

RecoverPoint for Virtual Machines

Basic Configuration Installation Guide

P/N 302-004-369

Rev 01

Version 5.1.1

Copyright © 2017 Dell Inc. or its subsidiaries. All rights reserved. Published in USA.

Published October, 2017

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

The information in this publication is provided as is. Dell makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any Dell software described in this publication requires an applicable software license.

Dell, EMC, and other trademarks are registered trademarks of Dell Inc. or its subsidiaries.

Published in the USA.

Dell EMC

Hopkinton, Massachusetts 01748-9103

1-508-435-1000 In North America 1-866-464-7381

www.DellEMC.com

Contents

Purpose of the basic configuration installation.....	4
Terminology	4
Reference architecture.....	5
Assumptions.....	7
Before you begin	7
Overview of the basic configuration installation flow	9
Data preparation	10
Create VMkernel ports	11
Deploy vRPA OVA	15
Configure vRPA cluster.....	21
Verify deployment success	27
Deploy and configure remote site	27
Connect Site A and Site B vRPA clusters.....	28
Register ESXi cluster for replication	30
Protect virtual machine.....	32
Verify Replication Status	38
Final RecoverPoint for VMs installation steps.....	40
Appendix A: Install vSphere Client Integration Plug-in	41

Purpose of the basic configuration installation

Thank you for choosing RecoverPoint for Virtual Machines, a virtualized solution that provides data replication, protection, and recovery within the VMware vSphere environment. This guide provides a step-by-step explanation for installing a basic configuration to replicate virtual machines between two sites. This installation is targeted at low-scale environments with minimal redundancy and can be used to show proof of concept.

Using this guide, you create RecoverPoint for VMs clusters and connect them to form a working replication system.

For more advanced installations, refer to the *RecoverPoint for Virtual Machines Installation and Deployment Guide*. For larger scale configuration information, refer to the *RecoverPoint for Virtual Machines Scale and Performance Guide*. These guides may be found at <https://support.emc.com/>.

Terminology

If you are new to RecoverPoint or VMware architectures, review these definitions before continuing.

vRPA

The virtual RecoverPoint Appliance is a data appliance that manages data replication. You create the vRPAs you need by using the vSphere web Client from the vCenter Server.

vRPA cluster

A group of up to 8 vRPAs in the same cluster that work together to replicate and protect data. In this installation process, you create 2 vRPA clusters with 2 vRPAs in each cluster.

RecoverPoint for VMs plug-in

The vSphere web client user interface for managing VM replication. The plug-in is installed automatically after the vRPA cluster has been installed.

RecoverPoint for VMs splitter

Proprietary software that is installed on every ESXi host in an ESXi cluster that is involved in RecoverPoint for VMs replication or running virtual RPAs. The splitter splits every write to the VMDK and sends a copy of the write to the vRPA and then to the designated storage volumes. It is installed automatically after you register the ESXi cluster.

RecoverPoint for VMs system

One or more connected vRPA clusters. Use this document to configure a RecoverPoint for VMs system that connects two vRPA clusters, one at each of two sites.

vRPA LAN network

vRPA cluster management network that is used for communication between vRPAs and to other entities such as the vCenter.

vRPA WAN network

vRPA inter-cluster network that is used for communication between vRPA clusters. The basic configuration uses a unified topology: a single WAN+LAN network. This topology does not require an additional gateway to communicate between sites.

vRPA data (IP) network

Data network that is used for traffic between the vRPA and the ESXi splitter.

Reference architecture

For ease of deployment and proof of concept, this guide uses a reference architecture that features a basic configuration that includes:

- Two vRPA clusters, one at Site A and one at Site B
- Two vRPAs at each vRPA cluster
- Two vCenters
- Two ESXi hosts on each vCenter
- Two vNICs per vRPA, one for LAN + WAN and the other for Data (IP) communication
- One VMkernel port for each ESXi

Note

Other network topologies can be deployed. For more information, refer to the *RecoverPoint for Virtual Machines Installation and Deployment Guide*.

The goal of this guide is to establish replication by deploying a single consistency group between Site A and Site B.

The steps include creating two vRPAs in a vRPA cluster at each site and connecting these two vRPA clusters to enable replication between the sites.

The examples in this guide are based on the reference architecture in [Reference architecture for the basic configuration](#).

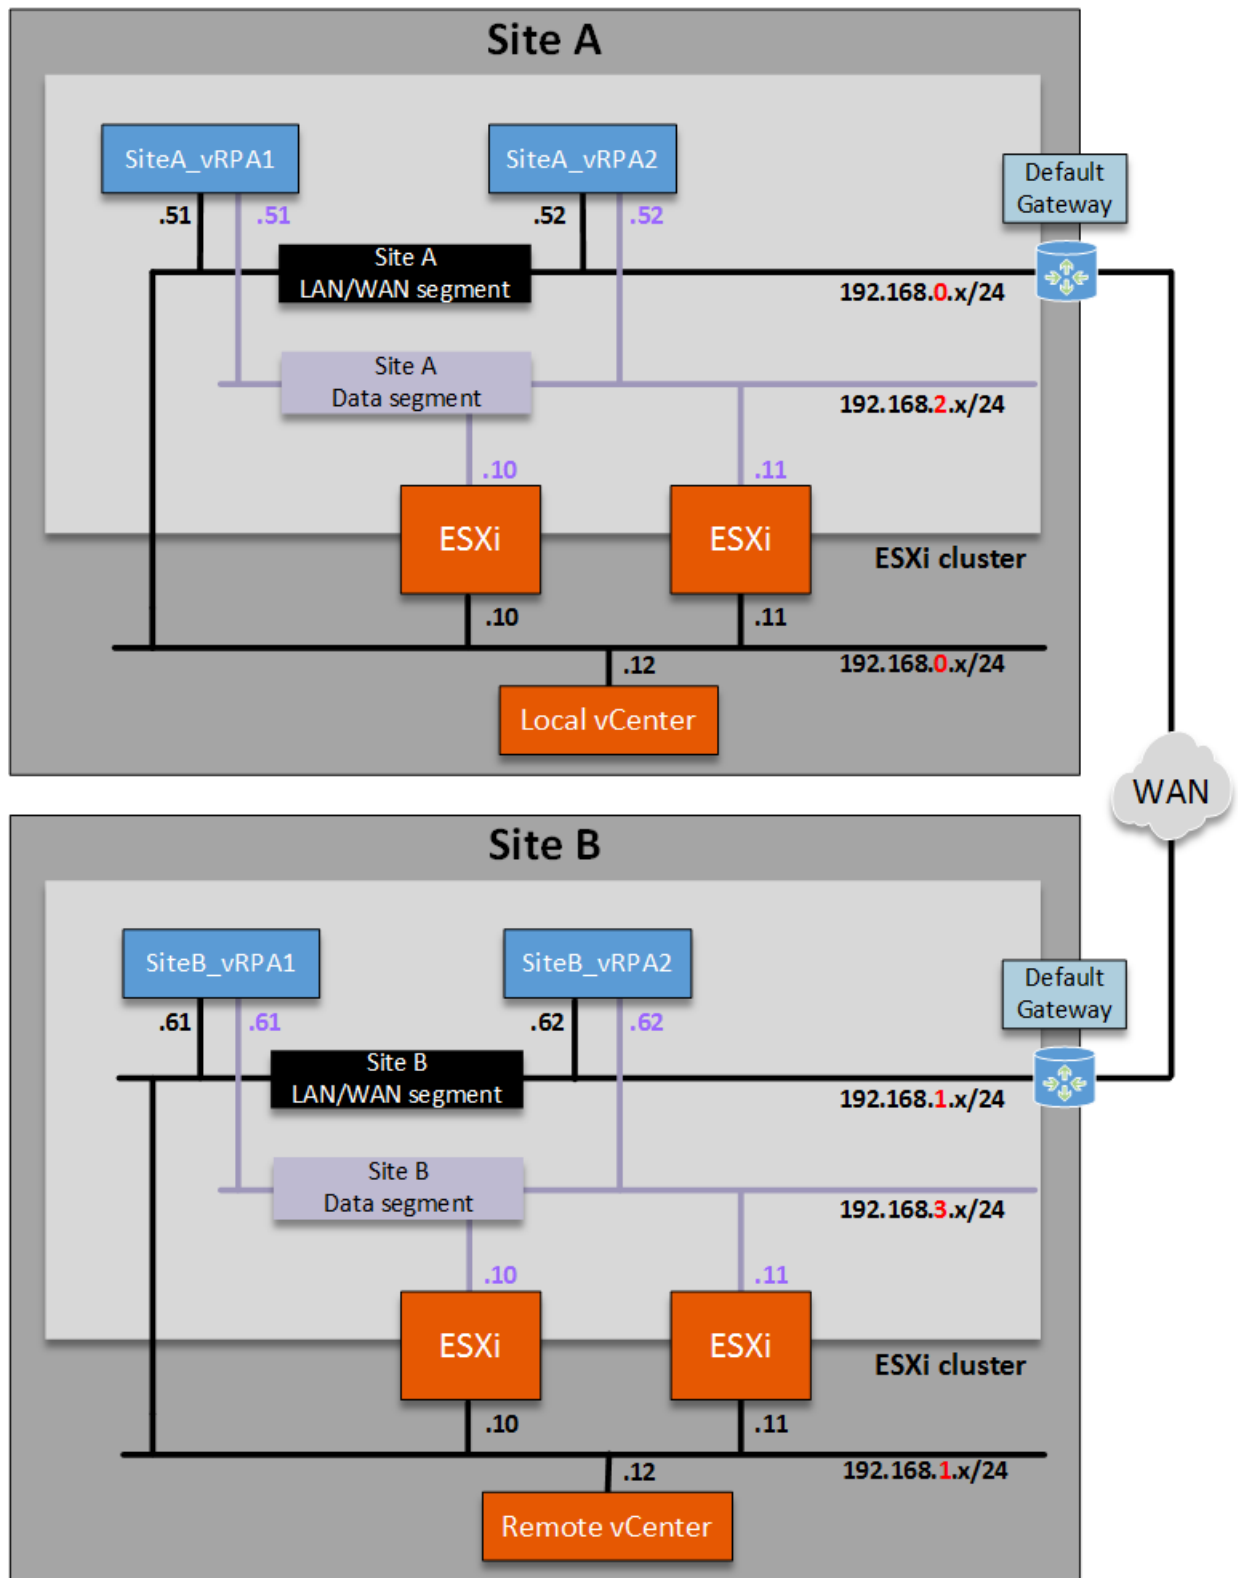


Figure 1 Reference architecture for the basic configuration

Assumptions

This procedure provides a limited installation of RecoverPoint for proof of concept. The procedure assumes various environmental characteristics.

- You have installed the latest version of Google Chrome, Firefox, or Internet Explorer browser on the workstation that you use for this installation.
- Screenshots in this document are taken from a vSphere 6 Update 2 environment. Screens may appear slightly different when using a different version.
- RecoverPoint for VMs Release 5.1.1 supports the vSphere versions that are listed in the Simple Support Matrix, which is available at <https://elabnavigator.emc.com> and contains the most up-to-date information.
- This guide assumes IP communication between the ESXi hosts and the vRPAs. The IP communication mode requires vSphere 6.0 and later.
- vSphere Web Client is available on both vCenters.

Before you begin

Before you begin this installation ensure that you prepare for the installation.

- Download the vRPA OVA file for RecoverPoint for VMs from <https://support.emc.com/downloads>.
 - a. Perform a search in the **Type a Product Name** text box for **RecoverPoint for Virtual Machines**.
 - b. Locate and download **RecoverPoint for Virtual Machines 5.1.1 Installation Kit**. Ensure that you select version 5.1.1.

Example of downloaded file:
`RecoverPoint_for_Virtual_Machines_5.1.1_Installation_Kit_md5_checksum.zip`
 - c. Extract the `.zip` file that contains the OVA file that is needed for the installation.
- Select an IP segment from which IPs are allocated for ESXi hosts and vRPAs (recommended on a dedicated VLAN).
- Select a network segment— any standard VLAN with three unassigned IPs for each cluster. This network segment is used for management of vRPAs, vRPA cluster management, and data transfer between the clusters.
- For each vRPA, open firewall ports for inbound and outbound TCP/IP communication. Ports 5020, 5040, 5042, 5044, and 5050 are required for communication between the vRPAs and the ESXi splitter. For more detailed information about these and other ports, refer to the *RecoverPoint for Virtual Machines Security Configuration Guide*.
- If you have a network firewall that is installed in the environment, open firewall ports 443 and 7225 between vCenter and vRPA.
- Ensure that you have at least one VM for protection in Site A.

- Ensure that you have the following datastore free space:
 - Site A
 - 100 GB for system usage (not including protected VM)
 - 10 GB per protected VM, for journaling

These datastores must be shared between the ESXi hosts within the ESXi cluster that hosts the vRPAs.
 - Site B
 - 100 GB for system usage
 - Size of Site A protected VMs
 - 10 GB per protected VM, for journaling

These datastores must be shared between the ESXi hosts within the ESXi cluster that hosts the vRPAs.
- Plan for 8 GB memory and 8000 MHz CPU in each ESXi cluster for each site. This requirement assumes using a Bronze vRPA (2 vCPU and 4GB). For larger deployments, use the *RecoverPoint for Virtual Machines Installation and Deployment Guide*.
- Ensure that vSphere Web Client with Client Integration Plug-in is installed (see high-level steps in Appendix A: Install vSphere Client Integration Plug-in).
- Ensure that all server clocks, including ESXi hosts and vCenters, are synchronized using NTP. Time differences may affect VMware vCenter certificate authentication resulting in possible installation failure.
- A persistent scratch location on each ESXi host is required for storing splitter configuration information. The scratch location (/scratch/log) requires at least 50 MB of free storage space on a permanently available persistent storage device.
- Related RecoverPoint for VMs guides are available at <https://support.emc.com/>:
 - *RecoverPoint for Virtual Machines Scale and Performance Guide*
 - *RecoverPoint for Virtual Machines Security Configuration Guide*
 - *RecoverPoint for Virtual Machines Installation and Deployment Guide*

Overview of the basic configuration installation flow

The basic installation flow that is described in this guide includes these high-level steps:

Table 1 Installation flow

Order of installation flow	Section of this guide
Gather the necessary data.	Data preparation
Create VMkernel ports.	Create VMkernel ports
Deploy the OVA for Site A vRPA 1, and repeat for Site A vRPA 2.	Deploy vRPA OVA
Configure the vRPA cluster for Site A and verify successful deployment.	<ul style="list-style-type: none">• Configure vRPA cluster• Verify deployment success
Deploy the OVA and configure the vRPA cluster for the remote site, Site B.	Deploy vRPA OVA <ul style="list-style-type: none">• Configure vRPA cluster• Verify deployment success
Connect the vRPA clusters between Site A and Site B.	Connect Site A and Site B vRPA clusters
Register the ESXi cluster for Site A, and repeat for Site B.	Register ESXi cluster for replication
Establish and verify replication between production copy at Site A and remote copy at Site B.	<ul style="list-style-type: none">• Protect virtual machine• Verify Replication Status
Register the RecoverPoint for VMs system	Final RecoverPoint for VMs installation steps

Data preparation

Before beginning this installation, use these forms to collect data that is required for the installation.

