) to make final preparations for protecting VMs, including registering and licensing the system, and creating any required VMkernel ports (Completing installation of the RecoverPoint for VMs system on page 25).

Installing the RecoverPoint for VMs system

Installing the RecoverPoint for VMs system involves deploying the vRPAs, installing the vRPA clusters, and connecting the vRPA clusters together. In the RecoverPoint for VMs plugin for vSphere Client (HTML), it also involves installing and configuring a plugin server.

Topics:

Download the installation package

Download the installation software kit and uncompress the `.zip` file.

About this task

ⓘ **NOTE:** Try & Buy customers can download the latest RecoverPoint for VMs trial version from https://www.delltechnologies.com/en-us/software-downloads/index.htm#collapse=0&tab0=1.

Steps

1. Browse to https://www.dell.com/support.
2. Perform a search in the **Type a Product Name** text box for **RecoverPoint for Virtual Machines**.
3. Locate and download **RecoverPoint for Virtual Machines *<version>* Installation Kit**.
   Example of downloaded file:
   `RecoverPoint_for_Virtual_Machines_<version>_Installation_Kit_<md5_checksum>.zip`
4. Uncompress the `.zip` file.

   The `.zip` file contains the OVA and ISO files that are needed for the installation.

   In RecoverPoint for VMs 5.3.1 and later, the OVA and ISO files are protected by SRO and GPG signatures, respectively. If you want to verify a signature, see Verifying signatures on page 19
5. (Recommended) Obtain documentation for RecoverPoint for VMs.
6. Continue to Deploy vRPAs on page 20.

Verifying signatures

Use the relevant procedure to verify a signature that is attached to an OVA or ISO file that you downloaded from the RecoverPoint for VMs downloads page.

Verifying the SRO signature of an OVA file

During deployment, the OVA deployment wizard automatically verifies the signature that was applied to the OVF package using the certification file provided by SRO.

You can verify the signature also manually in either of the following ways:

- Inspect the file, and compare with the signature.

- Run: **ovftool --machineOutput file.ova | grep -A40 CertValidate**

Verifying the GPG signature of an ISO file

Follow this procedure to verify the GPG signature on your ISO file:

1. To import the Public Key, run:

   bash-4.4# gpg --import pubkey.gpg

   GPG acknowledges successful import of the Public Key.

2. To verify the signature, run:

   bash-4.4# gpg --verify rp.iso.sig rp.iso

   GPG confirms that the signature is good.

Deploy vRPAs

Deploy a standard OVA to create vRPAs for RecoverPoint for VMs.

Prerequisites

Ensure that you have completed:

- Preparations for installation.
- Installation data form or spreadsheet to facilitate entering requested information (recommended). See Installation data forms on page 56.

Steps

1. In the vSphere Client (HTML5) or Web Client (Flex), right-click an ESX cluster, and select **Deploy OVF Template...**.
2. In the **Select an OVF template** screen, either enter a URL from which to download the OVF package from the Internet, or choose a location from which to access the file locally.
3. In the **Select name and folder** screen, type a name for this vRPA and select a folder or data center.

   If you type the name of an existing vRPA, you are not permitted to continue.
4. In the **Select resource** screen, specify the vRPA OVF package location.
5. In the **Review details** screen, review the general properties of the OVF template. To accept, click **Next**.
6. In the **Accept License Agreements** screen, if you accept the terms of the End-User License Agreement, click **Accept** and **Next**.
7. In the **Select configuration** screen, select the desired vRPA performance profile.
8. If prompted to select a resource, in the **Select a resource** screen, select a cluster, host, or resource pool.
9. In the **Select storage** screen, select a disk format, storage policy, and high-performance datastore (best practice) to host the vRPA virtual machine files.

   All ESXi hosts in the cluster where the vRPAs reside must share the datastore where the repository VMDK resides.

   Do not deploy the vRPA on a local datastore.
10. In the **Setup networks** screen, select a destination network for the **RecoverPoint Management Network**, and select an IP protocol.
11. In the **Customize template** screen, type these vRPA LAN settings: IP address, subnet mask, and gateway.

    Follow instructions on the screen for using DHCP or static IP addresses depending on the network adapter topology.
12. The **Ready to Complete** screen summarizes all the selections. Select **Power on after deployment**. To create the vRPA, click **Finish**.
    The **Deploying vRPA** screen appears, showing the progress.
13. Power on the vRPA VM.
14. To create additional vRPAs, repeat this procedure.

Results

When a vRPA is created, the **vRPA Summary** tab shows the vRPA package contents as specified. The selected IP policy is implemented automatically when the vRPA is powered on.

Next steps

To enable redundancy in case an ESXi host or datastore fails, ensure that vRPAs do not share the same ESXi host or datastore.

When you finish creating vRPAs:

● If you are running vSphere 6.7 Update 1, or later, continue to "Deploy the plugin server".
● Otherwise, continue to "Install vRPA clusters".

Deploy the plugin server

Deploy an OVA to create a RecoverPoint for VMs plugin server.

Prerequisites

Ensure that you:

● Have completed preparations for installation.
● Are running vSphere 6.7 Update 1 or later.
   (i) **NOTE:** In systems running vSphere versions earlier than 6.7 U1, you can use the plugin server to leverage use of the REAPER API (even though it does not support use of the RecoverPoint for VMs plugin for the vSphere Client).

Steps

1. Navigate to vCenter.
   You are prompted to launch **VSPHERE CLIENT (HTML5)** or **VSPHERE WEB CLIENT (FLEX)**
2. Click **VSPHERE CLIENT (HTML5)** to launch the HTML5-based vSphere client.
3. In the vSphere Client, right-click the ESX cluster and select **Deploy OVF Template**.
   The Deploy OVF Template screen opens.
4. In the **Select an OVF template** screen, select an OVF template for the RecoverPoint for VMs vSphere plugin server from a remote URL or local file system.
5. In the **Select a name and folder** screen, specify a unique name for the plugin server, and select the location of the virtual machine on which you want to deploy your plugin server.
6. In the **Select a compute resource** screen, select the destination compute resource for this operation.
7. In the **Review details** screen, review the general properties of the OVF template. To accept, click **Next**.
8. In the **Select storage** screen, select the storage for the configuration and disk files including the virtual disk format and VM storage policy.
9. In the **Select networks** screen, select a destination network for each source network.
10. In the **Customize template** screen, customize the deployment properties of this software solution.
    a. Select the time zone.
    b. The network management default is DHCP. If you are managing your network using static IP configuration, values for the gateway, netmask, and IP address are mandatory. Values for DNS and and FQDN are optional, but entering a DNS address is recommended.
11. In the **Ready to complete** screen, review the details of the deployment and, when satisfied, click **FINISH** to create the plugin server.
    To track that the plugin server creation is successful, check for an OVF deployment in the Recent Tasks window of the vSphere Client.
12. Power on the plugin server VM.

Next steps

After installing a vRPA cluster, you should configure this plugin server with the vCenter on which the vRPA cluster resides.

To create a plugin server for another vCenter, either at this site or a different site, repeat this procedure.

Install vRPA clusters

Follow the **Install a vRPA cluster** wizard to create one or more vRPA clusters for RecoverPoint for VMs.

Prerequisites

- If you have changed the default plugin server certificate and have not yet configured the plugin server, perform Changing the plugin server certificate on page 52.
- If you have changed the default plugin server certificate and have already configured the plugin server, perform Changing a registered plugin server certificate on page 53.
- Get the Installation data forms on page 56 that you created when planning the system ready so that when you are prompted to type data, you can consult them.

About this task

ⓘ **NOTE:** During vRPA cluster installation, the RecoverPoint for VMs splitters and JAM VIBs are automatically installed on all ESXi hosts that belong to ESXi clusters on which vRPAs are running.

Steps

1. In a web browser, type **https://<LAN-ip-address>** where *<LAN-ip-address>* is the LAN IP address of vRPA 1 or vRPA 2 in the cluster you are installing. In the home page, click **RecoverPoint for VMs Deployer**.

   If you are using DHCP, obtain the LAN IP address from the vSphere Client or vSphere Web Client by selecting the vRPA and clicking the **Summary** tab.

2. If prompted, type the login credentials for the admin user, and click **Sign in**.
   The **RecoverPoint for VMs Deployer** home page appears.

3. Select the **Install a vRPA cluster** wizard.

4. On the **Version Requirements** page, the version requirements file is automatically downloaded and validated to ensure that the system meets the requirements. If you have a `.json` configuration file that you want to import, click the **Settings** icon and then click **Import**.

   If version requirements verification is successful, click **Next** to continue. If issues are found, analyze and fix blocking issues before continuing.

   If the version requirements file fails to download, you are prompted to select one of these options:

   - **Retry downloading the up-to-date requirements from Dell Online Support**
   - **Provide version requirements file manually**
   - **Do not check version requirements**

   ⓘ **NOTE:** To obtain the version requirements file for offline installation, browse to https://rplicense.emc.com/download. This page provides an option to download or email the `a-cca.xml` file. If this option is not available, open a Service Request with Customer Support Services (severity level 3). In the request, ask for the latest version requirements file for the RecoverPoint for VMs Deployer. The file is provided within one (1) business day and must be used within 30 days.

5. On the **Installation Prerequisites** page, type the requested information for the vCenter on which the current vRPA is running, and then click **Connect**.

   If the **SSL Certificate** window appears, verify the vCenter's SSL certificate and click **Confirm**.

6. Review the **Pre-installation Validation Results** area. If validation errors are listed, fix them before proceeding.

   If an error can be automatically fixed, the **Fix** button appears in the **Auto-Fix** column.

7. On the **Environment Settings** page, define the required settings.

   - Type a name for the vRPA cluster.
   - To align with security best practices, replace the default admin user password with a new unique password.
     ⓘ **NOTE:** The admin user (with the administrator role) is authorized with all access permissions for managing your RecoverPoint for VMs system. Use the same admin password for all vRPA clusters in a system, and all vRPA clusters under the same vCenter Server. The password for the admin user serves also as the password for the root user across all vRPAs in the system.
   - Type IP addresses for DNS and NTP servers.

> (i) **NOTE:** If you have a cloud copy, all vRPAs must be able to resolve amazonaws.com addresses, so all vRPA clusters will require an appropriately configured DNS server. See the *RecoverPoint for VMs Cloud Solutions Guide* for more information.

8. On the **vRPA Settings** page:
   a. Select the vRPAs for the vRPA cluster and click the **Apply Selection** button.
   b. Select a repository volume from the list. All ESXi hosts in the cluster where the vRPAs reside must share this volume.

9. On the **Network Settings** page, provide the requested settings for the vRPA cluster and its vRPAs.
   - In the **Network Adapters Configuration** area, keep the default setting or click **Edit** to choose a different network adapter topology.
   - In the **Network Mapping** area, for each network adapter, select a value and whether to use DHCP. Type a Cluster Management IP address.
   - In the **vRPA Settings** area, type the requested IP addresses. If the network configuration requires gateways to communicate with remote vRPA clusters, click **Add** to insert each gateway. For each gateway that you add at the current cluster, add a gateway at the remote cluster.
   - In the **Advanced Settings** area, change the **MTU** values only if required. MTU values must be consistent across the communication interface from source to target. See KB article 484259 for more information.

10. Click **Install** to initiate completion of the vRPA cluster installation.

    You can follow progress of the installation on the **Deployment progress** page.

    During successful installation:
    - The Flex-based RecoverPoint for VMs vSphere plugin is installed automatically. Installation of the plugin usually occurs immediately, but it might take some time for the vCenter to identify the plugin. If you experience issues with the RecoverPoint for VMs plugin, log out and log in again to the vSphere Web Client as described in "Troubleshooting the RecoverPoint for VMs plugin".
    - The HTML5-based RecoverPoint for VMs vSphere plugin is installed after the plugin server is deployed and configured (see Configure the plugin server on page 23).
    - Splitters and Journal Access Modules (JAMs) are pushed to all ESXi hosts in the ESXi cluster where the vRPAs are installed.

    If installation fails:
    - To identify the cause of failure, review the displayed error messages.
    - To return to the step in the wizard where you can fix the problem, click **Back**. Fix the problem and retry the installation.
    - Alternatively, you can retry the operation that failed by clicking **Retry the operation**.
    - If installation continues to fail, contact Customer Support.