Table 2 Common to all vRPAs

Property	Site A	Site B	Notes
vCenter IP address			
Cluster name			
Datastore name			System usage - use of several datastores is optional.
Time zone			
Local domain (optional)			
Primary DNS server IP (optional)			
Secondary DNS Server IP (Optional)			
NTP server (recommended)			
vRPA cluster management IP			
vCenter credentials			
WAN + LAN/subnet mask			
IP subnet mask			
ESXi1 VMKernel IP			
ESXi2 VMKernel IP			

Table 3 For each vRPA

Property	SiteA_vRPA1	SiteA_vRPA2	SiteB_vRPA1	SiteB_vRPA2	Notes
WAN + LAN IP					
Data IP					

Create VMkernel ports


The basic configuration for RecoverPoint for VMs uses IP protocol for communication between the ESXi-based splitter and the vRPAs. To enable data transfer, you must create a VMkernel port.

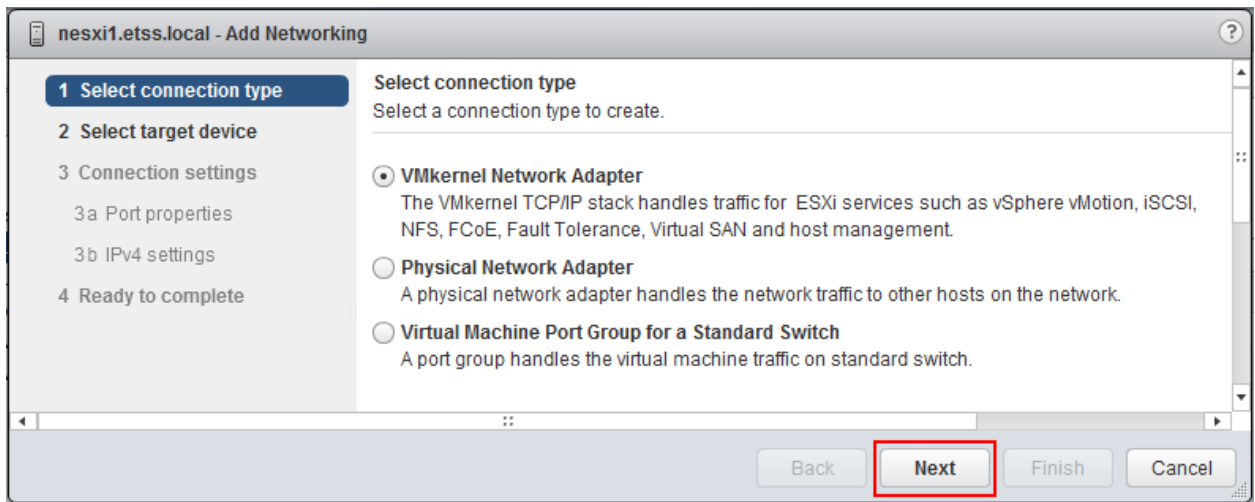
About this task

Perform these steps on each ESXi host, from each vCenter.

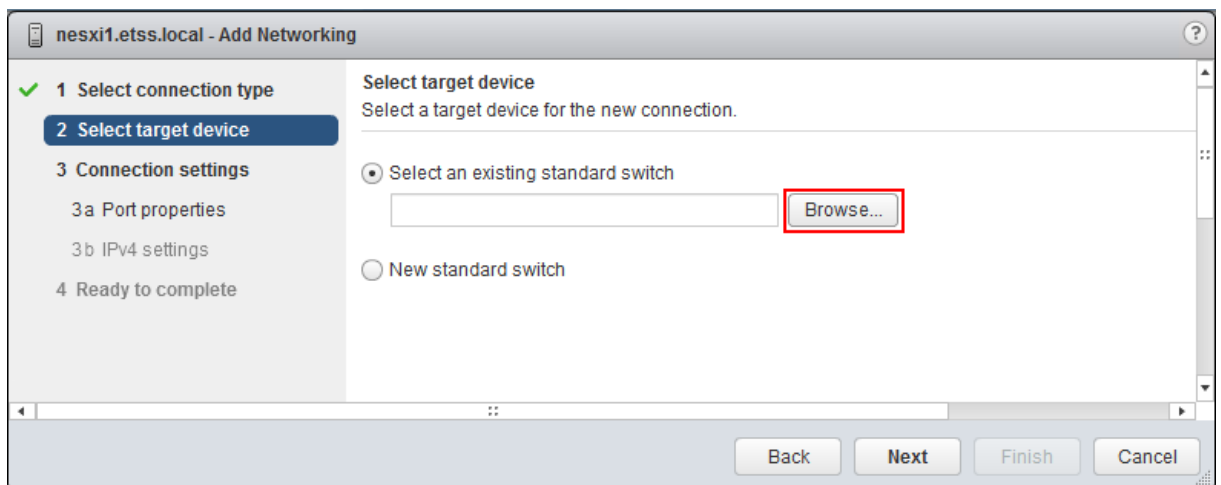
Procedure

1. Connect to the vSphere Web Client.
2. Create a VMKernel port:

Select **Hosts and Clusters** > *host* > **Manage** > **Networking** > **VMkernel adapters** and click .




3. Do one of the following:
 - a. To select an existing vSwitch, click **Browse**.

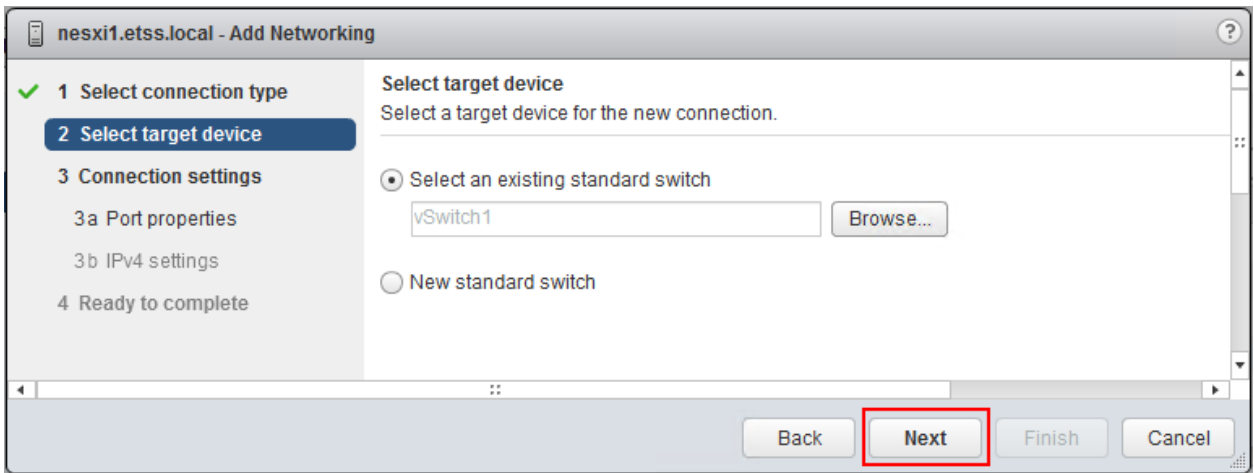


- b. To create a vSwitch select **New standard switch**.

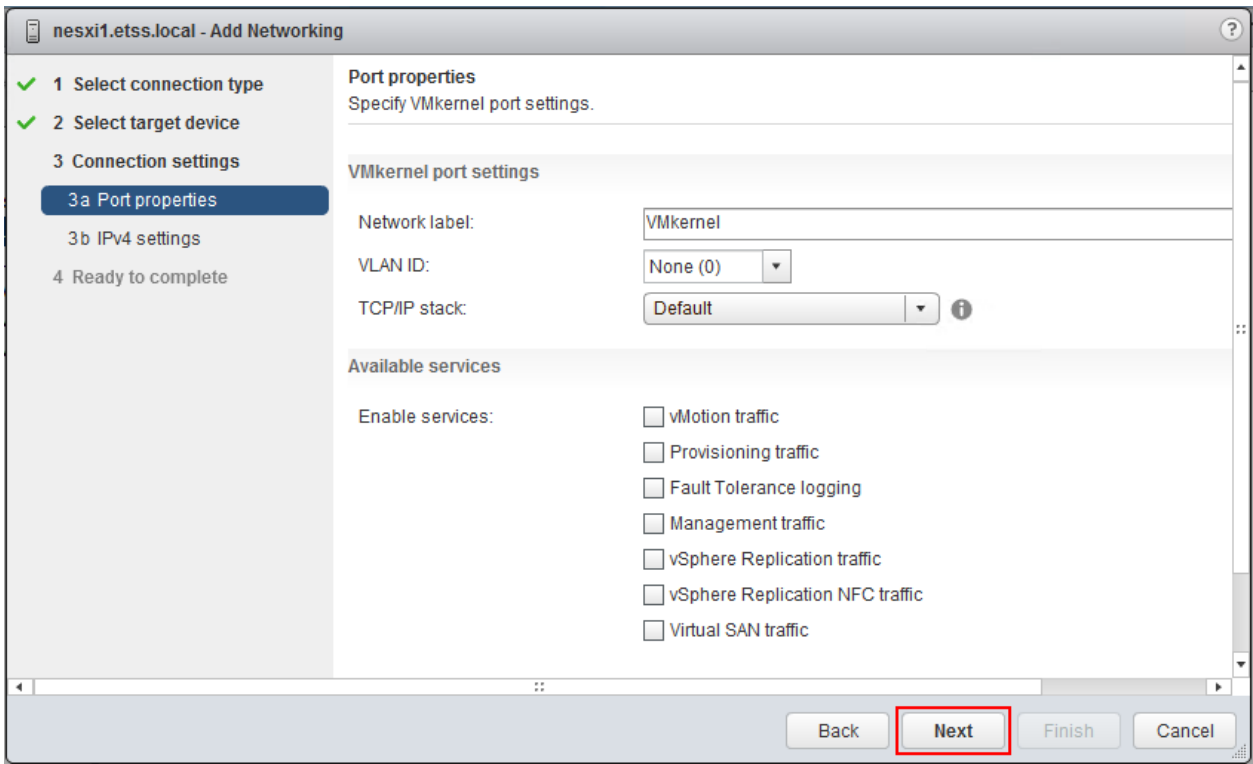
4. Select a vSwitch from the list. Click **OK**.



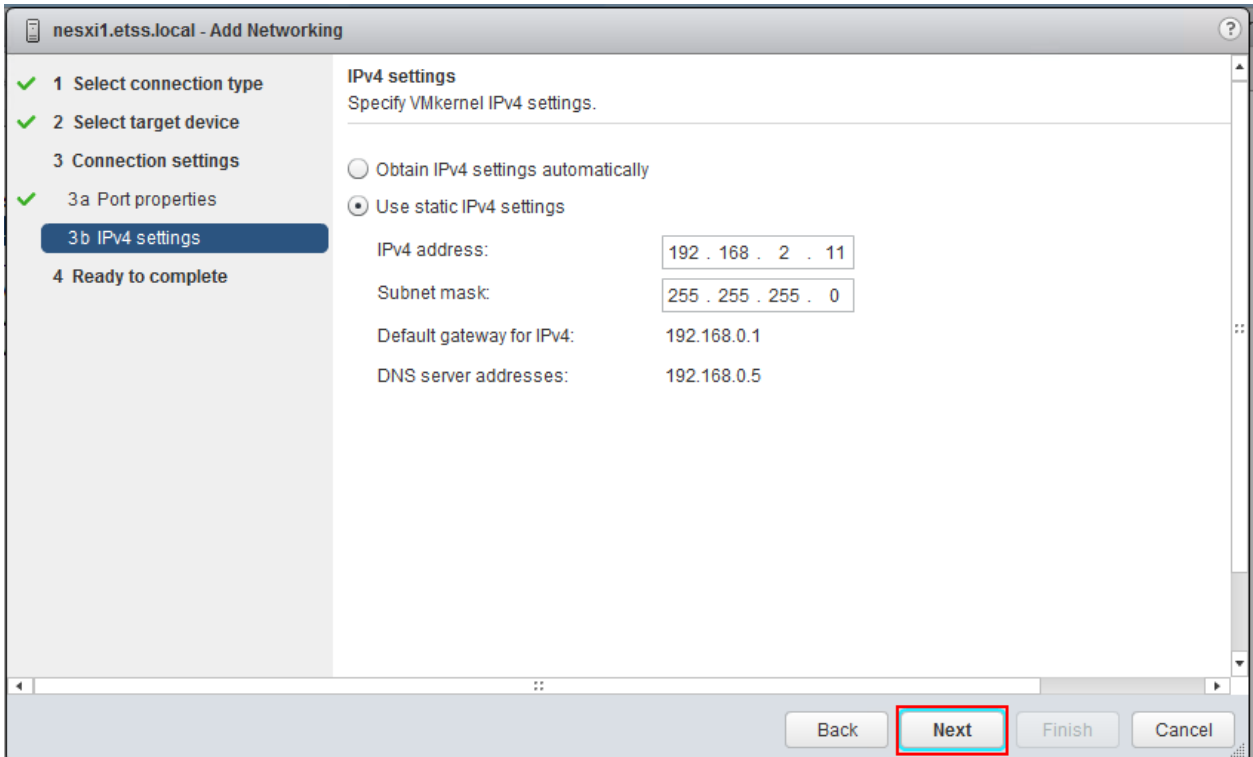
5. Click **Next**.



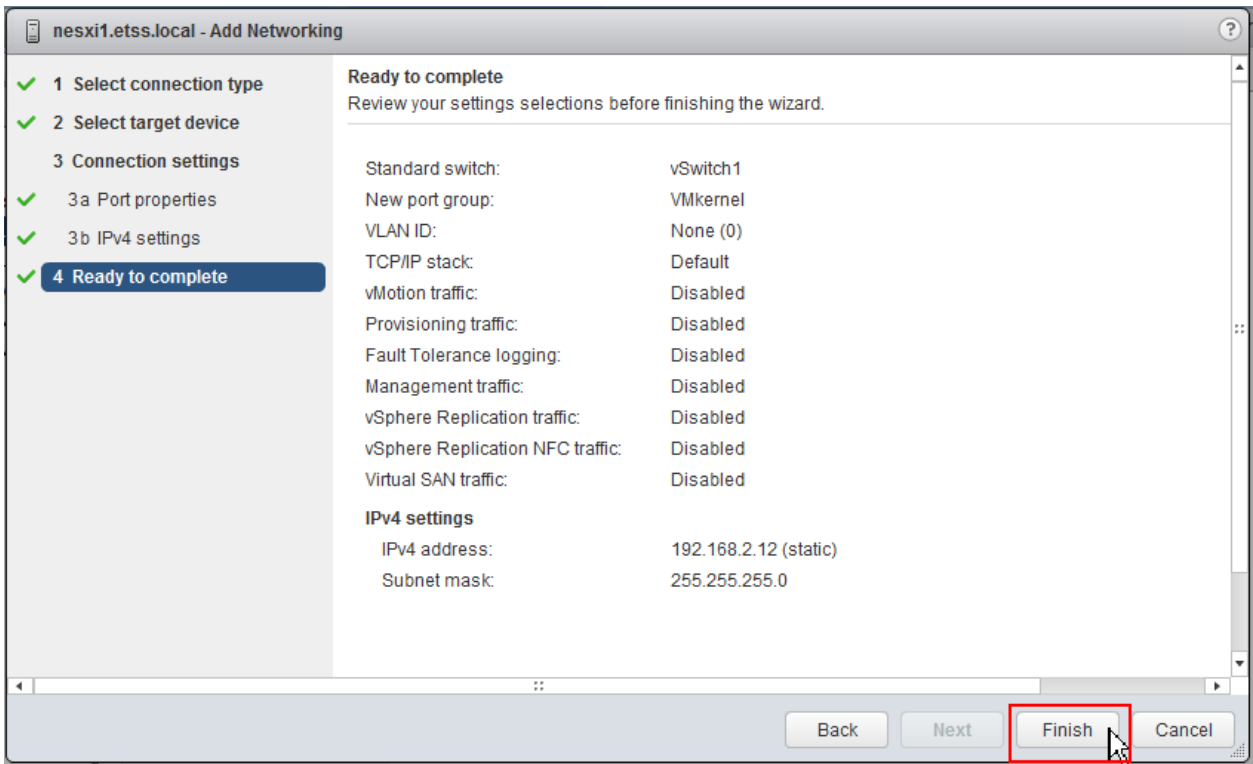
6. Type a **Network label** for the new VMkernel port. Click **Next**.



7. Type IP settings for the VMkernel port. Use the data that you collected in "[Data preparation.](#)" Click **Next**.



8. In the Ready to complete window, check that the information is correct and click **Finish**.



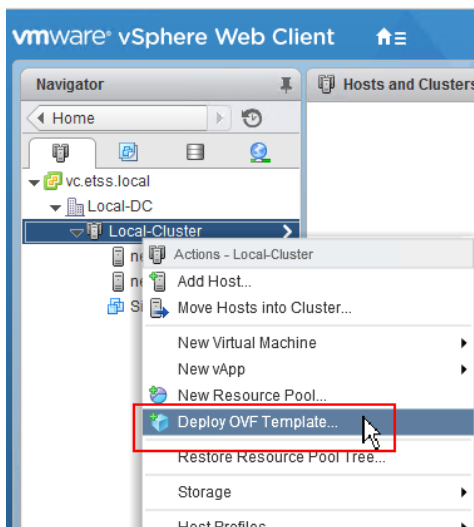
9. Repeat the steps in this procedure for each ESXi host, from each vCenter.

Deploy vRPA OVA

Deploy the OVA on two virtual Machines at Site A, then power on the vRPAs.

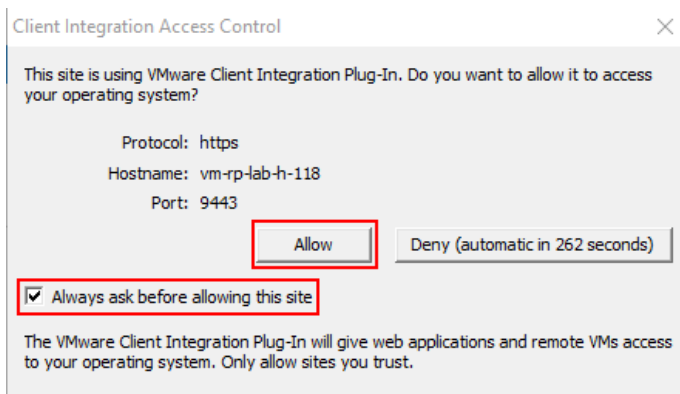
Procedure

1. Connect to the vSphere Web Client.
2. In the Navigator, click **Hosts and Clusters**.
3. Select the ESXi cluster or specific ESXi host on which to deploy the vRPA VM.
4. To deploy the OVA, right-click the cluster or host, and select **Deploy OVF template**.



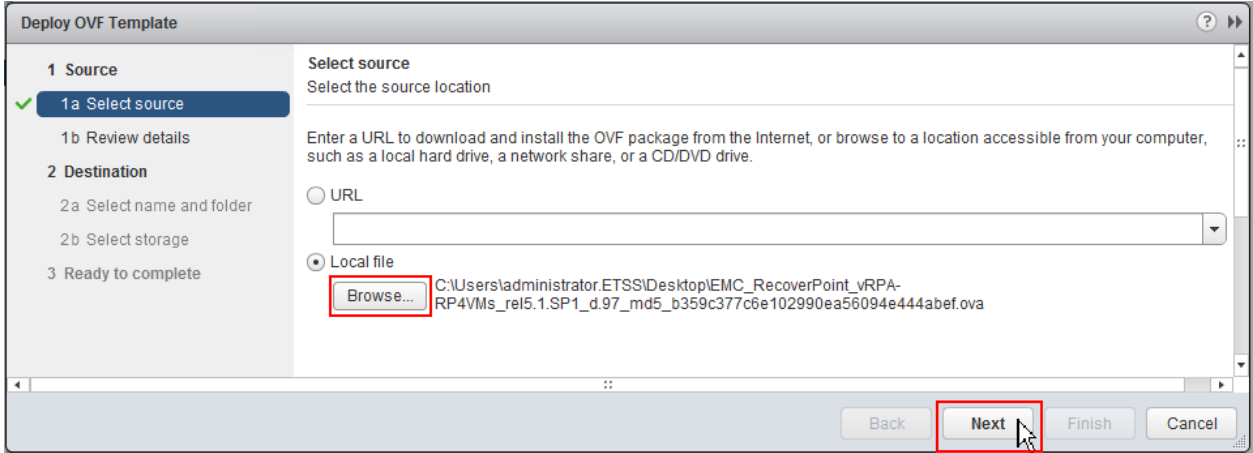
The vRPA reserves the virtual machine resources to ensure that the targeted environment has enough resources.

5. After selecting **Deploy OVF template**, you might receive the following pop-up. Clear the **Always ask before allowing this site** checkbox, and then click **Allow**.

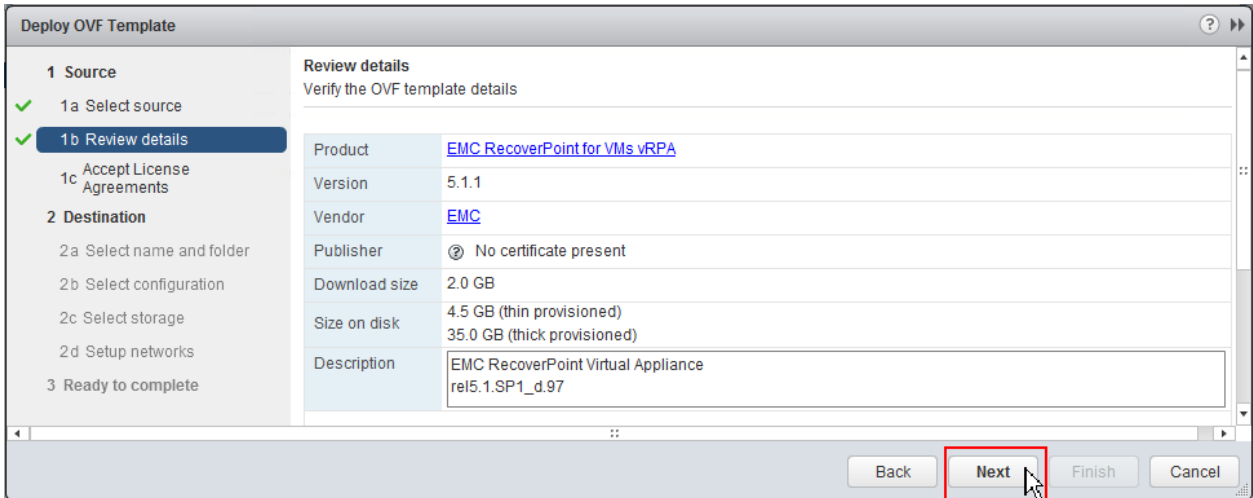


6. In the Select Source step of the Deploy OVF Template wizard:

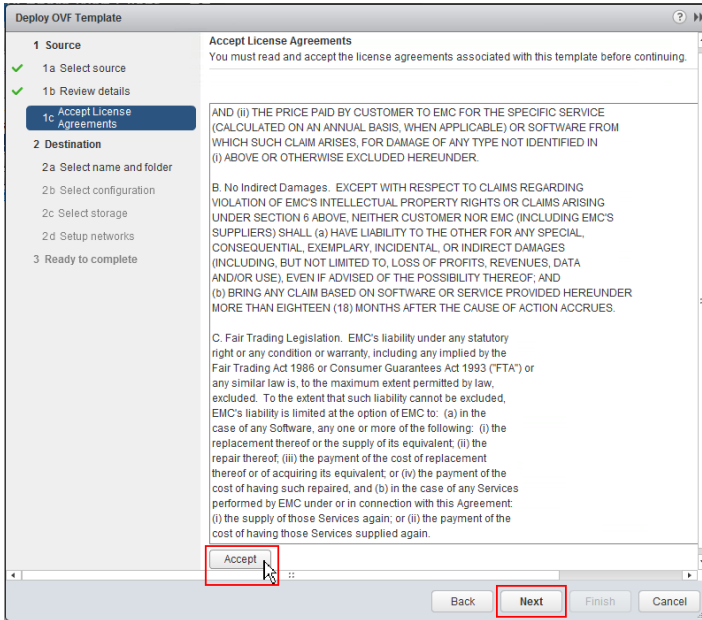
- a. Click **Browse**.
- b. Locate the OVA file, and then click **OK**.
- c. Click **Next**.



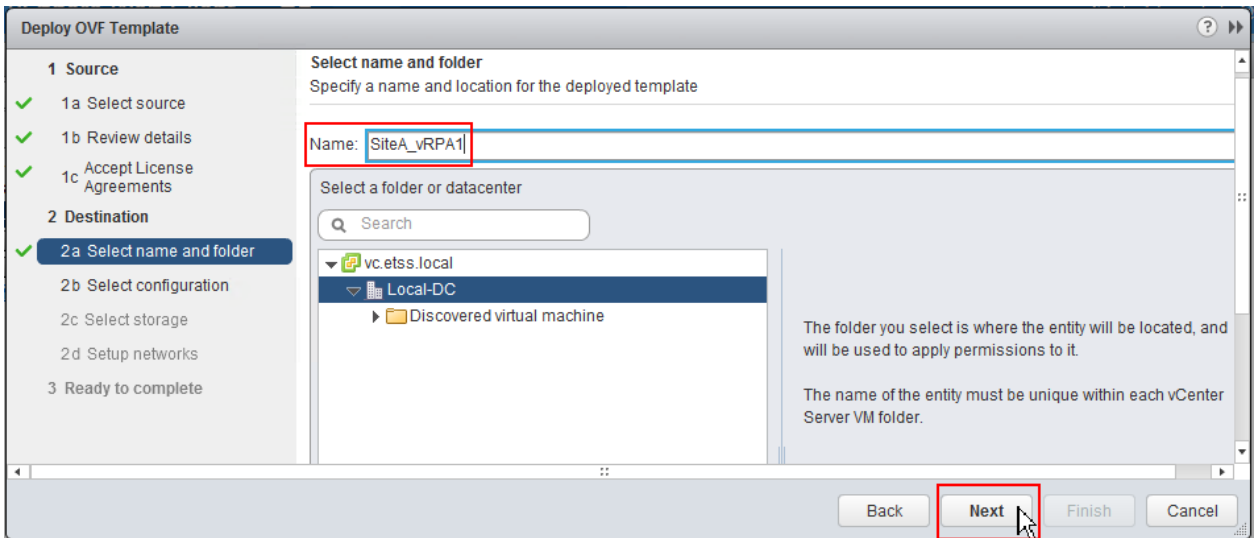
7. In the Review details step, click **Next**.