Results

If you are running vCenter 6.7 U1, or later, and this is the first vRPA cluster installed on the vCenter, click **Proceed to configure plugin server** now to configure your RecoverPoint for VMs plugin server for this vCenter. If you don't configure the plugin server now, you can configure it later using the **Configure plugin server** option in the RecoverPoint for VMs Deployer home screen.

To enable multi-site replication, create additional vRPA clusters by repeating this procedure for each site.

To export a configuration file of the vRPA cluster settings, click the **Settings** icon (upper right), and then click **Export**. This file provides a record of the vRPA cluster configuration for the major version you have installed. You use it to restore the vRPA cluster settings after an installation failure (requiring the installation to be repeated).

When all vRPA clusters are created, continue to "Connect the vRPA clusters".

Configure the plugin server

You can configure the RecoverPoint for VMs vSphere plugin server from the **Configure plugin server** option on the RecoverPoint for VMs Deployer home screen. Alternatively, after using the Deployer installation wizard to successfully install a vRPA cluster on a vCenter, you can proceed directly to configure the plugin server in the Deployer.

Prerequisites

- Ensure that the admin user credentials are the same for all vRPA clusters on the vCenter that you are registering to the plugin server.

- If you want to use a certificate that has been altered, for example, one that has been signed by your organization's internal certificate authority, see "Changing the plugin server certificate" in the *RecoverPoint for VMs Administrator's Guide*.

About this task

Configure the plugin server for only one vRPA cluster that is hosted by the vCenter Server.

Steps

1. To start configuration of the RecoverPoint for VMs plugin server:
   a. Select **Configure plugin server** on the Deployer home page.
   b. Alternatively, immediately after successfully installing a vRPA cluster, click **Proceed to Configure plugin server**.
2. Enter the IP address for the plugin server.

   (i) **NOTE:** Use an IPv4 address only.

3. Confirm the SSL certificate of the vCenter Server (remote host).
4. Click **Configure**

Results

- The plugin server installs the HTML-based RecoverPoint for VMs plugin on the vCenter Server. Installation of the plugin usually occurs immediately, but it might take some time for the vCenter Server to identify the plugin. It is recommended to log out from the vSphere Client, and then to log in again. To access the plugin, navigate to RecoverPoint for VMs in the vSphere Client menu.
- The plugin server discovers and registers all vRPA clusters that are hosted by the vCenter Server and have the same admin user password, including vRPA clusters that may be added later. As a result, the plugin server knows how to direct API calls to the correct vRPA cluster.

Connect vRPA clusters

To enable replication between any two vRPA clusters, use the **Connect vRPA clusters** wizard to establish a connection between them.

Prerequisites

- Do not exceed the maximum number of five vRPA clusters per system.
- (Recommended) Create Installation data forms on page 56.
- Ensure the remote vRPA cluster is *not*:
  - in maintenance mode.
  - an existing, configured vRPA cluster.
  - previously connected to a vRPA cluster.
- Ensure the remote vRPA cluster does *not*:
  - have protected VMs, consistency groups, or group sets.
  - have user or journal volumes.
  - have a license other than a vCenter license.
- If required, add a gateway for communication between vRPA clusters; add a gateway at each vRPA cluster before connecting the vRPA clusters.
- If you have changed the default plugin server certificate and have not yet configured the plugin server, perform Changing the plugin server certificate on page 52.
- If you have changed the default plugin server certificate and have already configured the plugin server, perform Changing a registered plugin server certificate on page 53.

(i) **NOTE:** A remote vRPA cluster that meets these requirements is called a "clean" cluster.

About this task

In this procedure, the "current" cluster is defined as the vRPA cluster to which the **Connect vRPA clusters** wizard is currently pointed. The "remote" cluster is the vRPA cluster at a remote site. This wizard helps you to connect a remote vRPA cluster to the current vRPA cluster.

Steps

1. In a web browser, type **https://<cluster_management-ip-address>** for the vRPA cluster that you want to connect.

2. In the RecoverPoint for VMs Deployer home page of the current cluster, select the **Connect vRPA clusters** wizard.

3. On the **Environment Settings** page, type the requested information for the remote cluster. It is important to enter the WAN IP of one of the remote vRPAs.

4. In the **Current Cluster Settings** area, review the list of gateways that are configured for this vRPA cluster. If required, add one or more gateways on the current vRPA cluster. Remember that for each additional gateway at the current cluster, you must add a gateway at the remote cluster.

5. On the **Add Cluster Progress** page, the remote cluster is connected to your RecoverPoint for VMs system, and IP communication is enabled between the remote cluster and the current cluster.

   ⓘ **NOTE:** This does not enable communication between the remote cluster and any other clusters in your system. To enable communication between the remote cluster and additional clusters, follow the procedure in Enable communication between vRPA clusters on page 27.

6. Continue to the next section, "Register and license the system".

Completing installation of the RecoverPoint for VMs system

Installation is almost complete. Once you register and license your system, and create any VMkernel ports that are needed, you will be able to begin protecting VMs.

Register and license the system

Use the RecoverPoint for VMs plug-in for the vSphere Client (HTML5) or vSphere Web Client (Flex) to register and license your system. Registration and licensing enables support and provides important product updates to keep your system running optimally.

If you are using the HTML5-based vSphere Client (running vSphere 6.7 U1, or later), adding a license automatically enables support. Refer to the "Before you begin" section of the *RecoverPoint for VMs HTML5 Plugin Administrator's Guide*.

If you are using the Flex-based vSphere Web Client (running any vSphere 6.7 or 6.5 version), refer to the "Getting Started wizard" section of the *RecoverPoint for VMs Flex Plugin Administrator's Guide*, or the "License and register RecoverPoint for VMs" section of the *RecoverPoint for Virtual Machines Cloud Solution Guide*.

Create VMkernel ports

Before protecting VMs, you may need to create VMkernel ports for the ESXi hosts in the cluster.

See "Creating VMware ports" section in the *RecoverPoint for VMs HTML5 Plugin Administrator's Guide*.

ⓘ **NOTE:** Use this procedure even If you are otherwise using the Flex-based vSphere Web Client. There is one slight difference: In the first step of the procedure, select **Administration > vRPA Clusters > ESX Clusters**, before clicking the **Settings** icon for an ESXi cluster.

Protect VMs

The RecoverPoint for VMs system is ready for operation. Use the RecoverPoint for VMs plugin (HTML5-based or Flex-based) to begin protecting VMs. See the *RecoverPoint for Virtual Machines Administrator''s Guide* for instructions on protecting VMs and monitoring the system.

4

Maintaining RecoverPoint for VMs

Maintaining the RecoverPoint for VMs system involves tasks such as collecting logs, modifying vRPA cluster network settings and topology, and adding, removing, or replacing vRPAs.

The topics in this chapter provide procedures for system maintenance.

Topics:

Register ESX clusters

By default, ESX clusters are registered automatically as part of the Protect VM procedure. Alternatively, you can register ESX clusters manually, using the RecoverPoint for VMs plug-in to the vSphere Client or vSphere Web Client.

Steps

1. Access the vSphere Client or Web Client.

   In the vSphere Client, click **System** > **ESX Clusters**.

   In the vSphere Web Client, click **RecoverPoint for VMs Management** > **Administration** > **vRPA Clusters**, and select the **ESXi Clusters** tab.

2. Select the ESX cluster to be registered, and click **Add**. Verify that the connectivity status is OK. If there are connectivity issues with the cluster, click **Troubleshoot**.

Add ESXi host to a cluster

Upon adding an ESXi host to an ESX cluster, the RecoverPoint for VMs splitter VIB is installed automatically on the new ESXi host.

The RecoverPoint for VMs JAM VIB is also installed automatically on the new ESXi host, provided that there are vRPA VMs running on the ESX cluster to which the ESXi host has been added.

Configure VMkernel ports

You can configure the VMKernel ports from the RecoverPoint for VMs plug-in UI, and they will be automatically created for all the ESXi hosts in the ESXi cluster. Alternatively, you can use this procedure to manually configure VMkernel ports.

Prerequisites

An ESXi must be registered (see Register ESX clusters on page 26) before you can configure VMkernel adapters on it.

Steps

1. For each ESXi host, click **Manage** > **Networking** > **VMkernel adapters**.
2. Add the VMkernel adapters.
   - Assign IP addresses that are on a routable subnet or on the same subnet as the vRPA data interfaces.

     It is recommended also that the VMkernel and vRPA data ports be on the same L2 network.
   - For a standard vSwitch, create a VMkernel port with the network label: `RP-VM-Kernel-Port-Group`
   - For a distributed vSwitch:
     - Create a VMkernel port on the relevant port group.
     - On the Ports tab of the Distributed Ports Group page, label each VMkernel port that is to be used for splitter-to-vRPA communication as `RP-VM-Kernel-Port-Key`.

   The vRPA data IP addresses are assigned when deploying the vRPA cluster.

Enable communication between vRPA clusters

Use this procedure to enable communication between pairs of vRPA clusters in your system.

About this task

Use the **Connect vRPA clusters** wizard in the RecoverPoint for VMs Deployer to add a vRPA cluster to your system, and to enable IP communication between that new cluster and one of the existing clusters. For more information, see Connect vRPA clusters on page 24.

Use the following procedure to enable communication between additional pairs of clusters.

Steps

1. Use an SSH client to connect as admin user to a cluster management vRPA of one of the pair of vRPA clusters between which you want to enable communication.
2. From the Boxmgmt CLI **Main Menu**, select **Cluster operations** > **Configure connection types to other clusters in the system** > **Configure cluster connection types**.
3. Select the vRPA cluster with which you want to enable communication.

Results

Bi-directional IP communication is enabled between the designated pair of vRPA clusters.

Repeat this procedure for each pair of vRPA clusters for which you want to enable communication.

Modify vRPA cluster network settings

Use the Modify vRPA cluster network wizard to change network settings.

Prerequisites

To modify the network adapter topology, refer to Modify the network topology on page 28.

Steps

1. In a web browser, type `https://<cluster_management-ip-address>` for the vRPA cluster that you want to modify.
2. In the home page, click **RecoverPoint for VMs Deployer**.
3. If prompted, type the login credentials for the admin user and click **Sign in**.
4. Under **More actions**, select **Modify vRPA cluster network**.
5. Make the desired changes to the **Environment Settings** page. If you have a `.json` configuration file that you want to import, hover over the **Settings** icon and click **Import**.
6. Make the desired modifications changes to the **Network Settings** page.

   Some settings cannot be modified.
7. To apply the changes, click **Modify**. To export a configuration file of the vRPA cluster settings, click the **Settings** icon (upper right), and then click **Export**. This file provides a record of the vRPA cluster configuration.

Change the RPA communication security level

For information about the RPA communication security level, including a procedure for changing it, see the *RecoverPoint for VMs Security Configuration Guide*.

Modify the network topology

Use this procedure to modify the existing network topology.

Steps

1. Pause transfer between the production and copies of the consistency groups for the vRPA cluster that you are modifying.
2. From the vSphere Client or Web Client, add the vNIC on all vRPA VMs. Ensure that the type is VMXNET3.
3. Use an SSH client to log in to the vRPA as the admin user.
    a. Detach the vRPA from the vRPA cluster. From the **Main** menu, select **Cluster operations** > **Detach RPA from cluster**.
    b. From the **Main** menu, select **Setup** > **Modify settings** > **Enter cluster details** > **Network Interface and IPs Configuration**.
    c. Select the network topology that you want to use.
    d. Attach the vRPA back to the cluster. From the **Main** menu, select **Cluster operations** > **Attach RPA to cluster**.
4. Repeat step 3 on page 28 for each vRPA in the vRPA cluster.
5. Start transfer between the production and copies of the consistency groups for the modified vRPA cluster.

Installing the splitter with Boxmgmt CLI

Ensure that all vRPAs reside on ESXi hosts with a splitter installed. Splitters are installed automatically during system installation. You can use this procedure to install splitters manually.

Steps

1. Use an SSH client to connect as admin user to a cluster management vRPA.
2. From the boxmgmt **Main Menu**, select **Setup** > **Advanced options** > **Splitter actions** > **Upgrade Splitter**.
3. Provide the vCenter Server credentials.
4. Let the system provide the vCenter certificate automatically and, if it's correct, approve the certificate.
5. Select one or more of the ESXi cluster(s) in this vCenter on which you want to deploy the splitters.