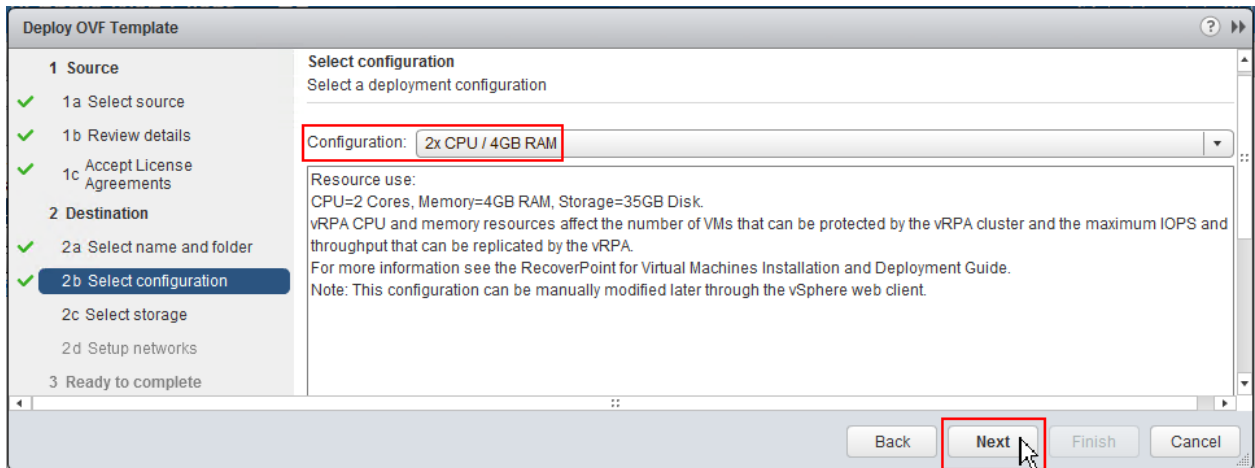
8. Scroll to the bottom of the **Accept License Agreements** window. If you accept the terms, click **Accept**, and click **Next**.



9. Select a VM folder in which to locate the vRPA, and type a meaningful name. For example:

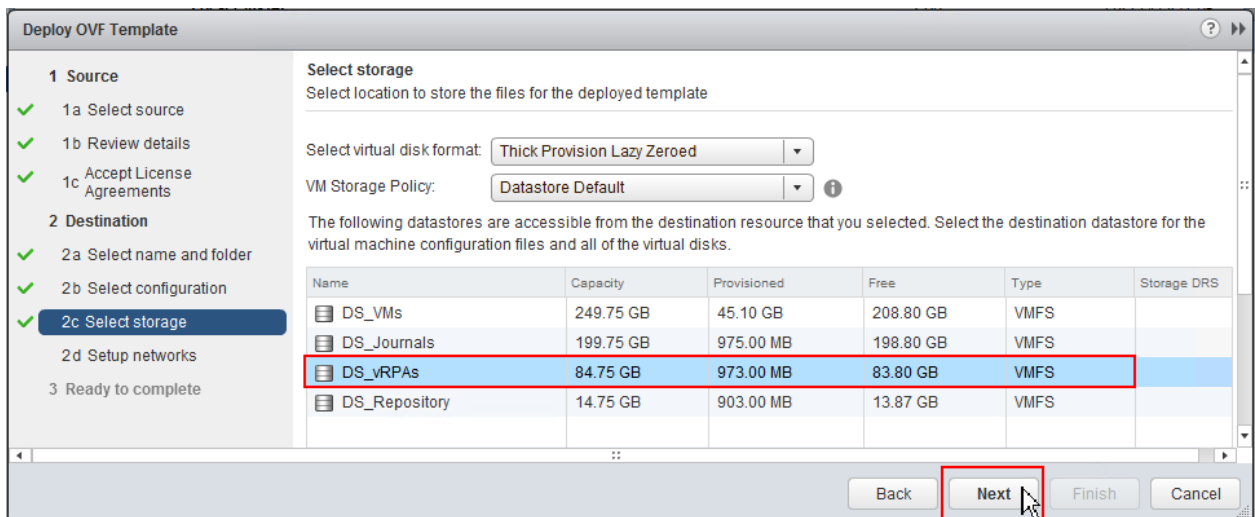


10. Select the desired deployment configuration. For this installation, the smallest configuration (as shown) is adequate. Click **Next**.

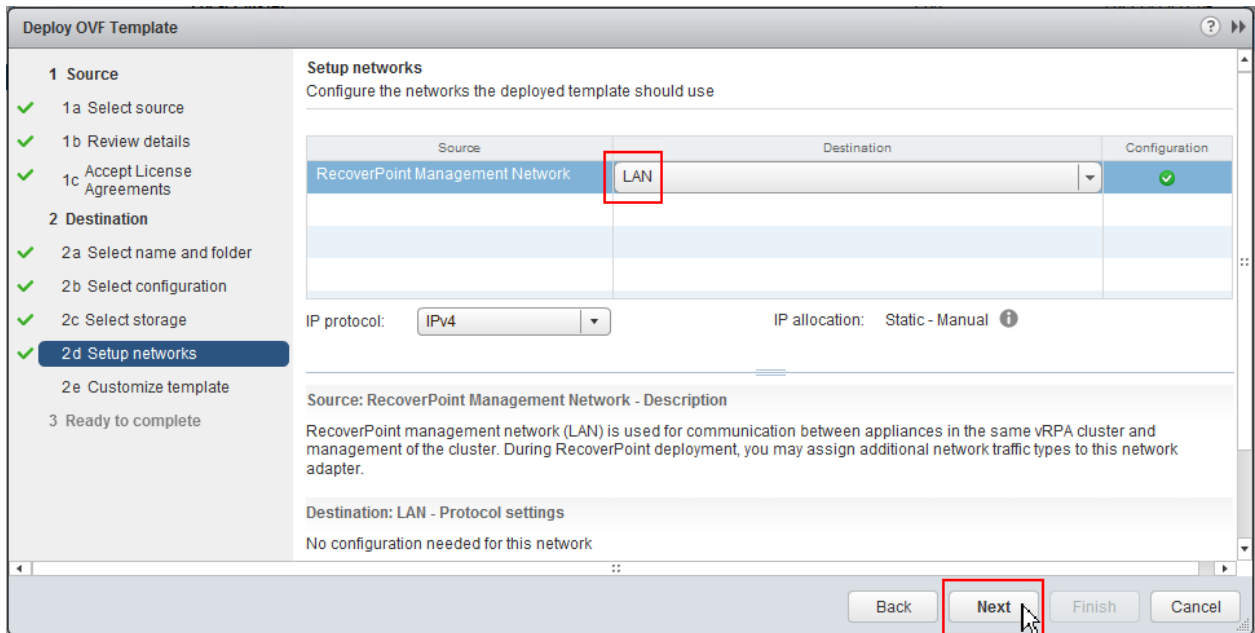


11. In the Select Storage step:

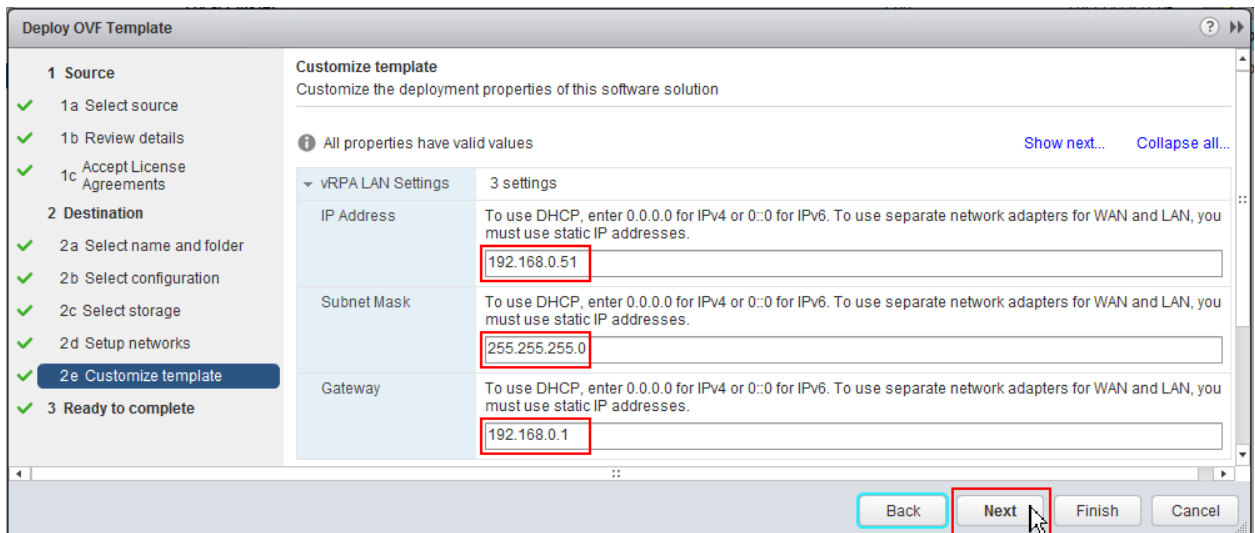
- Select a datastore (see ["Data Preparation"](#)) on which to deploy the vRPA.
- If you are using storage policies, select a policy that is based on the requirements. These requirements might include redundancy, speed, and back-up, among others.
- Click **Next**.



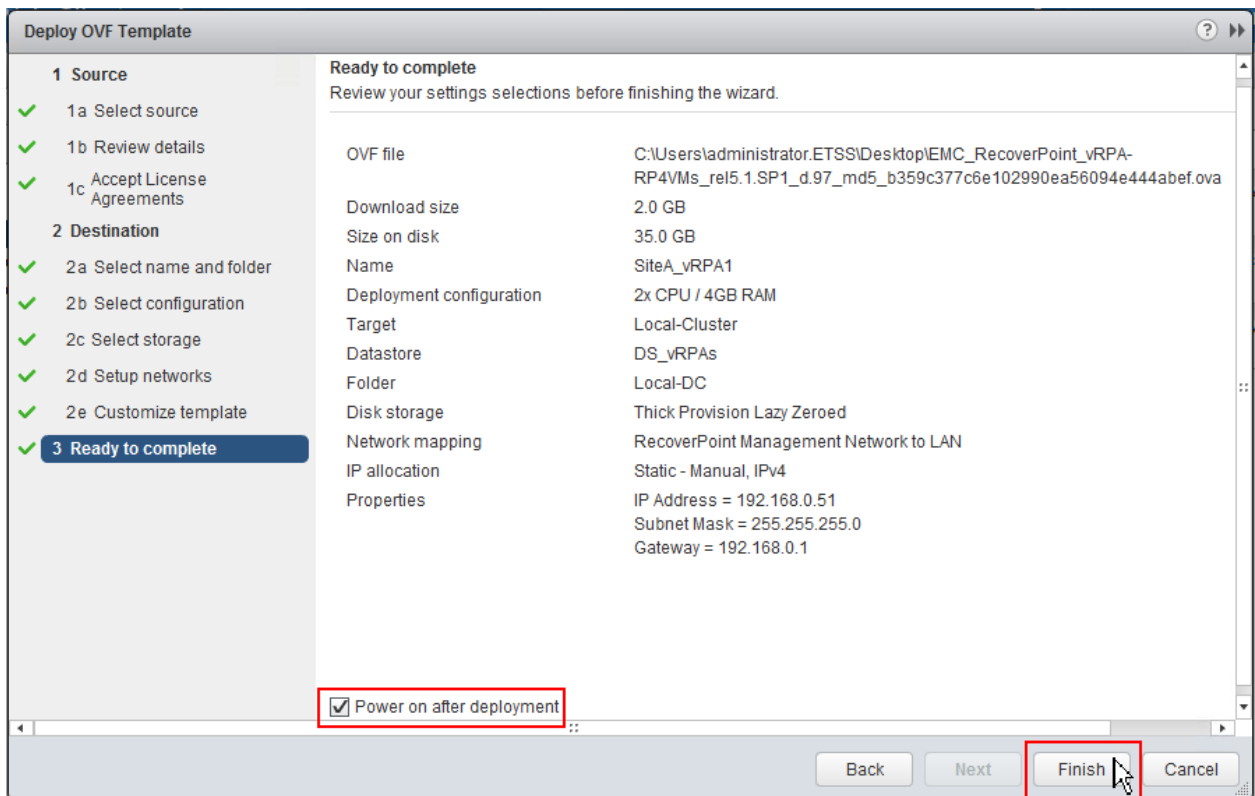
12. Select the port group for the vRPA management LAN network, also known as the LAN network. This network is used also for the WAN network (see reference architecture). Click **Next**.



13. Type the **IP address**, **Subnet Mask**, and **Gateway**. Click **Next**.



14. Review the summary, select the **Power on after deployment** checkbox, and click **Finish**.



15. Repeat steps 4-14 for SiteA_vRPA2.

16. When finished, ensure that both vRPAs are powered on. If not, power them on.

Results

You have finished deploying the vRPAs and are ready to configure the vRPA clusters.

Configure vRPA cluster

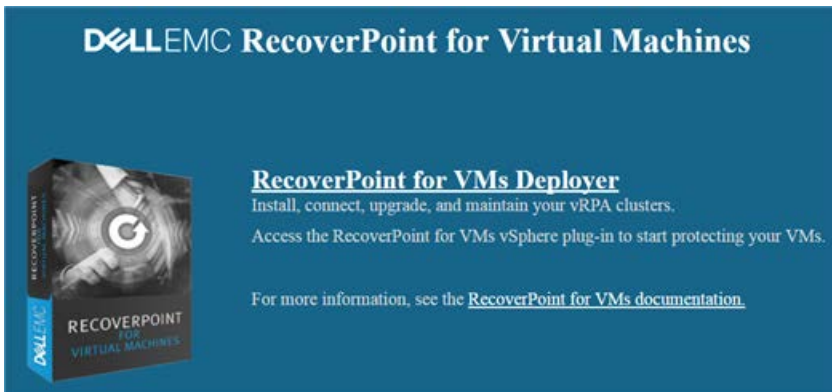
Use the RecoverPoint for VMs Deployer to configure both vRPAs of the cluster.

Procedure

1. Connect to the RecoverPoint for VMs Deployer by opening a web browser and entering the address of the first vRPA:

`https://first_vRPA_IP`

2. In the initial RecoverPoint for VMs Deployer window, click the **RecoverPoint for VMs Deployer** link.



3. At the bottom of the page, click **Install a vRPA cluster**.



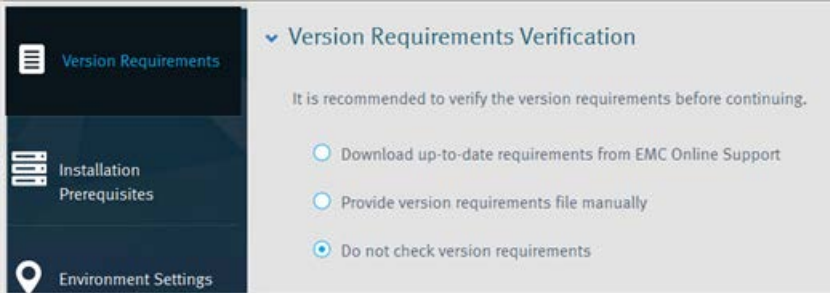
If possible, the Deployer automatically connects to the repository for the Version Requirements file.

4. If verification is successful, in the Version Requirements Verification window, click **Next**.

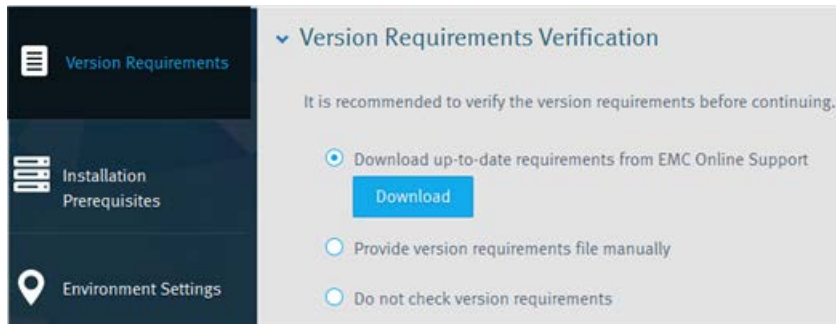


5. If prompted to verify version requirements, do one of the following:

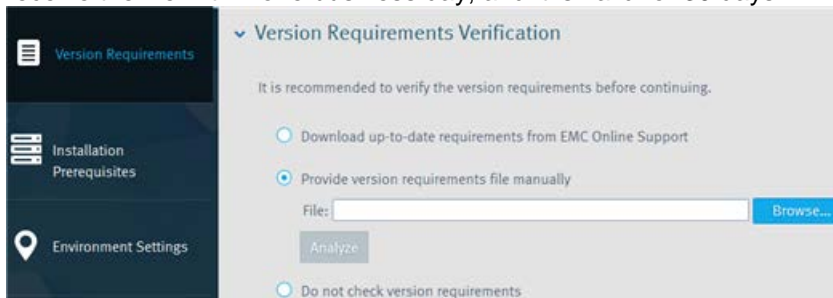
- If this installation is a proof-of-concept installation, select **Do not check version requirements**.



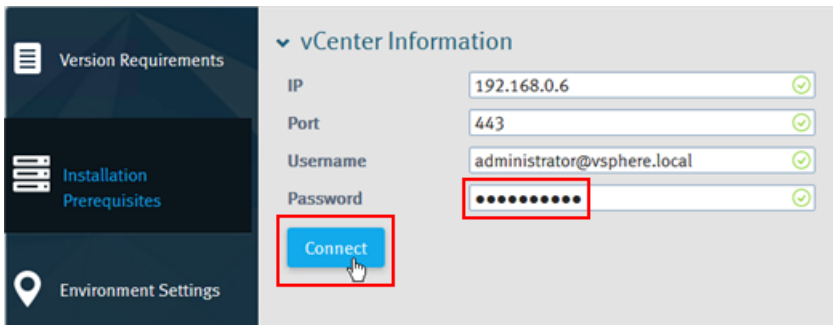
- If you are a licensed user with a direct connection to the Internet but you are encountering Internet connectivity issues, fix the issues and select **Download up-to-date requirements from the EMC Online Support**.



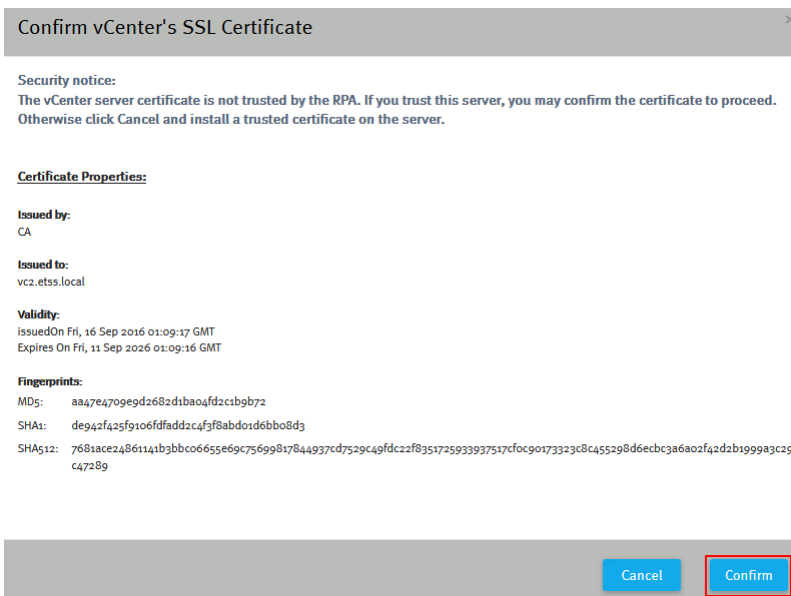
- If you are a licensed user without a direct connection to the Internet:
 1. Select **Provide version requirements file manually**.
 2. From an Internet-connected station, browse to <https://rplicense.emc.com/download/download.asp>.
 3. Download the a-cca.xml file or request that it be sent to you by email.
 4. If this option is not available, open a Service Request with Customer Support (severity level 3). Request the latest Upgrade Change Management XML file, indicating the installation date. You receive the file within one business day, and it is valid for 30 days.



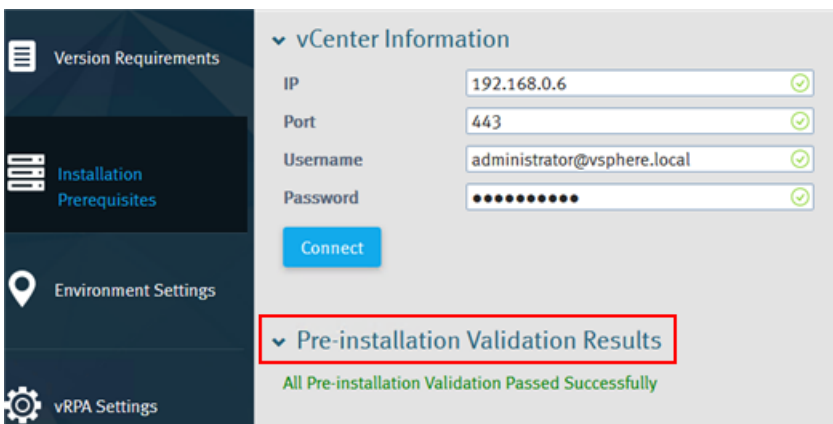
6. Type user credentials with administrator privileges, and connect to the vCenter by clicking **Connect**.



7. A pop-up may prompt you to confirm the SSL certificate. Accept the certificate by clicking **Confirm**.



8. Verify the **Pre-installation Validation Results**, and ensure that all results are positive (as shown). Click **Next**.

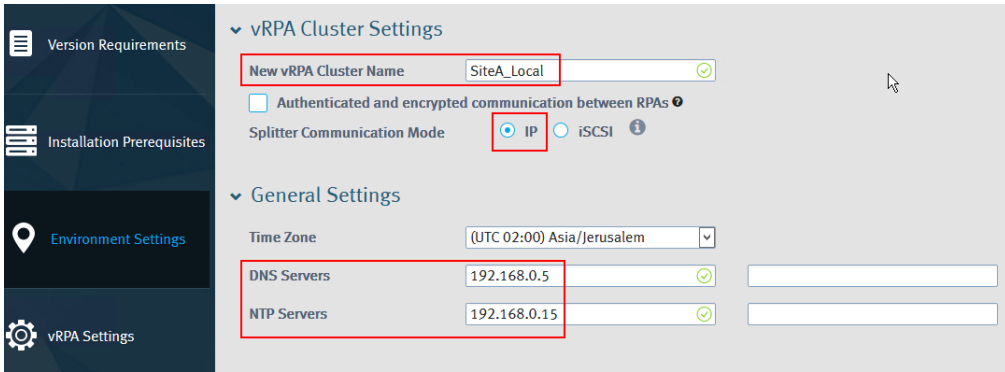


9. Configure the **Environment Settings** according to the data from "[Data preparation.](#)"
 - a. Type a name for the new vRPA cluster. For example: Site A_Local.
 - b. For this procedure, clear the **Authenticated and encrypted communication between RPAs** checkbox.

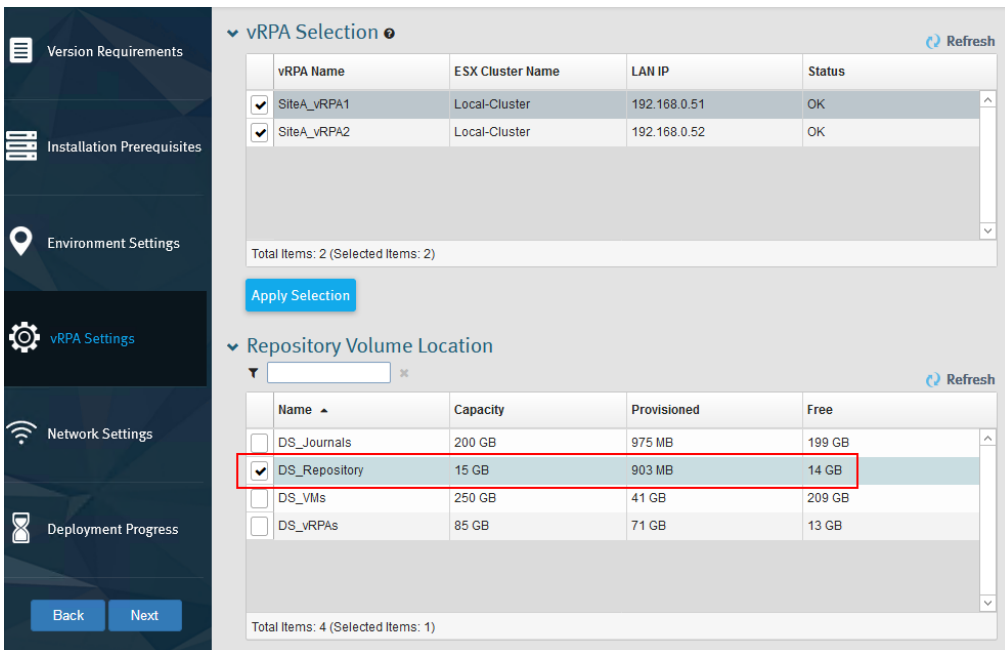
Note

For more information about communication between vRPAs, review the *RecoverPoint for Virtual Machines 5.1 Security Configuration Guide*.

- c. Ensure that IP splitter communication mode is selected (this mode is the default mode if verification checks pass).
- d. Type DNS and NTP server IP addresses. Click **Next**.



10. Configure **vRPA Settings** according to the data from "[Data preparation.](#)"
 - a. Select both vRPAs by clicking the checkboxes at the left of the vRPA names.
 - b. Click **Apply Selection**.
 - c. Select a datastore for the shared repository data between the vRPAs. Click **Next**.



This action creates a VMDK file under a folder that is named **RPvStorage**, which is used for shared persistent storage by the vRPAs.

11. Configure the **vRPA Cluster Topology** section according to the data from "[Data preparation.](#)"
 - a. Use the default network adapter topology, which places WAN and LAN on the same network adapter and Data on a separate network adapter.
 - b. Select the Data network (in this example, Replication_Network).

▼ vRPA Cluster Topology

▼ Network Adapters Configuration

WAN and LAN on same network adapter
Data (IP) on separate network adapter from WAN and LAN

Edit

▼ Network Mapping

WAN + LAN: LAN

Data: Replication_Network

IPv4 DHCP

IPv4 DHCP

12. Configure the **vRPA and Cluster Settings** section according to the data from "[Data preparation.](#)"
 - a. In the **WAN + LAN IP address** section, in the **Default Gateway** text box, type the default gateway IP address of the vRPA management network.
 - b. In the **Data IP address** section, in the **IPv4** text boxes, type the IP address of the vRPAs.
 - c. In the **Netmask** text box, type the netmask of the IP network.

▼ vRPA and Cluster Settings

Cluster Management: IPv4
192.168.0.50

▼ WAN + LAN IP address

IPv4

vRPA 1: SiteA_vRPA1: 192.168.0.51

vRPA 2: SiteA_vRPA2: 192.168.0.52

Netmask: 255.255.255.0

Default Gateway: 192.168.0.1

Additional Gateways: Add gateways for WAN connectivity to remote vRPA clusters.
Add

▼ Data IP address

IPv4

vRPA 1: SiteA_vRPA1: 192.168.2.51

vRPA 2: SiteA_vRPA2: 192.168.2.52

Netmask: 255.255.255.0

Note

The **IPv4** text boxes in the **WAN + LAN IP address** section are disabled because the IP address is automatically retrieved from vRPA properties.

- d. When finished typing the data, click **Install**.