Install the splitter with the RecoverPoint for VMs VIB installer

During system installation, the RecoverPoint for VMs splitter VIB is installed on ESXi hosts automatically. Use this procedure if you need to install the splitter VIB manually.

Prerequisites

Note the Secure boot considerations on page 18.

About this task

(i) **NOTE:** Do not erase the `/scratch` directory during installation.

Steps

1. To copy the splitter VIB, use an SSH client with secure copy protocol:

   `scp <vib name> <username>@<ESXi host IP>:/scratch`

   For example, run:

   `scp EMC_bootbank_RP-Splitter_RPESX-00.5.3.2.0.0.m.222.000.vib root@10.10.10.10:/scratch`

2. To install the splitter VIB in the ESXi host console:
   - In RecoverPoint for VMs version 5.3 SP2 and later, run the command:

     `esxcli software vib install -v /<vib_full_path>`

   - In RecoverPoint for VMs versions prior to 5.3 SP2, run the command:

     `esxcli software vib install -v /<vib_full_path> --no-sig-check`

   If installation is successful, the following message appears:

   ```
   Installation Result
       Message: Operation finished successfully.
       Reboot Required: false
       VIBs Installed: EMC_bootbank_RP-Splitter_RPESX-<version>.vib
       VIBs Removed:
       VIBs Skipped:
   ```

3. Confirm installation of the splitter VIB in the ESXi host console by using SSH to run the following command:

   `esxcli software vib list`

Results

The RecoverPoint for VMs splitter name is displayed at the top of the list.

Next steps

Repeat this procedure for every ESX on which you want to install the RecoverPoint for VMs splitter VIB.

Install JAM

During new installations of RecoverPoint for VMs (but not upgrades), the RecoverPoint for VMs Journal Access Module (JAM) VIB is installed automatically on all ESXi hosts that belong to ESXi clusters on which vRPAs are running. Use this procedure if you need to install the JAM VIB manually.

Prerequisites

Note the Secure boot considerations on page 18.

About this task

In RecoverPoint for VMs versions prior to 5.3 SP2, there is only one JAM VIB. In RecoverPoint for VMs 5.3 SP2 and later versions, there is one JAM VIB for ESXi version 6.x hosts and another for ESXi version 7.x hosts.

(i) **NOTE:** Do not erase the `/scratch` directory during installation.

Steps

1. To copy the JAM VIB, use an SSH client with secure copy protocol:

   `scp <vib name> <username>@<ESXi host IP>:/scratch`

   For example:
   - On ESXi 6.x hosts, run:

     `scp EMC_bootbank_emcjiraf_5320.m.186-1OEM.650.1.17.4598673.vib root@10.10.10.10:/scratch`

   - On ESXi 7.x hosts, run:

     `scp EMC_bootbank_emcjiraf_5320.m.180-1OEM.700.1.0.15843807.vib root@10.10.10.10:/scratch`

     .

2. To install the JAM VIB in the ESXi host console:
   - In RecoverPoint for VMs version 5.3 SP2 and later, run the command:

     `esxcli software vib install -v /<vib_full_path>`

   - In RecoverPoint for VMs versions prior to 5.3 SP2, run the command:

     `esxcli software vib install -v /<vib_full_path> --no-sig-check`

   If installation is successful, the following message appears:

   ```
   Installation Result
       Message: Operation finished successfully.
       Reboot Required: false
       VIBs Installed: EMC_bootbank_emcjiraf_<RP version><ESX version>.vib
       VIBs Removed:
       VIBs Skipped:
   ```

3. Confirm installation of the JAM VIB(s) in the ESXi host console by using SSH to run the following command:

   `esxcli software vib list`

Results

The RecoverPoint for VMs JAM installation name is displayed at the top of the list.

Next steps

Repeat this procedure for every ESX on which you want to install the RecoverPoint for VMs JAM VIB(s).

Add vRPAs to a vRPA cluster

Use this procedure to add a vRPA to an existing vRPA cluster. A vRPA cluster can have up to 8 vRPAs, and all vRPAs in a cluster must run the same RecoverPoint for VMs version.

Prerequisites

If you are not using the default vRPA web server certificate, ensure that your certificate is the same one used for all vRPAs in the vRPA cluster.

Steps

1. In a web browser, type `https://<cluster_management-ip-address>` for the vRPA cluster to which you want to add vRPAs.
2. In the home page, click **RecoverPoint for VMs Deployer**.
3. If prompted, type the login credentials for the admin user and click **Sign in**.
4. Under **More actions**, click **Add vRPAs to vRPA cluster**.
5. In the **Add Prerequisites** step, acknowledge that you have met the listed conditions by selecting the checkbox.
6. In the **Add vRPAs** step, select one or more VMs/vRPAs to add to the cluster.
   - New vRPAs must have the same RecoverPoint software ISO image as the existing vRPAs in the cluster.
   - A cluster can have a maximum of 8 vRPAs.
7. In the **vRPA Cluster Settings** and **vRPA Settings** sections, type required information for the vRPAs you are adding.
8. In the **Add vRPAs Progress** step, on reaching 100%, click **Finish** to return to the **Home Page**.

   If adding a vRPA fails:
   - To identify the cause of failure, review the displayed error messages.
   - To return to the step in the wizard where you can fix the problem, click **Back**. Fix the problem, and then retry the installation wizard from that point.
   - Alternatively, you can retry the operation that failed by clicking **Retry the operation**.
   - If adding a vRPA continues to fail, contact Customer Support.

Remove a vRPA from a vRPA cluster

Use this procedure to remove a vRPA from a vRPA cluster. You cannot remove a vRPA if the cluster has 2 or fewer vRPAs.

Steps

1. In a web browser, type `https://<cluster_management-ip-address>` for the vRPA cluster from which you want to remove a vRPA.
2. In the home page, click **RecoverPoint for VMs Deployer**.
3. If prompted, enter the login credentials for the admin user and click **Sign in**.
4. Under **More actions**, click **Remove vRPA from vRPA cluster**.
   - The highest numbered vRPA (the last one added) will be removed.
   - The consistency groups of the removed vRPA will be non-disruptively moved to a different vRPA.
   - The preferred vRPA setting for those consistency groups will be automatically updated.

Replace a vRPA

Use this procedure and wizard to replace a vRPA with a different vRPA.

Prerequisites

This wizard does not support replacing a vRPA within a vRPA cluster that has only one vRPA. If you must replace a vRPA in a single-vRPA cluster, contact Customer Support.

About this task

Deploy the new, replacement vRPA with the same IP settings as the faulty vRPA you want to replace. Ensure that the replacement vRPA is shut down. To shut down the replacement vRPA, login as admin user and select **Main Menu** > **Shutdown / Reboot operations** > **Shutdown RPA**.

Steps

1. In a web browser, type `https://<cluster_management-ip-address>` for the vRPA cluster in which you want to replace a vRPA.
2. In the home page, click **RecoverPoint for VMs Deployer**.
3. If prompted, type the login credentials for the admin user and click **Sign in**.

4. Under **More actions**, click **Replace vRPA**.

5. In the **Prerequisites** step, acknowledge that you have met the listed conditions by selecting the checkbox.

6. In the **Replace vRPA** step, select the vRPA that you want to replace.

7. Select the vRPA you want to add as a replacement.

8. In the **Replacement Progress** step, on reaching 100% click **Finish** to return to the home page.

   If replacing a vRPA fails:

   ● To identify the cause of failure, review the displayed error messages.

   ● To return to the step in the wizard where you can fix the problem, click **Back**. Fix the problem, and then retry the installation wizard from that point.

   ● Alternatively, you can retry the operation that failed by clicking **Retry the operation**.

   ● If replacing a vRPA continues to fail, contact Customer Support.

Change default passwords

To align with security best practices, when installing a vRPA cluster, you must replace the default password for the predefined admin user (admin) with a unique password.

About this task

The admin user password serves also as the password for the root user.

If you have upgraded your system from RecoverPoint for VMs 5.1.1.4, you could still be using the default admin (and root) password. If so, run this procedure to change the admin password,

Steps

1. Create an SSH connection to the vRPA management IP address, using your RecoverPoint for VMs admin username and password to log into the Boxmgmt CLI. Then select **System management CLI** to open the Sysmgmt CLI.

   Alternatively, if you have created a user with the sysmgmt role, use that user to log in directly to the Sysmgmt CLI.

2. In the Sysmgmt CLI, run the `set_password` command to change the password for the current user, or run the `set_user` command to change the password of another user, provided that your user/role includes the security permission.

   ⓘ **NOTE:** Keep passwords in a place where they are secure and available to you.

Collect logs

During deployment, collecting logs for the current vRPA cluster and its vRPAs provides information that may be helpful in troubleshooting the installation.

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 5.3 SP2 and later versions, you can initiate the vRPA cluster log collection process from within the RecoverPoint for VMs vSphere plugin. See "Collecting logs from vRPA clusters" in the *RecoverPoint for VMs Administrator's Guide* for more information.

Steps

1. In a web browser, type **https://<LAN-ip-address>** where *<LAN-ip-address>* is the vRPA cluster management IP address. In the vRPA cluster home page, click **RecoverPoint for VMs Deployer**.

2. If prompted, type the login credentials for the **admin** user and click **Sign in**.

3. Display the log collection settings in the **RecoverPoint for VMs Deployer** home page.

   ● In RecoverPoint for VMs 5.3 SP2 and later versions, click **Collect Logs**.
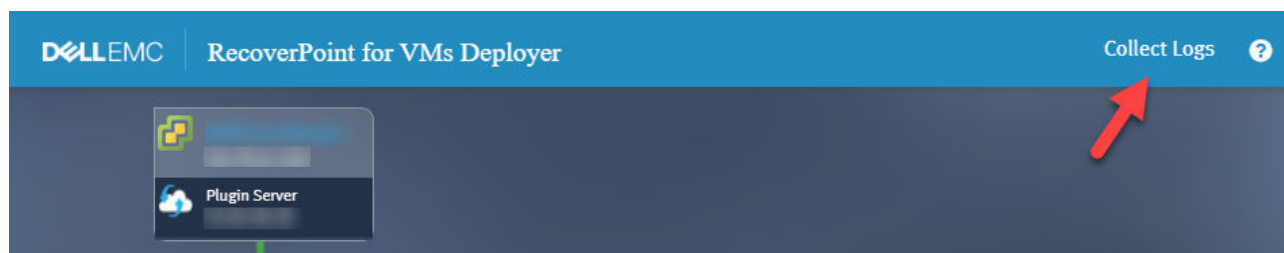
- In RecoverPoint for VMs versions prior to 5.3 SP2, at the upper right of the home page, click the **Settings** icon, and then click **Collect Logs**.

4. In the **Collect Cluster Logs** dialog box:

a. Enter a **Start time** and **End time** for log collection.
b. (Optional) To simultaneously collect logs from vRPAs in other vRPA clusters, expand the **Advanced** section and add another vRPA *<LAN-ip-address>*.
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.

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

Upgrading RecoverPoint for VMs

Upgrading RecoverPoint for VMs involves downloading the upgrade package and sequentially upgrading the vRPA clusters, the splitters and JAM VIBs, and the RecoverPoint for VMs plugin.

Topics:

Upgrade overview

Upgrade the RecoverPoint for VMs system from 5.2.x to 5.3.x, or from 5.3.x to a later 5.3.x, by downloading the desired upgrade package and then upgrading the relevant system components.

You can use the RecoverPoint for VMs vSphere plugin (**System** > **Administration** screen), API, or CLI (`get_versions` command) to determine what version of RecoverPoint for VMs system is currently installed.

In general, upgrading RecoverPoint for VMs includes the following activities:

- Download the upgrade package.
- Upgrade the vRPA clusters.
- Upgrade the RecoverPoint for VMs splitters.
- Upgrade the RecoverPoint for VMs JAM VIBs (only when the original installation was RecoverPoint for VMs 5.2 or later)
- Upgrade the Flex-based RecoverPoint for VMs plugin (when using vSphere 6.7 and earlier)

You may upgrade the RecoverPoint for VMs plugin server (when using vSphere 6.7 U1 and later) whenever a later version of the plugin server is available.