13. Observe the installation progress as the settings are applied. The installation may take a while to complete.

Installing the new cluster

53% Installing vRPAs cluster

- Installing vRPAs cluster, and verifying all vRPAs are attached. 53%
- Selecting the correct installation mode. 66%
- Setting up and installing vRPA cluster. 55%
- Installing vRPAs cluster. 69%
- Performing system checks for installation. 100% ✓
- Validating installation input. 100% ✓
 - Validating installation basic cluster input. 100% ✓
 - Validating installation basic cluster input without advanced parameters. 100% ✓
 - Validating advanced parameters input. 100% ✓
 - Validating installation vRPA details input. 100% ✓
 - Validating vCenter and resource pool input. 100% ✓
 - Checking for duplicate IPs. 100% ✓
- Validating if all IPs are in same network. 100% ✓
 - Validating whether all LAN IPs of version IPv4 are in the same network. 100% ✓
 - Validating whether all LAN IPs of version IPv6 are in the same network. 100% ✓
 - Validating whether all WAN IPs of version IPv4 are in the same network. 100% ✓
 - Validating whether all WAN IPs of version IPv6 are in the same network. 100% ✓
- Validating if all IPs are in range. 100% ✓
 - Checking if IP is in range. 100% ✓

14. When the vRPA cluster configuration is complete, in the **Deployment Progress** window, click **Finish**.

Installation completed successfully

Access the RecoverPoint for VMs plugin and follow the [Next Steps](#).
Change the default passwords, as described in the *RecoverPoint for VMs Installation and Deployment Guide*.

[Back to the home page](#)

Results

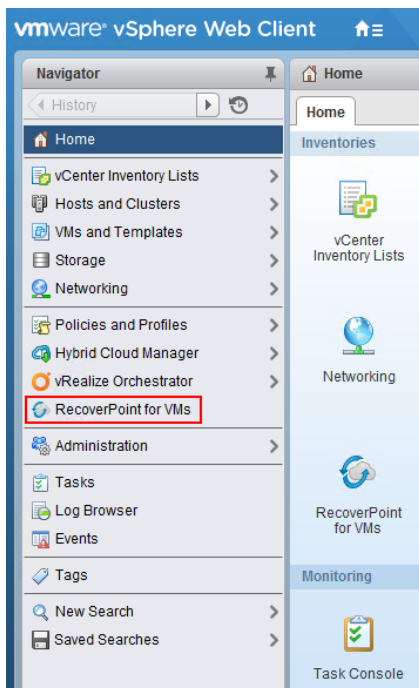
You have finished configuring the vRPA cluster and are ready to verify the deployment success.

Verify deployment success

After configuring the vRPAs, ensure that the deployment was successful by finding the RecoverPoint for VMs plug-in in the vSphere Web Client.

Procedure

1. If you are already logged in to vSphere Web Client, log out and log in again to force vSphere Web Client to refresh.
2. In the vSphere Web Client Navigator, ensure that **RecoverPoint for VMs** appears.



Results

You have finished verifying deployment success and are ready to deploy vRPAs and configure the cluster at the remote site.

Deploy and configure remote site

At the remote site, repeat all steps in the [“Deploy vRPA OVA”](#) and [“Configure vRPA cluster”](#) sections for SiteB_vRPA1 and SiteB_vRPA2. Use the data that you collected in [“Data preparation.”](#)

Results

The remote site is deployed and configured. Go to [“Connect Site A and Site B vRPA clusters”](#).

Connect Site A and Site B vRPA clusters

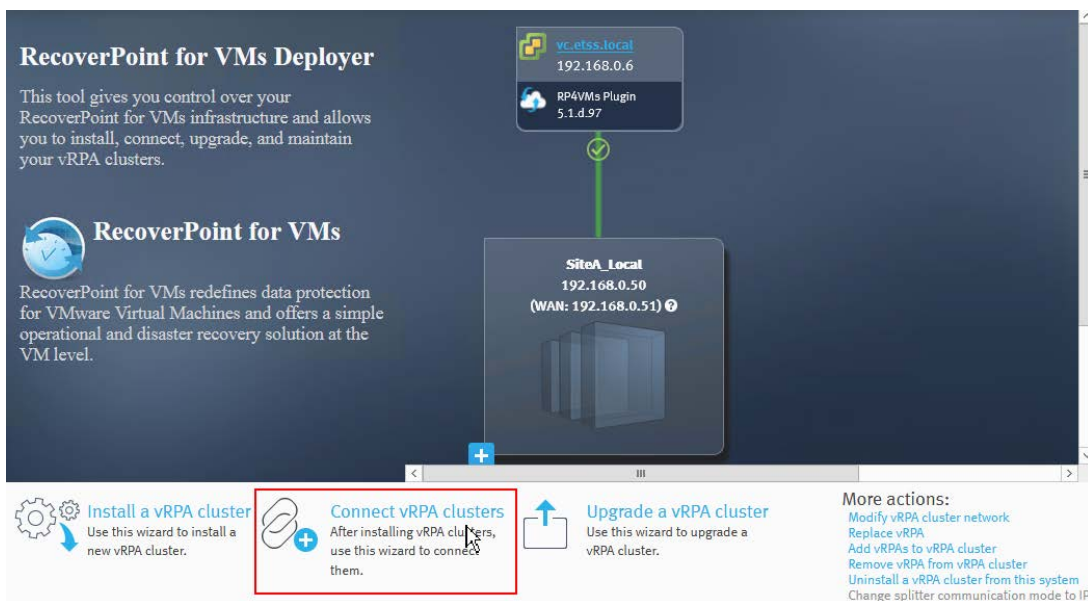
Use the RecoverPoint for VMs Deployer to connect the two sites so that RecoverPoint for VMs can replicate data from Site A to Site B.

Procedure

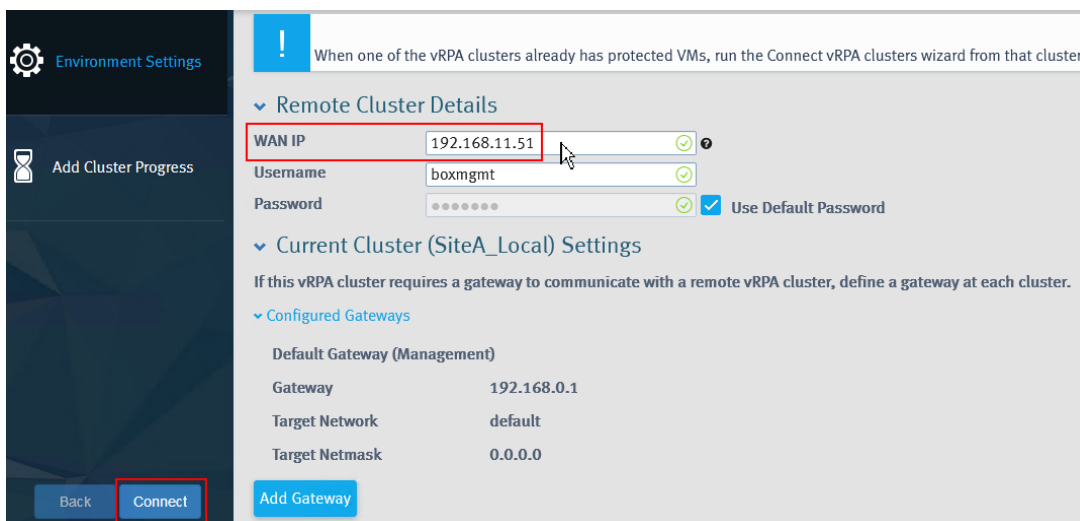
1. Connect to the RecoverPoint for VMs Deployer home page by browsing to `https://SiteA_cluster_management_IP`.

For example, `https://192.168.0.50`

2. In the RecoverPoint for VMs Deployer home page, click **Connect vRPA clusters**.

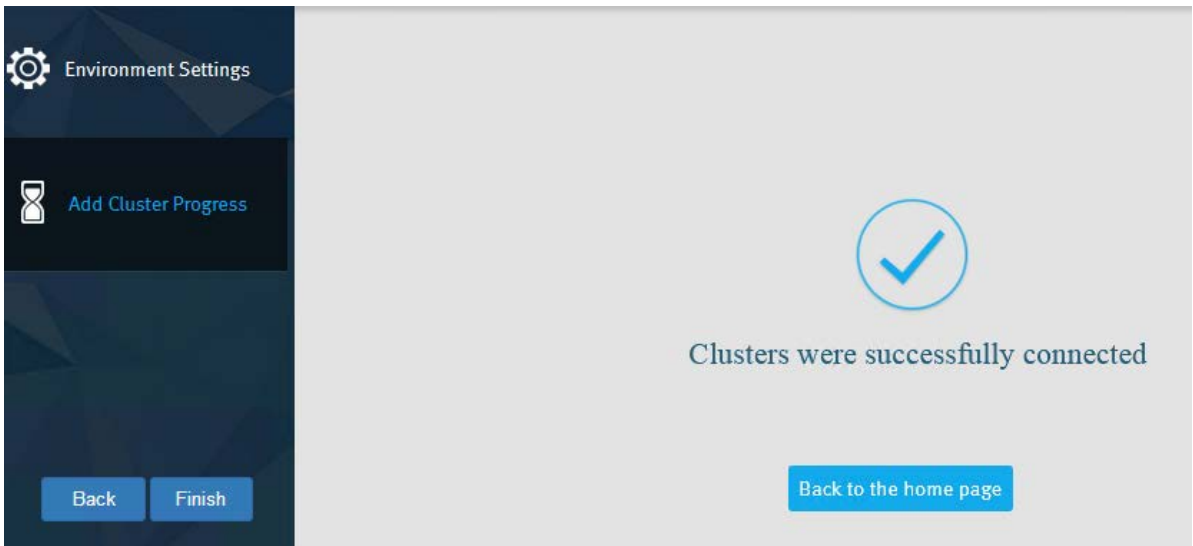


3. In the Remote Cluster Details section, type the WAN IP address of one of the vRPAs at the remote (Site B) cluster. Use the data that you collected in "[Data preparation.](#)" Then click **Connect**.



Observe the progress of the connection between the vRPA clusters.

4. Ensure that the clusters connect successfully. Then click **Finish**.



Results

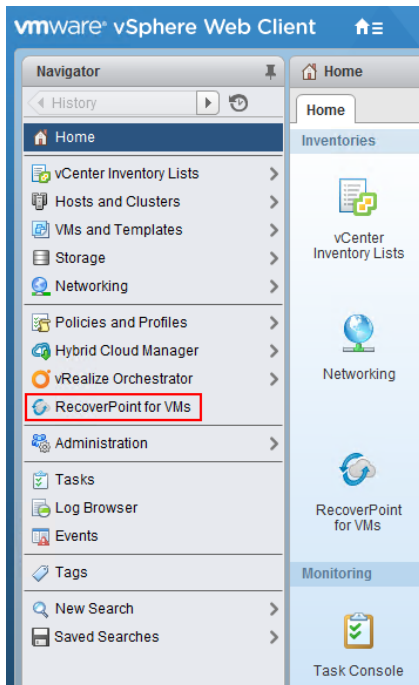
You have finished connecting Site A and Site B vRPA clusters and are ready register the ESXi cluster for replication.

Register ESXi cluster for replication

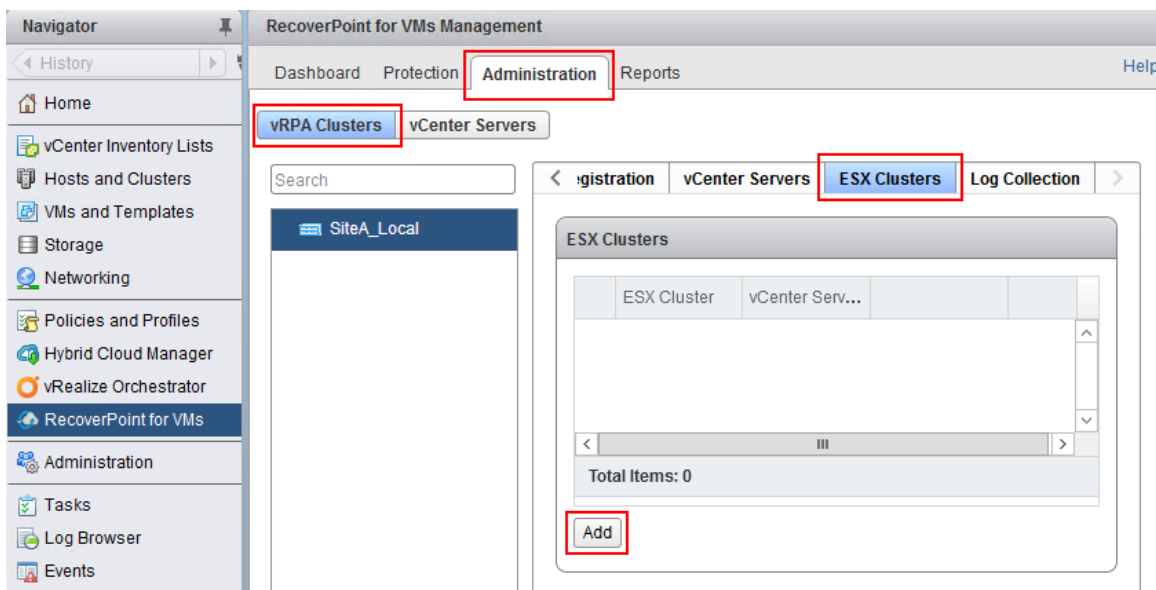
Replication occurs only for ESXi clusters that are registered for replication. Use the vSphere Web Client to register the ESXi clusters.

Procedure

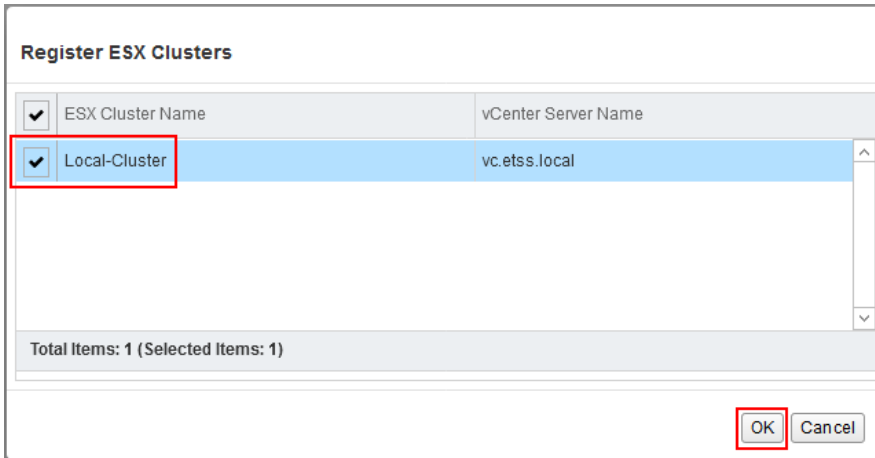
1. Connect to the vSphere Web Client of the local site (Site A).
2. Select **RecoverPoint for VMs** from the home page.