If you are using (or upgrading to) vCenter Server 6.7 U1, and upgrading from 5.2.x to 5.3.x, and planning to manage your system with the vSphere HTML5 plugin or the new REST API, Deploy the plugin server on page 21 and Configure the plugin server on page 23 for each vCenter Server.

ⓘ **NOTE:** When using RecoverPoint for VMs 5.3 and later versions with vCenter Server 6.5.x and 6.7.x, you may continue to manage your system with the vSphere Flex plugin.

When upgrading RecoverPoint for VMs, all existing RecoverPoint for VMs settings are preserved. There is no journal loss and no full sweep.

The Upgrade and Maintenance package

Download the RecoverPoint for VMs Upgrade and Maintenance Kit. The Upgrade and Maintenance Kit is a zip file that consists of multiple components required for the upgrade.

Download the RecoverPoint for VMs Upgrade and Maintenance Kit from https://www.dell.com/support.

In RecoverPoint for VMs 5.3.1 and later, the OVA and ISO files are protected by SRO and GPG signatures, respectively. If you want to verify a signature, see Verifying signatures on page 19

Upgrade a vRPA cluster

The RecoverPoint for VMs Deployer supports non-disruptive upgrades for vRPA clusters with two or more vRPAs and enables upgrading an ISO image without re-protecting VMs.

Prerequisites

If the vCenter Server SSL certificate has been changed, ensure that the new certificate is valid, and that RecoverPoint for VMs has been updated with it before beginning the upgrade.

About this task

- If you are upgrading a vRPA cluster with only one vRPA, the upgrade is disruptive to replication, but the upgrade occurs without full sweep or journal loss. Also, during the vRPA restart, the upgrade progress report may not update, and Deployer may become temporarily unavailable. When the vRPA completes its restart, you can log back into Deployer and observe the upgrade progress to completion.
- When you upgrade a cluster that has two or more vRPAs and is connected to a cluster with a single vRPA, a partially disruptive upgrade occurs. When the first vRPA is upgraded, all consistency groups move to another vRPA. However, for consistency groups that are replicated in the single vRPA, replication stops while the first vRPA is upgraded.

  ⚠ **CAUTION: Do not attempt to upgrade multiple connected clusters at the same time. This practice is not supported. Rather, upgrade connected vRPA clusters one cluster at a time until all of the connected vRPA clusters are upgraded to the same release.**

Steps

1. In a web browser, type `https://<cluster_management-ip-address>` for the vRPA cluster that you want to upgrade.
2. In the home page, click **RecoverPoint for VMs Deployer**.
3. If prompted, type the login credentials for the admin user, and click **Sign in**.
4. Click **Upgrade a vRPA cluster**.
   The wizard performs a system check.
5. In the **Upgrade Prerequisites** step, ensure that you meet the conditions that are listed on the screen. Select the checkbox: **I have fulfilled these conditions**.
6. In the **ISO** step, choose how you want to provide the ISO image for upgrading RecoverPoint for VMs.
7. In the **Change Version Requirements** step, the version requirements file is automatically downloaded and validated to ensure that the system meets the requirements. If the version requirements file fails to download, select one of these options:
   - **Retry downloading the up-to-date requirements from Online Support**
   - **Provide version requirements file manually**
   - **Do not check version requirements**

   Issues that are found are displayed for you to analyze. It is recommended that you fix blocking issues before continuing.
8. In the **System Diagnostics** step, Deployer checks for tweak modifications and signed scripts on the vRPAs. If discovered, these modifications are collected and the user is prompted to send the modifications file to Customer Support for analysis.
9. In the **Upgrade Progress** step, the progress bar displays the replacement progress. On reaching 100%, click **Finish** to return to the Deployer home page.

Results

All vRPAs in the vRPA cluster are upgraded.

Next steps

- If upgrade fails, review the displayed error message to identify the cause of the failure.
- To correct any issues and retry the upgrade, click **Back**.
- If upgrading a vRPA continues to fail, contact Customer Support.

Upgrade splitters and JAMs for entire ESX cluster

Use this procedure to upgrade splitters and JAM VIBs on all of the ESXi hosts in an ESX cluster.

Prerequisites

ⓘ **NOTE:** When upgrading to RecoverPoint 5.3 SP2 and later versions this procedure is not relevant. Instead:
- Upgrade splitter for single ESXi host on page 37, for each ESXi host.
- Upgrade JAM for single ESXi host on page 38, for each ESXi host.

- Ensure you have completed the procedure to Upgrade a vRPA cluster on page 36.
- To maintain uninterrupted replication throughout the upgrade of all splitters and JAMs, ensure that at least two ESXi hosts in the ESX cluster have an installed splitter.
- The ESX cluster must contain at least two ESXi hosts that are not already in maintenance mode prior to running the upgrade procedure.
- Ensure that DRS is enabled in automatic mode.

ⓘ **NOTE:** In VSAN environments:
- Upgrade splitter for single ESXi host on page 37, for each ESXi host.
- Upgrade JAM for single ESXi host on page 38, for each ESXi host.

Steps

1. Use an SSH client to log in to the vRPA as the admin user.
2. From the Main menu, select **Setup** > **Advanced options** > **Splitter actions** > **Upgrade Splitter**.
3. Enter the requested information: vCenter Server IP address and TCP port number (if other than the default, 443), and the vCenter credentials.
4. Let the system provide the vCenter certificate automatically and, if it's correct, approve the certificate.
5. Select the ESX cluster on which to upgrade the RecoverPoint for VMs splitters and JAM VIBs.
   The splitter and JAM VIB version currently installed on each of the ESXi hosts is listed, along with the splitter and JAM VIB version to be installed by the upgrade (that is, the VIBs for the currently installed RecoverPoint for VMs version).
6. Type **y** to begin the upgrade.
   Each ESXi host enters maintenance mode, in turn, as its splitter and JAM VIBs are upgraded. The table of splitter and JAM VIB versions is updated as the upgrade progresses.

Upgrade splitter for single ESXi host

Use this procedure to upgrade the RecoverPoint for VMs splitter on a single ESXi host.

Prerequisites

If you want vRPAs to keep replicating during splitter upgrade, ensure that at least two ESXi hosts have an installed splitter.

About this task

ⓘ **NOTE:** Do not erase the `/scratch` directory during upgrade.

Steps

1. On the ESXi host, vMotion all VMs to another ESXi host.
2. At the ESXCLI, enter maintenance mode. From the ESXi host console, use SSH to run the following command:
   ```
   esxcli system maintenanceMode set -e=true
   ```
   ⓘ **NOTE:** For VSAN environments, this command requires an additional switch (refer to the vSphere documentation for the vSphere version that you are using).

3. Remove the old RecoverPoint vSphere Installation Bundle on the ESXi host.
   ```
   esxcli software vib remove -n "RP-Splitter"
   ```

4. Install the splitter:
   a. To copy the RecoverPoint VIB to the `tmp` directory, use an SSH client with secure copy protocol:

   scp *<vib name> <username>*@*<ESXi host IP>*:/scratch

   For example, run:

   **scp EMC_bootbank_RP-Splitter_RPESX-00.5.3.2.0.0.m.222.000.vib root@10.10.10.10:/
   scratch**

   b. To install the splitter, in the ESXi host console:
      ● In RecoverPoint for VMs version 5.3 SP2 and later, run the command:

      esxcli software vib install -v /scratch/*<vib_full_path>*

      ● In RecoverPoint for VMs versions prior to 5.3 SP2, run the command:

      esxcli software vib install -v /scratch/*<vib_full_path>* --no-sig-check

      If installation is successful, the following message appears:

```
Installation Result
   Message: Operation finished successfully.
   Reboot Required: false
   VIBs Installed: EMC_bootbank_RP-Splitter_RPESX-<version>.vib
   VIBs Removed:
   VIBs Skipped:
```

   c. Confirm installation of the splitter in the ESXi host console by using SSH to run the following command:

   esxcli software vib list

   The RecoverPoint for VMs splitter VIB name is displayed at the top of the list.
5. On the ESXi host, exit maintenance mode by running the following command:

   esxcli system maintenanceMode set -e=false

6. vMotion the VMs back to this ESXi host.

Results

The splitter VIB is updated for the specified ESXi host.

Next steps

Repeat this procedure for each ESXi host.

Upgrade JAM for single ESXi host

Use this procedure to upgrade the RecoverPoint for VMs Journal Access Module (JAM) VIB for a single ESXi host. Upgrade the JAM VIB on all ESXi hosts that belong to ESX clusters on which vRPAs are running. In new RecoverPoint for VMs installations, the RecoverPoint for VMs JAM VIB is installed automatically on all ESXi hosts that belong to ESX clusters on which vRPAs are running.

Prerequisites

● Ensure a JAM VIB exists on the ESXi host prior to upgrade.
● Remove the **VM Storage Policy** of each vRPA VM on the ESXi cluster.
   1. Select the vRPA VM in vSphere, and click **Edit** > **VM Storage Policies**.

Edit VM Storage Policies | vRPA-2

Configure per disk ⬤

	Name	Disk Size	VM Storage Policy	Datastore		Datastore Type
>	🗀 VM home	-	RecoverPoint_JCD_6354af56ee0afe13 ⌄	UT-307-SITE1...		VMFS 5
>	Hard disk 1	35 GB	Datastore Default ⌄	UT-307-SITE1...		VMFS 5
>	Hard disk 2	1 MB	RecoverPoint_JCD_6354af56ee0afe13 ⌄	UT-307-SITE1...		VMFS 5

2. Write down and save the names of the policies applied to each VMDK.

   Some VMDKs will have a **RecoverPoint_JCD** policy and some will have a **Datastore Default** policy.

3. Set all **VM Storage Policies** to **Datastore Default**.

About this task

In RecoverPoint for VMs versions prior to 5.3 SP2, there was only one JAM VIB. In RecoverPoint for VMs 5.3 SP2 and later versions, there is one JAM VIB for ESXi version 6.x hosts and another for ESXi version 7.x hosts.

ⓘ **NOTE:** Do not erase the `/scratch` directory during upgrade.

Steps

1. On the ESXi host, vMotion all VMs to another ESXi host.
2. At the ESXCLI, enter maintenance mode. From the ESXi host console, use SSH to run the command:

   `esxcli system maintenanceMode set -e=true`

   ⓘ **NOTE:** For VSAN environments, this command requires an additional switch. Refer to the vSphere documentation for the vSphere version that you are using.

3. To remove the old JAM installation, run the command:

   `esxcli software vib remove –n emcjiraf`

4. Upgrade the JAM VIB on the ESXi host:

   You can either:
   - Manually install the JAM VIB (go on to step 5)
   - Have RecoverPoint for VMs automatically install the JAM VIB for you (skip to step 6).

5. (Optional) Manually install the new JAM VIB.

   a. Use an SSH client with secure copy protocol to copy the RecoverPoint JAM VIB:

   `scp <vib name> <username>@<ESXi host IP>:/scratch`

   For example:
   - On ESXi 6.x hosts, run: **`scp EMC_bootbank_emcjiraf_5320.m.186-1OEM.650.1.17.4598673.vib root@10.10.10.10:/scratch`**
   - On ESXi 7.x hosts, run: **`scp EMC_bootbank_emcjiraf_5320.m.180-1OEM.700.1.0.15843807.vib root@10.10.10.10:/scratch`**.

   b. To install the JAM VIB in the ESXi host console:
   - In RecoverPoint for VMs version 5.3 SP2 and later, run the command:

     `esxcli software vib install -v /scratch/<vib_full_path>`

   - In RecoverPoint for VMs versions prior to 5.3 SP2, run the command:

     `esxcli software vib install -v /scratch/<vib_full_path> --no-sig-check`

   If installation is successful, the following message appears:

   ```
   Installation Result
      Message: Operation finished successfully.
      Reboot Required: false
   ```

```
        VIBs Installed: EMC_bootbank_emcjiraf_<RP version><ESX version>.vib
        VIBs Removed:
        VIBs Skipped:
```

   c. Confirm installation of the JAM VIB in the ESXi host console by using SSH to run the command:

      esxcli software vib list

      The RecoverPoint for VMs JAM VIB name is displayed at the top of the list.

6. On the ESXi host, exit maintenance mode by running the command:

    esxcli system maintenanceMode set -e=false

    The RecoverPoint for VMs JAM VIB is updated on the ESXi host.

7. vMotion the VMs back to this ESXi host.

Next steps

Repeat this procedure for all ESXi hosts that belong to ESX clusters on which vRPAs are running.

Then, reset the **VM Storage Policy** of each vRPA VM on the ESXi cluster back to the state they were before upgrade:

1. Select the vRPA VM in vSphere, and click **Edit** > **VM Storage Policies**.
2. Change the names of the policies applied to each VMDK back to what they were in before upgrade.

Upgrade the RecoverPoint for VMs Web Client plugin

Use the vSphere Web Client to upgrade the Flex-based RecoverPoint for VMs Web Client plugin.

Prerequisites

The vRPA is backward-compatible, but the RecoverPoint for VMs plugin is not. New vRPAs work with older plugins, but a new RecoverPoint for VMs plugin might not be able to communicate with older vRPAs. Therefore, the RecoverPoint for VMs plugin version must correspond with the version of the oldest vRPA cluster.

About this task

Upgrade the Flex-based RecoverPoint for VMs Web Client plugin for each vCenter in the system.

Steps

1. Access the vSphere Web Client at: `https://<vCenter-ip-address>:9443/vsphere-client/`. In the vSphere Web Client home page, click the **RecoverPoint for VMs** icon.
2. Click the **Help...** link at the top right of the **RecoverPoint for VMs Management** screen, and select **Upgrade RecoverPoint for VMs**.
3. In the **Upgrade RecoverPoint for VMs** window, select the upgrade version and click **OK**.
4. Log out all active user sessions of the vSphere Web Client, and log back in. Verify that the RecoverPoint for VMs plugin is listed under **Inventories**.
5. If the RecoverPoint for VMs plugin is not listed under **Inventories**, restart the vCenter Web Client service to ensure that all active user sessions are disconnected.

   After upgrading the plugin, you will not be able to access vRPA clusters that are running earlier versions of RecoverPoint for VMs.

Upgrade 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. In the RecoverPoint for VMs HTML5 plugin for vSphere Client, click **System** > **Plugin Server**
   The **RecoverPoint for VMs Plugin Server** 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

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

Uninstalling RecoverPoint for VMs

You can uninstall a single vRPA cluster or all vRPA clusters from a vCenter.

The uninstaller tool scans the vCenter, datastores, and ESXi hosts. It removes vRPAs (production and copy VMs), configuration objects, and repository and journal volumes.

Topics:

Using the RecoverPoint for VMs uninstaller tool

The uninstaller tool removes vRPA clusters and their configuration entities from a vCenter.

Instructions for downloading the uninstaller tool are provided in Run the RecoverPoint for VMs uninstaller tool on page 44.

The uninstaller tool has the following options:

* `uninstall` - Uninstalls a single vRPA cluster from a vCenter. Use this option to:
  * Replace a Try and Buy or Beta version with a supported production version
  * Remove a vRPA cluster (after data migration)
  * Remove unwanted elements from the vCenter environment
* `full_rp_uninstall` - Uninstalls all vRPA clusters from a vCenter. Use this option to completely remove all RecoverPoint entities and clusters from the vCenter.

What the RecoverPoint for VMs uninstaller tool does

The RecoverPoint fo VMs uninstaller tool removes vRPAs, shadow VMs, configuration objects, and repository and journal volumes.

Running the RecoverPoint for VMs uninstaller tool does the following:

1. Scans the vCenter, datastores, and ESXi hosts.
2. Displays a list of all detected vRPA clusters and marks them either active or suspected inactive. Active clusters are clusters that have registered vCenter tokens during the last hour.
3. After you select which vRPA clusters the tool should uninstall, the tool removes the following from the selected vRPA clusters: production and replica VMs that were running vRPAs, shadow VMs (if they exist), RecoverPoint configuration objects, and the repository and journal volumes.

In addition to all the actions performed when uninstalling one vRPA cluster, uninstalling all vRPAs removes all vRPA clusters on the selected vCenter with all related elements. It also removes from the vCenter RecoverPoint elements not belonging to a specific vRPA cluster, such as the RecoverPoint vCenter plug-in.

The uninstaller tool does not remove plugin servers or splitter/JAM VIBs. For further instructions, see Finishing up the uninstall on page 45.

Preparing to uninstall vRPA clusters

Unprotect VMs

To stop replication for a vRPA, unprotect the associated VM.

Steps

1. In the vSphere Web Client home page, click the **RecoverPoint for VMs Management icon** > **Protection** tab. Click **Virtual Machines**.

   Alternatively, in the vSphere Client home page, open the **RecoverPoint for VMs** menu, and click **Protection** > **Protected VMs**.

2. Select the VM you wish to stop replicating. Click the **Unprotect** icon. Repeat for each protected VM.

Remove ESX clusters from vRPA clusters

Unregisters the ESX cluster of a production VM or copy VM, from a vRPA cluster.

Steps

1. Access the **RecoverPoint for VMs plugin** in your vSphere client.
   - In the **RecoverPoint for VMs HTML5 plugin**:
     a. Select **System** > **ESX Clusters**.
     b. If you are replicating remotely, select the **vRPA Cluster** from which you want to unregister the ESX cluster.
   - In the **RecoverPoint for VMs FLEX plugin**:
     a. Select **Administration** > **vRPA Clusters**, and select the **ESXi Clusters** tab.
     b. Select the relevant vRPA cluster.
2. Click the **Delete** icon next to each ESX cluster to unregister that ESX cluster from the specified vRPA cluster.

Results

The ESX cluster is unregistered from the specified vRPA cluster.

Uninstall a vRPA cluster

The Uninstall a vRPA cluster from this system wizard guides you in uninstalling a vRPA cluster. Use the uninstaller tool to uninstall the last vRPA cluster.

Prerequisites

- You cannot uninstall a cluster from a single-cluster system.
- After you uninstall a vRPA cluster, you cannot reuse the cluster or its vRPAs.
- When you uninstall a vRPA cluster, the vRPAs are shut down and cannot be restored.
- If required, collect logs before you uninstall the cluster. Log collection for the cluster is not possible later.

About this task

If you want to remove only one vRPA cluster from a system with two or more clusters, perform these steps from a vRPA cluster that is remaining in the system (and not from the cluster that you want to remove).

If you want to remove all of the vRPA clusters, perform these steps from one of the clusters. The last remaining cluster must be removed by using the uninstaller tool.

Steps

1. In a web browser, type `https://<cluster_management-ip-address>` for the vRPA cluster that you want to uninstall.
2. In the home page, click **RecoverPoint for VMs Deployer**.

3. If prompted, type the login credentials for the admin user and click **Sign in**.
4. Under **More actions**, click **Uninstall a vRPA cluster from this system**.
5. Select the vRPA cluster that you want to remove. Click **OK**.

   If cluster removal does not succeed, try again. If cluster removal fails, contact Customer Support.

Results

The vRPA cluster is successfully uninstalled. Continue to the next procedure.

Run the RecoverPoint for VMs uninstaller tool

Download the uninstaller tool, uncompress the .bat file, and run the tool in the command line.

Prerequisites

Obtain the IP and TCP port number of the vCenter (or vCenters) you want to scan, the vCenter username (and domain name, if one exists), and the vCenter password.

System requirements for the computer running the uninstaller tool:

● Microsoft Windows
● Java 7 or higher

ⓘ **NOTE:** If a time difference of more than 30 minutes exists between the vRPA and the computer running the uninstaller tool, the tool may recognize the vRPA cluster as inactive when it is not. (The time difference is not influenced by different time zones.)

Steps

1. Navigate to the RecoverPoint for Virtual Machines download page under the Tools & Utilities section and click the **RecoverPoint for Virtual Machines Uninstaller Tool** link. The link is also referenced from the Customer Installation kit.

   ⓘ **NOTE:** Use only the version of the uninstaller tool that is compatible with the RecoverPoint for VMs release that is installed on the vRPA cluster that you want to uninstall. For the latest support information, see the Simple Support Matrix for your version of RecoverPoint for VMs.

2. From a computer with IP connectivity to the vCenters managing the RecoverPoint VMs you want to uninstall, unzip the zip file.
3. Double click on `uninstaller.bat`.
   The RecoverPoint for VMs uninstaller tool opens in a command line.
4. Perform one of the following actions:

Option	Description
If uninstalling a single vRPA cluster from a vCenter	...type **uninstall**.
If uninstalling all vRPA clusters from a vCenter	...type **full_rp_uninstall**.

   Type **--h** after a command to view an explanation of that command. Type **help** to view a short explanation of all available commands.

5. Enter the IP address of the vCenter.
6. Enter the vCenter's TCP port number or press Enter for the default port (443).
7. Enter the vCenter's username.
8. Enter the vCenter's password.
   The tool tests connectivity and logs in to the vCenter.
9. Type **y** if you want to add another vCenter. Type **n** if you do not.

   If you have remote vRPA clusters connected to a different vCenter, type that vCenter's IP address if you want to uninstall that cluster as well.

   The tool displays a list of detected vRPA clusters.

10. Perform one of the following actions:

Option	Description
If uninstalling a single vRPA cluster from a vCenter	Type the index number of the vRPA cluster that you want to uninstall. To remove more than one cluster, type the index numbers separated by commas (for example: **1,4,9**).
If uninstalling all vRPA clusters from a vCenter	Type **y** to perform the uninstallation.

Results

The tool begins to scan and uninstall the cluster (or clusters).

If the process notifies you that it did not uninstall all objects, you may run the uninstall operation again.

RecoverPoint splitters and JAM are not removed by the uninstaller tool. They can be manually removed from each ESXi host. For instructions, see Uninstall the RecoverPoint for VMs splitters on page 45.

The uninstaller tool does not remove the RecoverPoint for VMs plugin server.

● Do not remove a plugin server that is still managing other vRPA clusters. Instead, *exclude* the vRPA cluster that you are uninstalling in order to stop the plugin server from managing that vRPA cluster. For instructions on excluding a vRPA cluster (and clearing excluded clusters), see "Managing vRPA clusters on vCenters registered with the plugin server" in the *RecoverPoint for VMs HTML5 Plugin Administrator's Guide*.
● If you want to uninstall the plugin server, see Uninstall the plugin server on page 45.

Finishing up the uninstall

Perform the following tasks to complete the uninstall procedure.

Uninstall the plugin server

Use this procedure to uninstall the RecoverPoint for VMs plugin server.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client, select **System** > **Plugin Server**, and then select the vCenter Servers tab.
2. Delete all vCenter Servers from the plugin server.
3. Power off and remove the plugin server VM.

Results

The plugin server is uninstalled.

Uninstall the RecoverPoint for VMs splitters

Use the ESXCLI to uninstall the RecoverPoint for VMs splitters.

Steps

1. Use ESXCLI to obtain a list of all installed vSphere Installation Bundles (VIBs):
   esxcli software vib list
2. Ensure that a bundle named *RP-Splitter* is installed.
3. On the ESXi host, enter maintenance mode:
   esxcli system maintenanceMode set -e=true

   (i) **NOTE:** For VSAN environments, this command requires an additional switch (refer to the vSphere documentation for the vSphere version that you are using).

4. To uninstall the splitter, type:

`esxcli software vib remove -n "RP-Splitter"`

(i) **NOTE:** If you want to remove the JAM VIB also (while still in maintenance mode), run the command:

`esxcli software vib remove -n emcjiraf`

5. On the ESXi host, exit maintenance mode:

`esxcli system maintenanceMode set -e=false`

Uninstall JAM

Use this procedure to uninstall the JAM VIB from an ESXi host.

Steps

1. On the ESXi host, vMotion all VMs to another ESXi host.
2. At the ESXCLI, enter maintenance mode. From the ESXi host console, use SSH to run the command:

`esxcli system maintenanceMode set -e=true`

(i) **NOTE:** For VSAN environments, this command requires an additional switch. Refer to the vSphere documentation for the vSphere version that you are using.

3. To remove the old JAM installation, run the command:

`esxcli software vib remove -n emcjiraf`

(i) **NOTE:** If you want to remove the splitter VIB also (while still in maintenance mode), run the command:

`esxcli software vib remove -n "RP-Splitter"`

4. On the ESXi host, exit maintenance mode:

`esxcli system maintenanceMode set -e=false`

Removing unused directories

From the vCenter, remove RecoverPoint for VMs Flex plugin directories that are no longer being used.