3. Select **Administration > vRPA Clusters > ESX Clusters**, and then click **Add**.

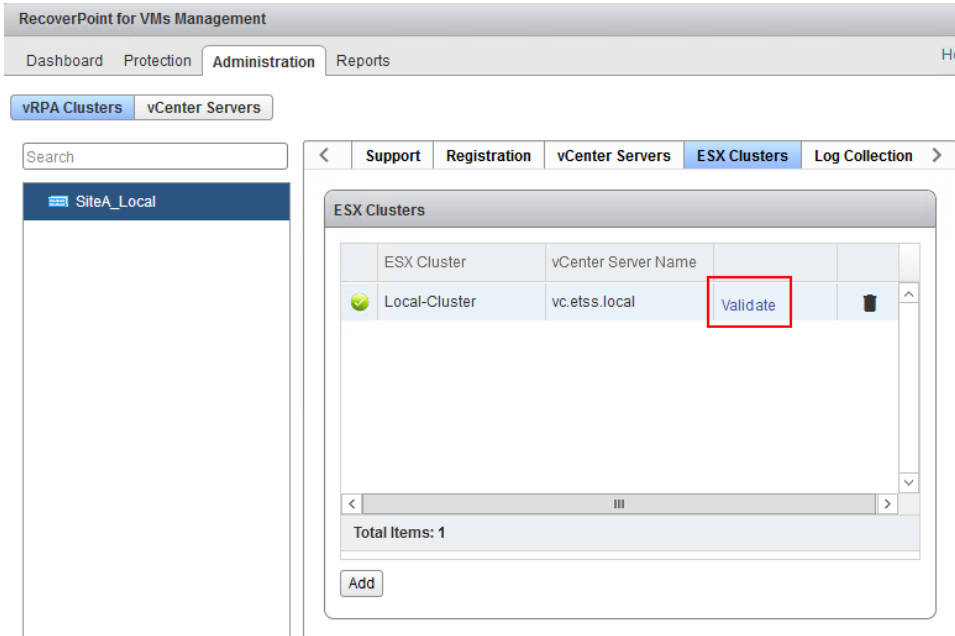


4. Select the ESX cluster from which to replicate virtual machines, and then click **OK**.



5. Wait 10 minutes for splitters to be auto-installed on the ESXi hosts and for the vRPA and splitters to initialize communication.

6. Click **Validate** to verify connectivity between the ESXi and the vRPA NICs.



7. Repeat the previous steps for Site B, selecting the remote ESXi cluster from the vSphere Web Client of Site B.

Results

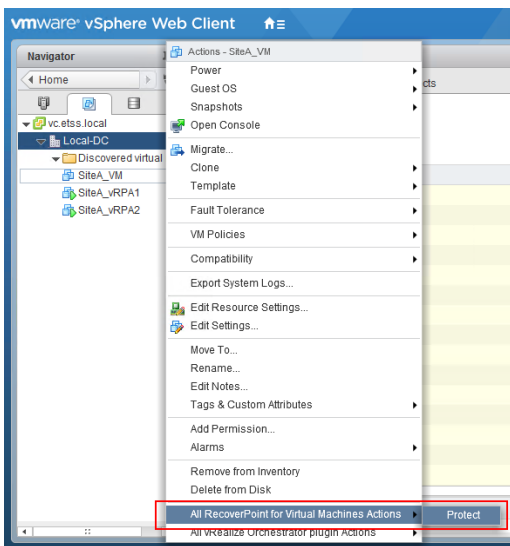
You have finished registering the ESXi cluster for replication and are ready to protect the virtual machine.

Protect virtual machine

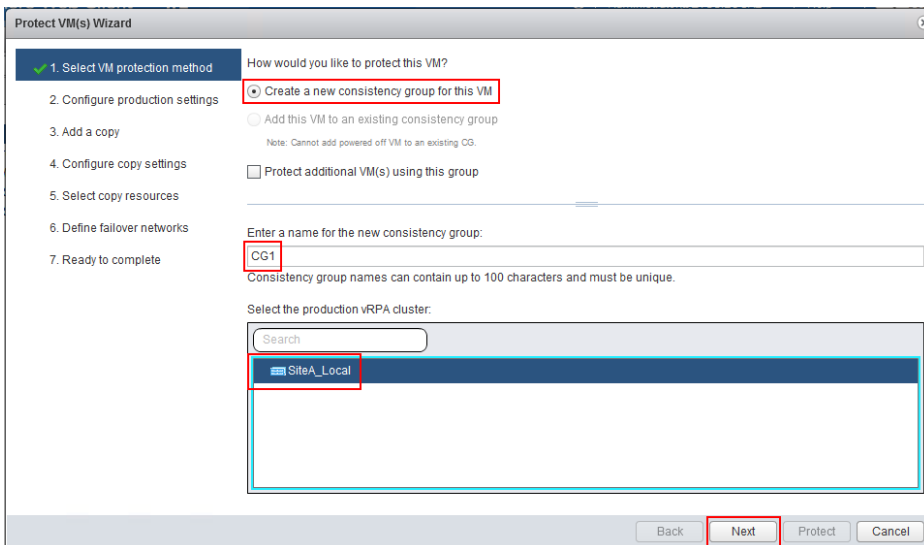
Production VMs can be protected using the vSphere Web Client.

Procedure

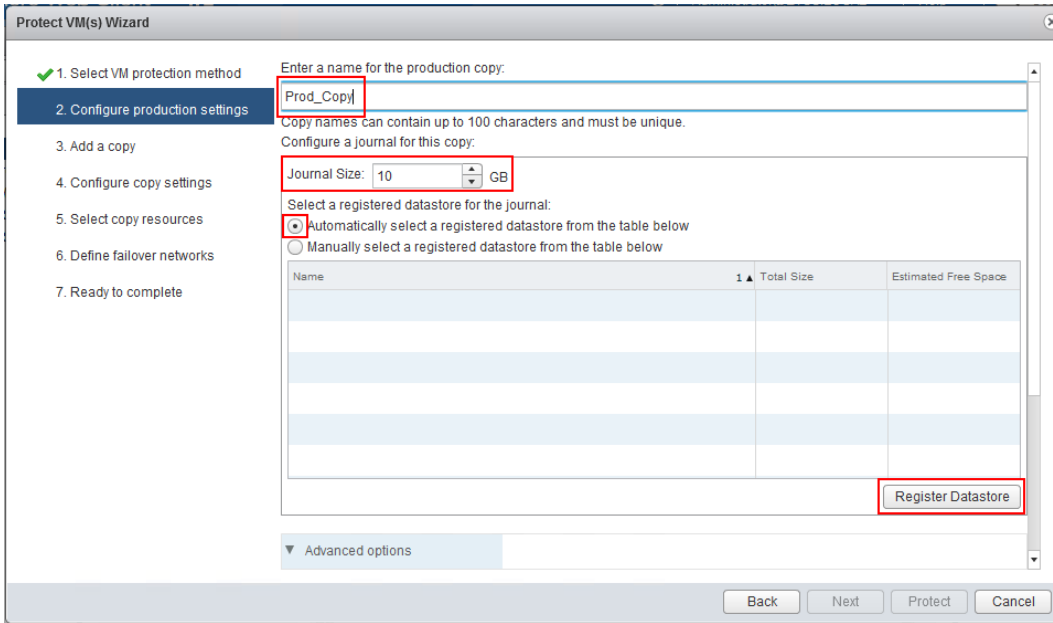
1. Connect to Site A vSphere Web Client.
2. Access the **VMs and Templates** view.
3. Power on the desired virtual machine for replication.
4. Right-click the VM, and select **All RecoverPoint for Virtual Machines Actions > Protect**.



5. Select VM protection method.
 - a. Select **Create a new consistency group for this VM**.
 - b. Name the new consistency group. It is recommended to use the VM or application name.
 - c. Select the vRPA cluster to use as a source for replication. Then click **Next**.

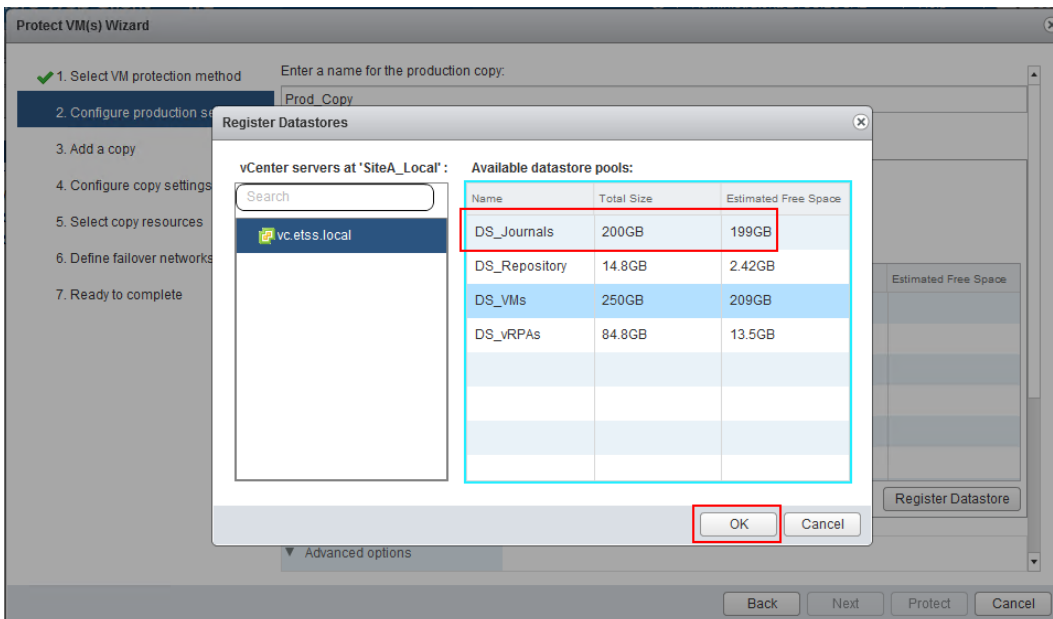


6. Configure production settings. Use the data that you collected in ["Data preparation."](#)
 - a. Type a name for the production copy. Differentiate the production copy from the remote copy (for example, "Prod_Copy"). Alternatively, use the site same (for example, "Site_A").
 - b. Configure a production journal using the default size of 10 GB.
 - c. Select a registered datastore for the production journal.



Note

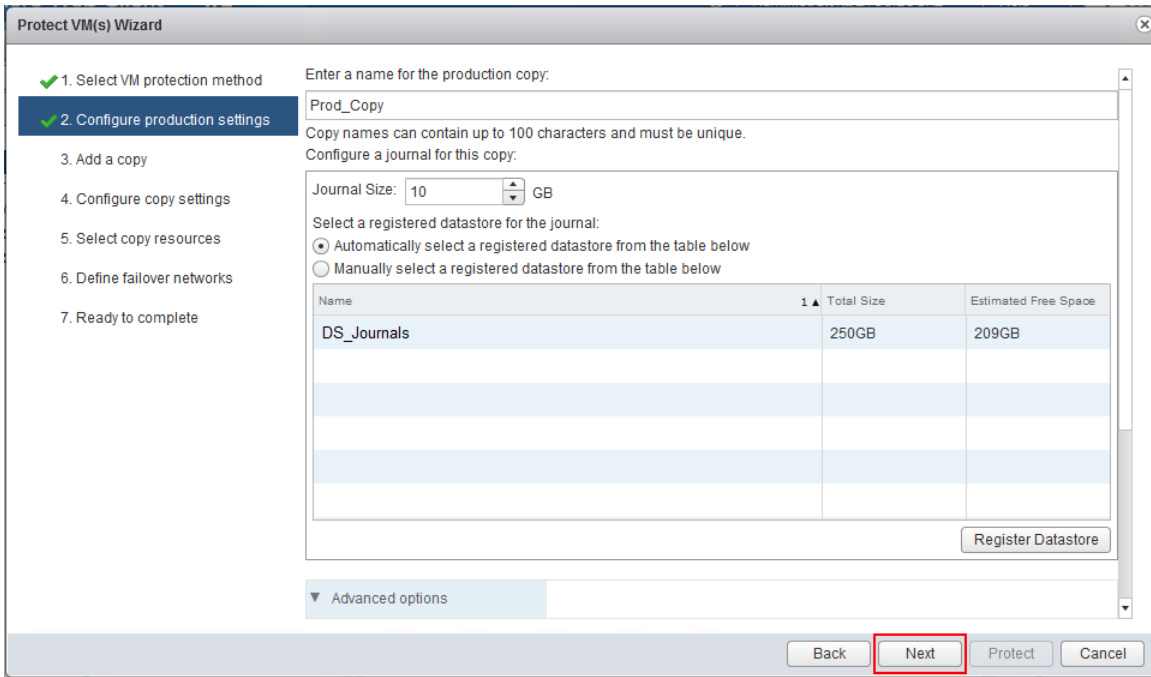
When protecting the first virtual machine within the RecoverPoint cluster, register a datastore by clicking **Register Datastore**. Select a datastore from the available datastore list. Click **OK**.