Steps

1. Connect to the vCenter server (using a local network mapping or Remote Desktop Connection). Delete the following folder:
   - vCenter 6.5/6.7 and Windows vCenter: `C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\vc-packages\vsphere-client-serenity\com.emc.recoverpoint.vwc-<version>`
   - vCenter 6.5/6.7 and vCSA: `/etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity/com.emc.recoverpoint.vwc-<version>`
2. Restart the vSphere Web Client. For instructions, refer to VMware KB1003895 (Windows), VMware KB2109887 (Linux - vCenter 6.x), or VMware KB2147152 (Linux - vCenter 6.5).

Installing in VxRail environments

Installing RecoverPoint for VMs in VxRail environments is similar to a standard installation, but includes a few specific requirements for preparing the network, configuring VMkernel ports, creating vRPAs and vRPA clusters, and adding VxRail appliances or nodes.

The topics in this chapter provide procedures needed for installing RecoverPoint for VMs with VxRail.

Topics:

* Deploying RecoverPoint for VMs in a VxRail™ environment

Deploying RecoverPoint for VMs in a VxRail™ environment

Follow specific guidelines when deploying Recoverpoint for VMs in a VxRail environment.

About this task

When deploying RecoverPoint for VMs on VxRail appliances, follow the guidelines in this chapter along with the instructions that are listed in Preparing the network on page 16, Deploy vRPAs on page 20, and Install vRPA clusters on page 22.

Downloading from the VxRail Marketplace

Download the latest qualified RecoverPoint for VMs release.

About this task

Download the latest RecoverPoint for VMs release from the VxRail manager marketplace or RecoverPoint for Virtual Machines download page.

Preparing the network for VxRail

Prepare the network for the VxRail environment by choosing a network adapter topology and defining the required ports.

About this task

VxRail supports adding a PCIe NIC to the node in E, P, S, and V Series. VxRail initialization does not impact the PCIe NIC. You can connect unused ports to the VxRail system vSphere Distributed switch. Alternatively, you can create a vSphere Standard Switch (VSS)/vSphere Distributed switch and connect the unused ports after initialization. The ports are available for uses such as RecoverPoint traffic. VxRail G-Series 2x10G models have only 2x10G ports.

In VxRail 4.0 and later, vSphere Network I/O Control (NIOC) is enabled during initialization, and vSAN traffic has the highest priority to consume the bandwidth in contention. If NIOC is enabled with the default VxRail setting, you can use the vSAN port for other traffic.

The configuration described here uses the G-Series 2x10G model uplink configuration and VxRail system vSphere Distributed switch default name ("VMware HCIA Distributed Switch") as an example. Choose a vSwitch and uplink name according to the VxRail model and uplink configuration.

Prepare the required port groups on the VMware HCIA Distributed Switch.

Steps

- Use the default configuration of two network adapters (WAN + LAN on one adapter and data on the other) unless required to use a different network adapter topology. For this configuration, define two port groups: RP_WAN+LAN and RP_DATA.
- If using a single network adapter, define one port group: RP_ALL.
- If using four network adapters, define four port groups: RP_WAN, RP_LAN, RP_DATA1, RP_DATA2.

Create vRPAs for VxRail

Use the OVA file and guidelines in this procedure to create vRPAs for VxRail environments.

About this task

When creating vRPAs:

Steps

1. In the **Select storage** screen, in the **VM Storage Policy** drop down, select **VxRail-Virtual-SAN-Datastore**. The compatible VSAN datastore will be selected.
2. Deploy two vRPAs and configure VM-Host affinity rules to avoid running both vRPAs on the same ESXi node (recommended).

Create and configure VMkernel ports for VxRail

Create and configure VMkernel ports for VxRail environments.

Steps

1. Create one or two VMkernel ports on each ESXi node by selecting an existing distributed vSwitch "VMware HCIA Distributed Switch."

   A single VMkernel port is required when using the default of two network adapters (WAN + LAN on one network adapter and data on the other network adapter). This configuration is standard. Two VMkernel ports are required when you are using two network adapters for data.
2. To select one network adapter (uplink) as active, override the NIC teaming policy. The other network adapter should be marked as unused.

   When using a single VMKernel port, assign uplink1 to the port.

   When using two VMkernel ports:
   - Assign uplink1 to one VMkernel port and uplink2 to the second VMkernel port.
   - For uplink2, use traffic shaping to limit bandwidth to no more than 1Gb/s (if NIOC is enabled with the default VxRail setting, traffic shaping is optional):
     a. Locate the port group, right-click it, and select **Edit Settings**.
     b. In the **Edit Settings** window, change traffic shaping for the port group:

Traffic shaping	Field	Value
Ingress	Peak bandwidth (kb/s)	1048576
	Burst size (KB)	102400
Egress	Status	Enabled
	Average bandwidth (kb/s)	1048576
	Peak bandwidth (kb/s)	1048576
	Burst size (KB)	102400

Create a vRPA cluster for VxRail

Use the Install a vRPA cluster wizard to create a vRPA cluster for the VxRail environment.

About this task

When creating a vRPA cluster:

Steps

1. In the **Environment Settings** step, select the vSAN datastore from the table of available datastores.
2. In the **Network Settings** step of the wizard, specify the vRPA data network addresses (and not the VMkernel port IP addresses that were created earlier).

Adding VxRail appliances or nodes

Adding a VxRail appliance or node requires verifying the node addition, configuring ESXi nodes, registering the new ESXi clusters, and adjusting VM-host affinity rules.

About this task

After adding a VxRail appliance or a node to an existing appliance:

Steps

1. Verify that the nodes are added into the same vSAN cluster and under the same vCenter.
2. Configure each ESXi node with the required data network adapters for enabling splitter-to-vRPA communication.
3. If you created a new ESXi cluster, register it within the vRPA cluster.
   This action installs the RecoverPoint for VMs splitters on the new ESXi nodes.
4. Adjust VM-host affinity rules for the vRPAs to ensure that they are running on separate ESXi servers.

8

Troubleshooting RecoverPoint for VMs installation

When the RecoverPoint for VMs installation is not successful, knowing how to troubleshoot the vRPAs, splitters, RecoverPoint for VMs plug-in, and replication helps you to fix the problem.

These troubleshooting procedures use the vSphere Web Client.

Some commands that may be useful in troubleshooting can be run from the root user, including: ethtool, kps.pl, ping6, uptime, date, ssh, telnet, arping32, switch utils (Dell EMC Customer Support only), netstat, arp, ping, top, and su.

Topics:

- Troubleshooting vRPAs
- Troubleshooting splitters
- Troubleshooting the RecoverPoint for VMs plug-in
- Troubleshooting RecoverPoint for VMs replication
- ESXi UUID duplication

Troubleshooting vRPAs

This section describes how to troubleshoot these vRPA conditions:

- vRPA is down
- vRPA is detached from cluster
- vRPA does not see storage or splitter

vRPA is down

If a vRPA is down (powered off), check for vRPA errors, vRPA cluster status, and conflicts in the vRPA resource reservation. To investigate the root cause, collect and analyze logs. From the vSphere Web Client, power on the vRPA.

Steps

1. Check the RecoverPoint for VMs dashboard for Error events indicating that the vRPA is not online.
2. Log in to a surviving vRPA and type the RecoverPoint admin username and password to log in to the Boxmgmt CLI. Then select **System management CLI** to open the Sysmgmt CLI. Alternatively, if you have created a user with the sysmgmt role, use that user to log in directly to the Sysmgmt CLI. To check the cluster status, use the `get_system_status` Sysmgmt CLI command. Choose to retrieve the status of all categories.
3. Confirm that the failed vRPA cannot be reached.
4. Check any conflicts in the vRPA resource reservation that might have led to the vRPA being powered off. Resolve any issues before proceeding.
5. In the vSphere Web Client, right-click the vRPA that is down and select **All vCenter Actions** > **Power** > **Power On**.
6. To ensure that the vRPA was powered on successfully, monitor the vRPA console in the vSphere Web Client.
7. To investigate the root cause of the vRPA failure, collect logs.

vRPA is detached from the vRPA cluster

If the vRPA is detached from the vRPA cluster, check for vRPA errors and cluster status.

Steps

1. Check the RecoverPoint for VMs dashboard for Error events indicating that the vRPA cannot access storage or communicate with the splitters.
2. Create an SSH connection to a surviving vRPA, using your RecoverPoint for VMs admin username and password to log into the Boxmgmt CLI. Then select **System management CLI** to open the Sysmgmt CLI. Alternatively, if you have created a user with the sysmgmt role (RecoverPoint for VMs 5.2.0.2 or later), use that user to log in directly to the Sysmgmt CLI. Use the `get_system_status` Sysmgmt CLI command to check the cluster status. Choose to retrieve the status of all categories.
3. Confirm that the detached vRPA cannot be reached from the surviving vRPA.
4. Log in to the Boxmgmt CLI of the detached vRPA using admin username and password, and select **Cluster operations** > **Attach RPA to Cluster**. To ensure that the vRPA was powered on successfully, monitor the vRPA console in the vSphere Web Client.
5. To investigate the root cause of the vRPA detachment from the cluster, collect logs.
6. If you are using a licensed version of RecoverPoint for VMs, contact Customer Support.

vRPA cannot detect storage or splitter

When a vRPA cannot detect the storage or the splitters, investigate the status of the vRPA and splitters. Collect and analyze the logs.

Steps

1. Ensure the vRPA is online.
2. Ensure the vRPA is attached to the cluster.
3. Verify that the splitters are running:
   a. Login to ESXi hosts.
   b. Run: `ps |grep kdriver`
   c. Ensure that splitter processes are running.
4. To investigate the root cause of why the vRPA went down, collect logs.
5. If you are using a licensed version of RecoverPoint for VMs, contact Customer Support.

Troubleshooting splitters

The section describes how to troubleshoot the splitter when it is not visible or is in error state.

Splitter is not visible or in error state

To determine why the splitter is not visible or in error, check splitter processes and investigate logs.

Steps

1. If possible, vMotion any protected VMs from ESXi hosts with splitters in error state continue or resume replication.
2. Ensure that the splitter processes are running on the host you are troubleshooting:
   a. Login to the ESXi host and use the following command: `ps |grep kdriver`
   b. If processes are not running, place the ESXi node in maintenance mode and restart.
3. To investigate the root cause of the splitter failure, collect logs.
4. If you are using a licensed version of RecoverPoint for VMs, contact Customer Support.

Troubleshooting the RecoverPoint for VMs plug-in

This section describes how to troubleshoot these conditions:

- vSphere Web client does not contain plug-in
- Plug-in does not see the vRPA cluster

vSphere Web client does not contain plug-in

About this task

Go through the following steps until the problem is resolved:

Steps

1. Log out of vSphere Web client and log back in. Check if the RecoverPoint for VMs plug-in is listed under **Inventories**.
2. If the RecoverPoint for VMs plug-in is not listed, close all active vSphere Web client user sessions. Then check if the RecoverPoint for VMs plug-in is listed under **Inventories**.
3. If the RecoverPoint for VMs plug-in is still not listed, restart the vCenter Web Client service.
4. If the plug-in is still not visible in the vSphere Web Client, validate the vCenter Credentials configuration. You may need to reconfigure vCenter credentials. Consult Customer Support if protected VMs exist.
5. If the plug-in is still not visible in the vSphere Web Client, collect logs to investigate the root cause of why the plug-in is not visible.

Plug-in does not detect the vRPA cluster

About this task

Go through the following steps until the problem is resolved:

Steps

1. Log out of the vSphere Web Client and log back in.
2. Refresh the vSphere Web Client.
3. Log out all users from the vSphere Web Client.
4. Restart the vSphere Web Client.
5. Log in to the Managed Object Browser at `https://<vSphere Web Client>/mob`. Ensure the vCenter credentials are configured correctly.

   Access to the Managed Object Browser is disabled by default in vSphere 6.0 and later. For instructions on how to enable access, refer to VMware KB2108405.
6. Restart vRPA1.
7. Restart vRPA2.
8. To investigate the root cause of the vRPA failure, collect logs.
9. If you are using a licensed version of RecoverPoint for VMs, contact Customer Support.

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

(i) **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.

Troubleshooting RecoverPoint for VMs replication

This section describes how to troubleshoot these conditions:

- Consistency group is in high-load transfer state or initialization is not completing
- Consistency group is in error state

CG in high-load transfer state or initialization not completing

Steps

1. If consistency groups are not balanced across vRPAs, create an SSH connection to the vRPA management IP address, and type the RecoverPoint admin username and password to log in to the Boxmgmt CLI. Then select **System management CLI** to open the Sysmgmt CLI. Alternatively, if you have created a user with the sysmgmt role (RecoverPoint for VMs 5.2.0.2 or later), use that user to log in directly to the Sysmgmt CLI. Run the `balance_load` Sysmgmt CLI command and change consistency group assignments. For more information about load balancing, see the *RecoverPoint for Virtual Machines Administrator's Guide*.

2. If the throughput required by a consistency group exceeds the availability on a single vRPA, review the vRPA profile to see if additional resources can be added to meet higher IOPS requirements.

3. Enabling deduplication when WAN compression is also enabled may overload the vRPA and consequently degrade replication performance. For more information, see the *RecoverPoint for Virtual Machines Scale and Performance Guide*.

   a. It is recommended to enable WAN and journal compression and disable deduplication.

   b. If the consistency group contains more than one VM, consider moving VMs to dedicated consistency groups and using group sets as needed.

   c. Consider adding vRPAs (up to 8) to the vRPA cluster for additional resources.

   d. Review ESXi resources to ensure that there is no contention.

4. I/O throttling is initially enabled at 400 MBps. If needed, create an SSH connection to the vRPA management IP address, and type the RecoverPoint admin username and password to log in to the Boxmgmt CLI. Then select **System management CLI** to open the Sysmgmt CLI. Alternatively, if you have created a user with the sysmgmt role, use that user to log in directly to the Sysmgmt CLI. Run the `config_io_throttling` command to slow down production storage reads during full sweep process.

Consistency group is in Error state

Steps

1. Perform all of the procedures suggested for a consistency group in high-load state.

2. If the consistency group is still in Error state, try the following:

   a. Check if the image access buffer is full. If so, disable image access.

   b. Resolve any WAN issues.

   c. Check if the consistency group is in a permanent high-load state.

3. To investigate why the consistency group is in error state, collect logs.
4. If you are using a licensed version of RecoverPoint for VMs, contact Customer Support.

ESXi UUID duplication

In the VMware environment, each ESXi host is assigned a Universally Unique ID (UUID). RecoverPoint for VMs uses these UUIDs to maintain the integrity of replicated copies and protect the ESXi hosts from data corruption.

However, in some cases, a UUID might change with results that include:

● More than one ESXi host within a cluster reporting the same UUID.
● A single ESXi host reporting a different UUID after host restart (or similar operations).
● A single ESXi host reporting a degenerated UUID with all 0's or F's.

These cases can occur when using hardware that is not certified by VMware because the UUID is based on the BIOS UUID reported by the underlying server hardware. For more information about duplicate UUIDs, see VMware Knowledgebase Article 2006865.

Duplicate or degenerated UUIDs can cause the following:

● The RecoverPoint cluster can experience reboot regulation (vRPAs restarting over and over again until they detach from the cluster).
● The RecoverPoint consistency groups may not be able to recognize, connect to, or communicate with the splitter on the affected ESXi hosts.

RecoverPoint for VMs replaces the use of VMware's ESXi host UUID and creates its own unique identifier, which ensures that no duplicate or degenerated UUIDs exist in the system. The substitution occurs only if the:

● vRPA cluster version supports this feature
● Splitter version supports this feature

For versions that do not support this feature, RecoverPoint for VMs displays a warning about the condition.

RecoverPoint for VMs installation form

The installation form is a data sheet or spreadsheet that lists the site-specific values that you require to successfully complete the installation.

To streamline the installation tasks, create the RecoverPoint for VMs installation form during the planning phase.

Topics:

•   Installation data forms

Installation data forms

The best practice for successful installations is to collect and document required data before the installation.

The forms that are provided in this section are examples of the types of information you should collect before installation. You can create a planning spreadsheet that matches specific requirements (number of vRPA clusters, network topology, and so forth).

You are directed to type the data from these forms (or similar data sheet) during the installation process.

Table 2. Example: vRPA cluster/site form

vRPA cluster	vRPA cluster 1	vRPA cluster 2	vRPA cluster 3	vRPA cluster 4	vRPA cluster 5
Cluster/site name					
Time zone					
Local domain					
Primary DNS server (optional)					
Secondary DNS server (optional)					
Primary NTP server (recommended)					
Secondary NTP server (recommended)					
Cluster management IP					
Management default gateway IP					
Management subnet mask IP					
WAN default gateway					
WAN subnet mask					
SMTP (optional)					
vCenter IP					
vCenter credentials					

Table 2. Example: vRPA cluster/site form (continued)

vRPA cluster	vRPA cluster 1	vRPA cluster 2	vRPA cluster 3	vRPA cluster 4	vRPA cluster 5
vCenter credentials					
Plugin server (per vCenter Server)					
VMkernel IP pool					
ESXi 1					
_Data1 IP					
_Data2 IP					
_Management IP					
ESXi 2					
_Data1 IP					
_Data2 IP					
_Management IP					

Table 3. Example: vRPA IP form

vRPA	vRPA IPs	Site: _____	Site: _____	Site: _____	Site: _____	Site: _____
vRPA_1	LAN IP					
	WAN IP					
	Data1 IP					
	Data2 IP					
vRPA_2	LAN IP					
	WAN IP					
	Data1 IP					
	Data2 IP					
vRPA_3	LAN IP					
	WAN IP					
	Data1 IP					
	Data2 IP					
vRPA_4	LAN IP					
	WAN IP					
	Data1 IP					
	Data2 IP					
vRPA_5	LAN IP					
	WAN IP					
	Data1 IP					
	Data2 IP					
vRPA_6	LAN IP					
	WAN IP					
	Data1 IP					

Table 3. Example: vRPA IP form (continued)

vRPA	vRPA IPs	Site: _____	Site: _____	Site: _____	Site: _____	Site: _____
	Data2 IP					
vRPA_7	LAN IP					
	WAN IP					
	Data1 IP					
	Data2 IP					
vRPA_8	LAN IP					
	WAN IP					
	Data1 IP					
	Data2 IP					

Table 4. Example: Site map

Site	Site1 (Prod)	Site2 (Remote)	Site2 (Remote)	Site3 (Remote)	Site3 (Remote)
Cluster	Cluster1	Cluster2	Cluster3	Cluster4	Cluster5
vCenter Server	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*
Plugin server (per vCenter Server)	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*
ESXi 1	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*
ESXi 2	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*	*<name>* *<ip_address>*
vRPA 1	*<ip_address>*	*<ip_address>*	*<ip_address>*	*<ip_address>*	*<ip_address>*
vRPA 2	*<ip_address>*	*<ip_address>*	*<ip_address>*	*<ip_address>*	*<ip_address>*
Cluster Mgmt	*<ip_address>*	*<ip_address>*	*<ip_address>*	*<ip_address>*	*<ip_address>*
NIC1 IP (WAN)	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2
Data 1	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2
Data 2	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2	*<ip_address>* RPA1 *<ip_address>* RPA2

Support procedures for uninstalling vRPA clusters

When the automated uninstaller tool is unavailable, you can manually uninstall a vRPA cluster by using support procedures to guide you.

The topics in this Appendix provide procedures for use in manually uninstalling a vRPA cluster.

Topics:

- Uninstalling a single vRPA cluster from a vCenter manually
- Uninstalling all vRPA clusters from a vCenter manually
- Unprotect VMs
- Remove ESX clusters from vRPA clusters
- Remove a vRPA from a vRPA cluster
- Detaching vRPAs
- Powering off vRPAs
- Deleting the repository folder
- Verifying that the configuration parameters are empty
- Removing custom tokens from the Managed Object Browser
- Unregistering the plugin from the Managed Object Browser
- Removing unused directories
- Uninstall the RecoverPoint for VMs splitters
- Uninstall JAM

Uninstalling a single vRPA cluster from a vCenter manually

Perform this procedure to uninstall one vRPA cluster from a vCenter.

Prerequisites

Obtain the internal cluster name of the vRPA you are uninstalling by connecting to the vRPA Boxmgmt CLI as the admin user, and selecting **Main Menu** > **Setup** > **View Settings**.

About this task

You must perform this procedure for each vRPA cluster you want to uninstall.

If removing the last vRPA cluster on the vCenter, use the procedure Uninstalling all vRPA clusters from a vCenter manually on page 60 instead of this one.

Steps

1. If the vRPA cluster is active:
   a. Unprotect the virtual machines. For more information, see Unprotect VMs on page 43.
   b. Remove all ESX clusters from the vRPA cluster. For more information, see Remove ESX clusters from vRPA clusters on page 43. Repeat this step for all ESX clusters in the vRPA cluster.
2. If you are removing just one vRPA cluster from a system with at least two clusters, perform the following procedure: Uninstall a vRPA cluster on page 43.
   If the procedure Uninstall a vRPA cluster on page 43 was successful, skip to step 6 on page 60. If the procedure Uninstall a vRPA cluster on page 43 failed, continue with the next step.

3. Detach the vRPAs from the cluster. For more information, see Detaching vRPAs on page 62.
4. Power off the vRPAs. For more information, see Powering off vRPAs on page 62.
5. Remove the custom tokens that correspond to the RecoverPoint for VMs cluster ID. For more information, see Removing custom tokens from the Managed Object Browser on page 63.
6. Delete from all datastores the repository folder of the cluster you are uninstalling. For more information, see Deleting the repository folder on page 62.
7. Verify that the configuration parameters are empty. For more information, see Verifying that the configuration parameters are empty on page 62.

   (i) **NOTE:** Perform this step only if you encountered problems when unprotecting the VMs. Performing this step requires downtime of the production VM.

8. Ensure that the vRPA virtual machines are powered off, and delete them.
9. If the ESX cluster you are removing is not registered to any other vRPA cluster, you can uninstall the RecoverPoint for VMs splitter on that ESXi host. For more information, see Uninstall the RecoverPoint for VMs splitters on page 45.

Uninstalling all vRPA clusters from a vCenter manually

Perform this procedure to uninstall all vRPA clusters from a vCenter.

Prerequisites

Obtain the internal cluster name of the vRPA you are uninstalling by connecting to the vRPA Boxmgmt CLI as the admin user, and selecting **Main Menu** > **Setup** > **View Settings**.

About this task

You must perform this procedure for each vRPA cluster you want to uninstall.

Steps

1. If the vRPA cluster is active:
   a. Unprotect the virtual machines. For more information, see Unprotect VMs on page 43.
   b. Remove all ESX clusters from the vRPA clusters. For more information, see Remove ESX clusters from vRPA clusters on page 43. Repeat this step for all ESX clusters in all vRPA clusters.
2. If you are removing just one vRPA cluster from a system with at least two clusters, perform the following procedure: Uninstall a vRPA cluster on page 43.
   If the procedure Uninstall a vRPA cluster on page 43 was successful, skip to step 6 on page 60. If the procedure Uninstall a vRPA cluster on page 43 failed, continue with the next step.
3. Detach the vRPAs from the cluster. For more information, see Detaching vRPAs on page 62.
4. Power off the vRPAs. For more information, see Powering off vRPAs on page 62.
5. Remove the custom tokens that correspond to the RecoverPoint for VMs Internal cluster name. For more information, see Removing custom tokens from the Managed Object Browser on page 63.
6. Delete from all datastores the repository folders of all clusters. For more information, see Deleting the repository folder on page 62.
7. Verify that the configuration parameters are empty. For more information, see Verifying that the configuration parameters are empty on page 62.

   (i) **NOTE:** Perform this step only if you encountered problems when unprotecting the VMs. Performing this step requires downtime of the production VM.

8. Ensure that the vRPA virtual machines are powered off, and delete them.
9. Unregister the plug-in from the Managed Object Browser. For more information, see Unregistering the plugin from the Managed Object Browser on page 63.
10. Uninstall the RecoverPoint for VMs splitter. For more information, see Uninstall the RecoverPoint for VMs splitters on page 45.

11. Unregister the RecoverPoint extension from the Managed Object Browser. For more information, see .
12. Remove the RecoverPoint datastore element. Delete the `RecoverPoint.flp` file located in the RecoverPoint folder.

Unprotect VMs

To stop replication for a vRPA, unprotect the associated VM.

Steps

1. In the vSphere Web Client home page, click the **RecoverPoint for VMs Management icon** > **Protection** tab. Click **Virtual Machines**.

   Alternatively, in the vSphere Client home page, open the **RecoverPoint for VMs** menu, and click **Protection** > **Protected VMs**.
2. Select the VM you wish to stop replicating. Click the **Unprotect** icon. Repeat for each protected VM.

Remove ESX clusters from vRPA clusters

Unregisters the ESX cluster of a production VM or copy VM, from a vRPA cluster.

Steps

1. Access the **RecoverPoint for VMs plugin** in your vSphere client.
   - In the **RecoverPoint for VMs HTML5 plugin**:
     a. Select **System** > **ESX Clusters**.
     b. If you are replicating remotely, select the **vRPA Cluster** from which you want to unregister the ESX cluster.
   - In the **RecoverPoint for VMs FLEX plugin**:
     a. Select **Administration** > **vRPA Clusters**, and select the **ESXi Clusters** tab.
     b. Select the relevant vRPA cluster.
2. Click the **Delete** icon next to each ESX cluster to unregister that ESX cluster from the specified vRPA cluster.

Results

The ESX cluster is unregistered from the specified vRPA cluster.

Remove a vRPA from a vRPA cluster

Use this procedure to remove a vRPA from a vRPA cluster. You cannot remove a vRPA if the cluster has 2 or fewer vRPAs.

Steps

1. In a web browser, type **`https://<cluster_management-ip-address>`** for the vRPA cluster from which you want to remove a vRPA.
2. In the home page, click **RecoverPoint for VMs Deployer**.
3. If prompted, enter the login credentials for the admin user and click **Sign in**.
4. Under **More actions**, click **Remove vRPA from vRPA cluster**.
   - The highest numbered vRPA (the last one added) will be removed.
   - The consistency groups of the removed vRPA will be non-disruptively moved to a different vRPA.
   - The preferred vRPA setting for those consistency groups will be automatically updated.

Detaching vRPAs

Steps

1. Use an SSH client to connect to a vRPA and enter login credentials for the admin user.
2. From the **Main Menu**, select **Cluster Operations** > **Detach from Cluster**.
   Replication is paused.
3. Repeat this procedure on all vRPAs in all vRPA clusters in the system.

Powering off vRPAs

Steps

1. In the vSphere Client or vSphere Web Client, select **VMs and Templates**.
2. In the vSphere Client, select each vRPA, right-click and select **Power** > **Power Off** for each vRPA that is powered on.
   In the vSphere Web Client, select each vRPA, right-click and select **All vCenter Actions** > **Power** > **Power Off** for each vRPA that is powered on.

Deleting the repository folder

Steps

1. At the vSphere Web Client, select **Home** > **Storage** > **Manage**.
2. Select the datastore where the repository folder was created.
3. In the list of files displayed in the **Files** subtab, locate and open the `RPvStorage` folder.
4. Within the `RPvStorage` folder, delete all folders and/or files that include the Internal cluster name.

Verifying that the configuration parameters are empty

About this task

(i) **NOTE:** Performing this task requires downtime of the production VM.

Steps

1. At the vSphere Web Client, in **Inventory**, select **Hosts and Clusters**. Select a VM that was protected by RecoverPoint for VMs. Power off the VM. Right-click and select **Edit Settings...**
2. In the **Edit Settings** dialog box, select the **VM Options** tab. Expand the **Advanced** column. In the **Configuration Parameters** row, click **Edit Configuration...** to edit the advanced configuration parameters.
3. In the **Configuration Parameters** window, ensure that all configuration parameters with "RecoverPoint" or "esx_splitter" in the name have empty values.
   The following parameters must not exist or have empty values:
   ● RecoverPoint RPA number
   ● RecoverPoint CGUID
   ● RecoverPoint Cluster ID
   ● esx_splitter.globalOptions
   ● esx_splitter.scsi0:1.options

Removing custom tokens from the Managed Object Browser

About this task

The custom tokens that correspond to the RecoverPoint for VMs cluster ID need to be removed from the cluster(s) being reinstalled for all previously used vCenters.

ⓘ **NOTE:** Access to the Managed Object Browser is disabled by default in vSphere 6.0. For instructions on how to enable access, refer to VMware KB2108405.

Steps

1. In a web browser, enter the fully-qualified domain name (or IP address) of the vCenter Server system:

   https://*<hostname.yourcompany.com>*/mob/?moid=CustomFieldsManager

2. Log in using your vCenter login credentials.

3. In the **Methods** table, select **RemoveCustomFieldDef**.
   A new browser window opens with the **void RemoveCustomFieldDef** command displayed.

4. In the **Parameters** table, enter the value of a custom field listed in the **Properties** table that corresponds to the Internal cluster name, RecoverPoint_TOKEN, for example, `config.RecoverPoint_TOKEN;3070371118132351610`.

5. Click **Invoke Method**.

6. If you are reinstalling several clusters, repeat steps 3 through 5 for each custom field listed in the **Properties** table that corresponds to the Internal cluster names.

Unregistering the plugin from the Managed Object Browser

About this task

Unregister the RecoverPoint for VMs plugin from the Managed Object Browser at each vCenter Server that contains ESXi clusters that are hosting vRPA clusters. Unregister the plugin while the vRPAs are detached. Use this procedure also to unregister a RecoverPoint extension from the Managed Object Browser.

ⓘ **NOTE:** Access to the Managed Object Browser is disabled by default in vSphere 6.0. For instructions on how to enable access, refer to VMware KB2108405.

Steps

1. In a web browser, enter the fully-qualified domain name (or IP address) of the ESXi or vCenter Server system:

   `https://<hostname.yourcompany.com>/mob/?moid=ExtensionManager`

2. Log in using your vCenter login credentials.

3. In the **Methods** table, select **UnregisterExtension**.
   A new browser window opens with **void UnregisterExtension** command displayed.

4. In the **Parameters** table, enter `com.dell.recoverpoint.vc.h5plugin` (HTML5 value), or `com.emc.recoverpoint.vwc` (FLEX value), and click **Invoke Method**.

Removing unused directories

From the vCenter, remove RecoverPoint for VMs Flex plugin directories that are no longer being used.

Steps

1. Connect to the vCenter server (using a local network mapping or Remote Desktop Connection). Delete the following folder:
   - vCenter 6.5/6.7 and Windows vCenter: `C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\vc-packages\vsphere-client-serenity\com.emc.recoverpoint.vwc-<version>`

- vCenter 6.5/6.7 and vCSA: `/etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity/com.emc.recoverpoint.vwc-<version>`

2. Restart the vSphere Web Client. For instructions, refer to VMware KB1003895 (Windows), VMware KB2109887 (Linux - vCenter 6.x), or VMware KB2147152 (Linux - vCenter 6.5).

Uninstall the RecoverPoint for VMs splitters

Use the ESXCLI to uninstall the RecoverPoint for VMs splitters.

Steps

1. Use ESXCLI to obtain a list of all installed vSphere Installation Bundles (VIBs):
   `esxcli software vib list`
2. Ensure that a bundle named *RP-Splitter* is installed.
3. On the ESXi host, enter maintenance mode:
   `esxcli system maintenanceMode set -e=true`

   ⓘ **NOTE:** For VSAN environments, this command requires an additional switch (refer to the vSphere documentation for the vSphere version that you are using).

4. To uninstall the splitter, type:
   `esxcli software vib remove -n "RP-Splitter"`

   ⓘ **NOTE:** If you want to remove the JAM VIB also (while still in maintenance mode), run the command:

   `esxcli software vib remove –n emcjiraf`

5. On the ESXi host, exit maintenance mode:
   `esxcli system maintenanceMode set -e=false`

Uninstall JAM

Use this procedure to uninstall the JAM VIB from an ESXi host.

Steps

1. On the ESXi host, vMotion all VMs to another ESXi host.
2. At the ESXCLI, enter maintenance mode. From the ESXi host console, use SSH to run the command:
   `esxcli system maintenanceMode set -e=true`

   ⓘ **NOTE:** For VSAN environments, this command requires an additional switch. Refer to the vSphere documentation for the vSphere version that you are using.

3. To remove the old JAM installation, run the command:
   `esxcli software vib remove –n emcjiraf`

   ⓘ **NOTE:** If you want to remove the splitter VIB also (while still in maintenance mode), run the command:

   `esxcli software vib remove -n "RP-Splitter"`

4. On the ESXi host, exit maintenance mode:
   `esxcli system maintenanceMode set -e=false`

C

vSphere upgrades

You may be required to upgrade a vCenter or an ESXi host that is used in the RecoverPoint for VMs system.

Information about these tasks helps you to successfully perform these upgrades.

Topics:

- Upgrading vCenter
- Upgrading an ESXi host

Upgrading vCenter

About this task

The vCenter upgrade is transparent to RecoverPoint for VMs provided that the upgrade process does not cause a change in the vCenter UUID.

Use this procedure when you need to upgrade the vCenter within a RecoverPoint for VMs system.

During the upgrade:

- vRPA clusters cannot be managed from this vCenter. Ensure you have access to other vCenters.
- Data replication and recovery point objective (RPO) might be affected.
- vCenters in Enhanced Linked Mode might be impacted (one vCenter at a time).
- RecoverPoint for VMs plug-in should remain intact.

⚠ **CAUTION: To avoid changing the vCenter UUID during the upgrade process, ensure that you select the Use existing inventory option.**

Steps

1. If you are upgrading to vCenter 6.5, follow these best practices: Upgrading to vCenter Server 6.5 best practices.
2. If you are upgrading to vCenter 6.7, follow these best practices: Upgrading to vCenter Server 6.7 best practices.

   When upgrading to vCenter Server 6.7 U1, or later (from vCenter Server earlier than 6.7 U1), deploy and configure a plugin server for each vCenter.

Results

ⓘ **NOTE:** If RecoverPoint for VMs is in an error state after you upgrade the vCenter, check if the vCenter UUID has changed. If it has, contact Customer Support.

Upgrading an ESXi host

Perform this procedure on every ESXi host that belongs to an ESX cluster running vRPA VMs.

Prerequisites

Before upgrading to ESXi 7.0U2 or later, ensure your vRPA clusters and splitter and JAM VIBs are all running RP4VMs 5.3 SP2 or later versions.

About this task

Steps

1. At the ESXCLI, enter maintenance mode. From the ESXi host console, use SSH to run the following command:

```
esxcli system maintenanceMode set -e=true
```

ⓘ **NOTE:** For VSAN environments, this command requires an additional switch (refer to the vSphere documentation for the vSphere version that you are using).

2. (Optional) If DRS is in automatic mode, vMotion is carried out automatically. If not, you must manually use vMotion to move all VMs to a different ESXi host in the ESX cluster.
3. If you are upgrading from an ESX version prior to 7.0, Uninstall the RecoverPoint for VMs splitters on page 45 and Uninstall JAM on page 46 now.
4. Follow the VMware instructions for ESXi upgrade.

Upgrading to 6.7	https://docs.vmware.com/en/VMware-vSphere/6.7/vsphere-esxi-67-upgrade-guide.pdf	
Upgrading to 7.0	https://kb.vmware.com/s/article/78205	

5. Exit maintenance mode. From the ESXi host console, use SSH to run the following command:
```
esxcli system maintenanceMode set -e=false
```
   The ESXi is updated with a new splitter and JAM VIBs.
6. (Optional) If DRS is not in automatic mode, use vMotion to manually move VMs to the upgraded ESXi host.

Results

The ESXi is upgraded.

Next steps

1. Repeat this procedure on every ESXi that has secure boot disabled and belongs to an ESX cluster hosting vRPAs.
2. Enable secure boot on every ESXi host in every ESX cluster hosting vRPAs.

D

Installing on Nutanix

There are no special instructions for installing RecoverPoint for VMs on Nutanix.

Topics:

- Installing RecoverPoint for VMs on Nutanix

Installing RecoverPoint for VMs on Nutanix

Perform a normal installation when installing RecoverPoint for VMs on Nutanix.

Prerequisites

RecoverPoint for VMs is supported by Nutanix AOS 4.7.5 or 5.0.2, or later. If necessary, upgrade your Nutanix installation to one of those versions.

(i) **NOTE:** You do not need to modify the Nutanix settings when installing RecoverPoint for VMs.