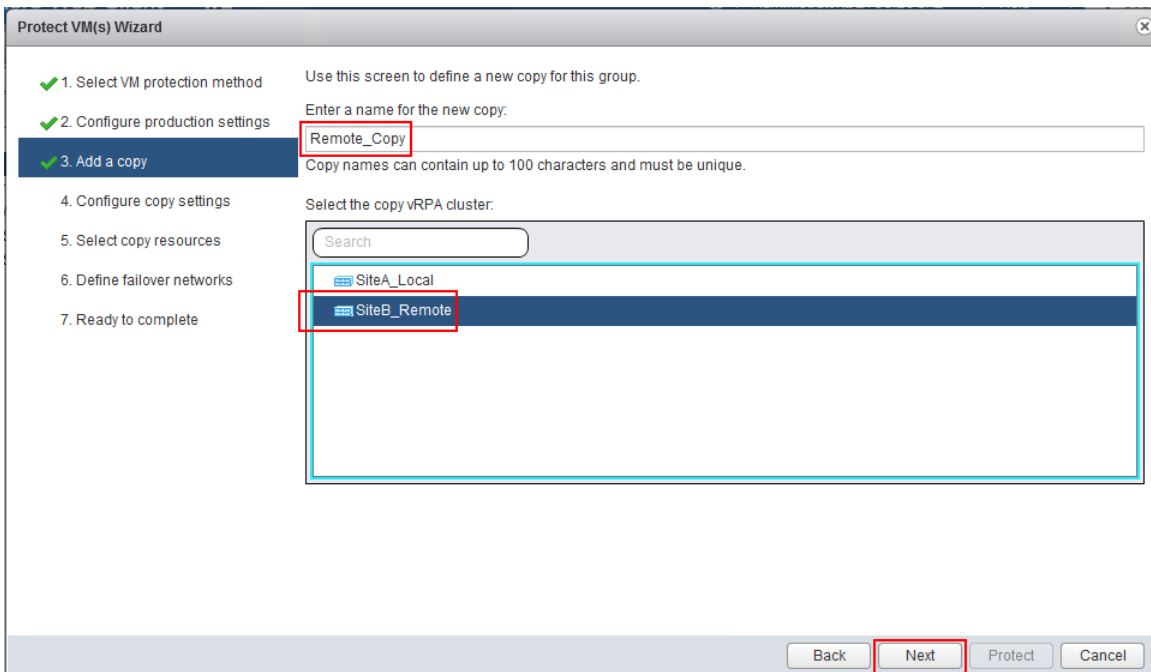
Note

The production journal tracks data that is not yet replicated to copy sites. It also tracks failovers. Changes that are made to the protected VM are stored on the journal VMDK before being replicated to copies.

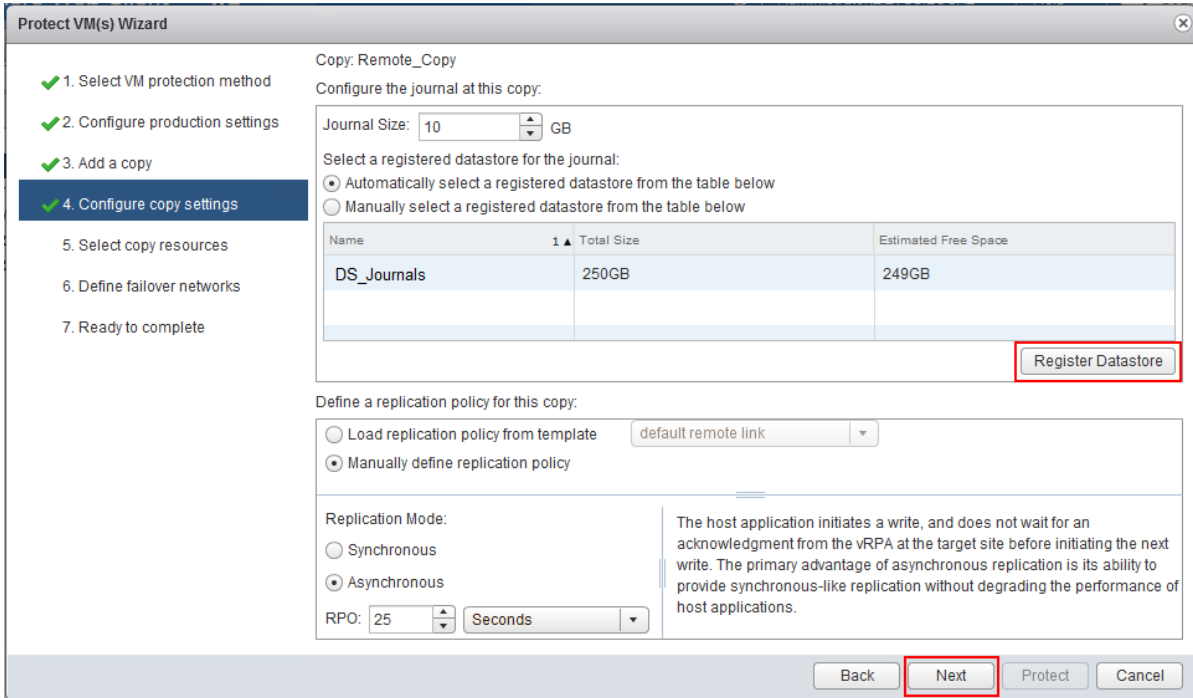
7. When finished with the configuration, click **Next**.



8. In the **Add a copy** step, type a name for the remote copy (for example, "Remote_Copy"), and select the vRPA cluster at Site B.



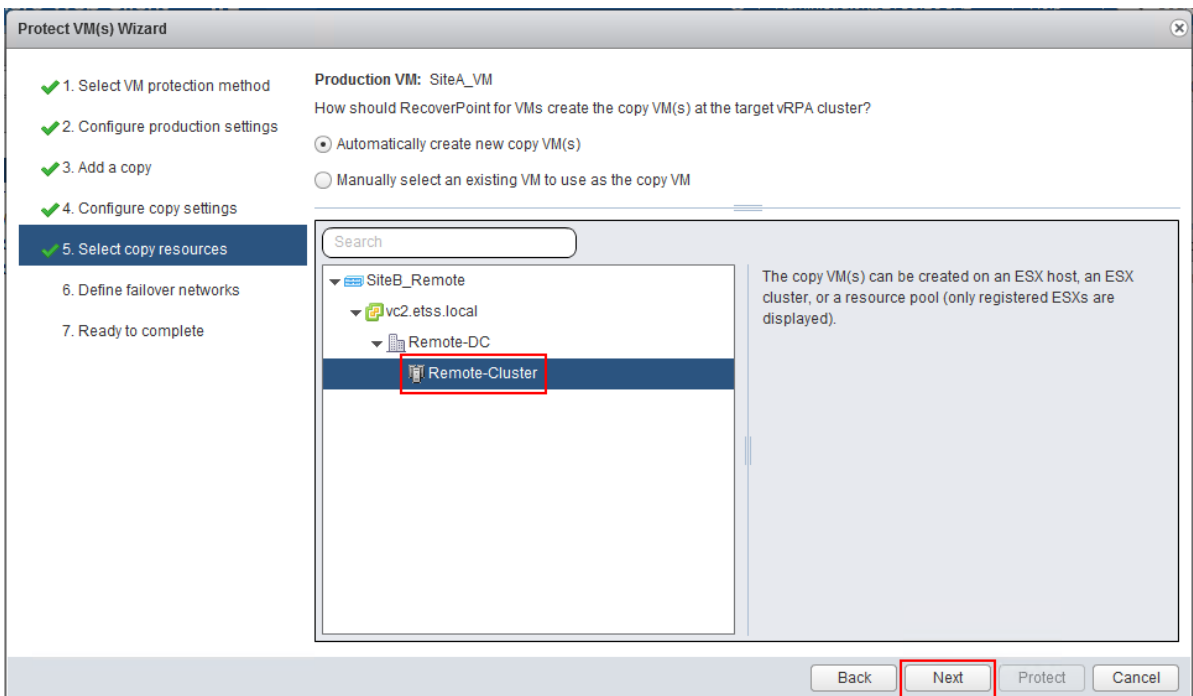
9. In the Configure copy settings step, use default settings, register a datastore for the replica journal, and then click **Next**.



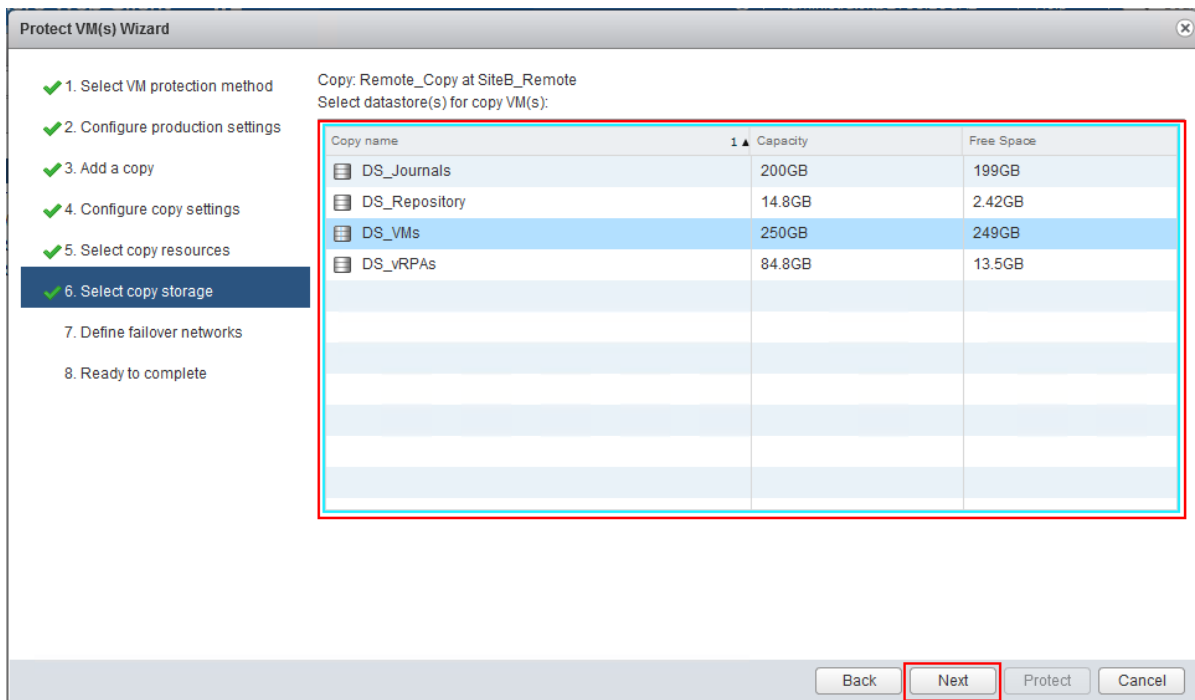
10. In the Select copy resources step, select the ESXi cluster on which to create the replica.

Note

If Distributed Resource Scheduler (DRS) is disabled, select a specific ESXi.



11. In the Select copy storage step, select a datastore on which the replica is created. Click **Next**.

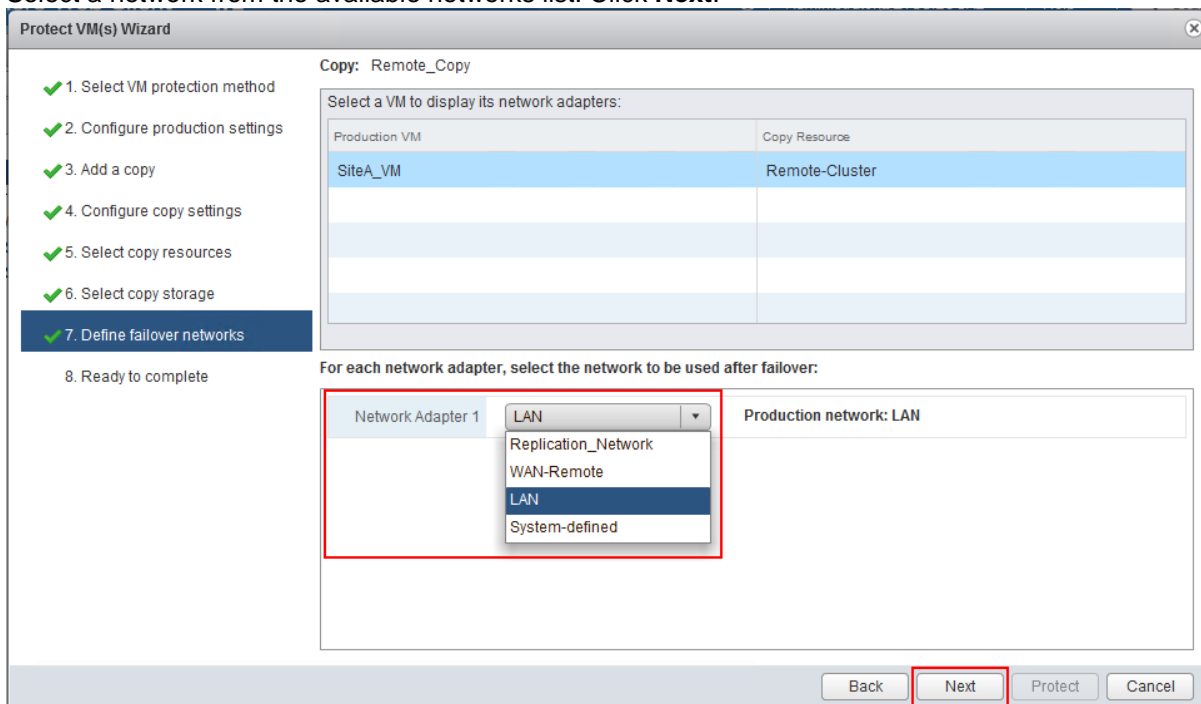


Note

This is the copy that you use at the remote site when failing over or when accessing an image (test a copy) without fail-over.

12. Select the network that the protected VM connects to in case of failover to remote site operation.

Select a network from the available networks list. Click **Next**.



13. Expand the **Production** and **Copy** settings to ensure that they are correct, and then click **Protect**.

The screenshot displays the 'Protect VM(s) Wizard' dialog box. On the left, a list of steps is shown, with step 8, 'Ready to complete', highlighted in dark blue. The main area shows a table for configuring a consistency group. The 'Group name' is 'CG1'. Under the 'Production (Prod_Copy)' and 'Copy (Remote_Copy)' sections, there are 'Edit...' and 'Delete' buttons. A checkbox at the bottom is checked, labeled 'Start replicating this consistency group when I click Finish'. At the bottom right, there are buttons for 'Back', 'Next', 'Protect', and 'Cancel', with the 'Protect' button highlighted by a red box.

Protection method	Create a new consistency group for this VM
Group name	CG1
Production (Prod_Copy)	Edit..
Copy (Remote_Copy)	Edit.. Delete

Start replicating this consistency group when I click Finish

Buttons: Back, Next, **Protect**, Cancel

Results

You have finished protecting the virtual machine and are ready to verify the replication status

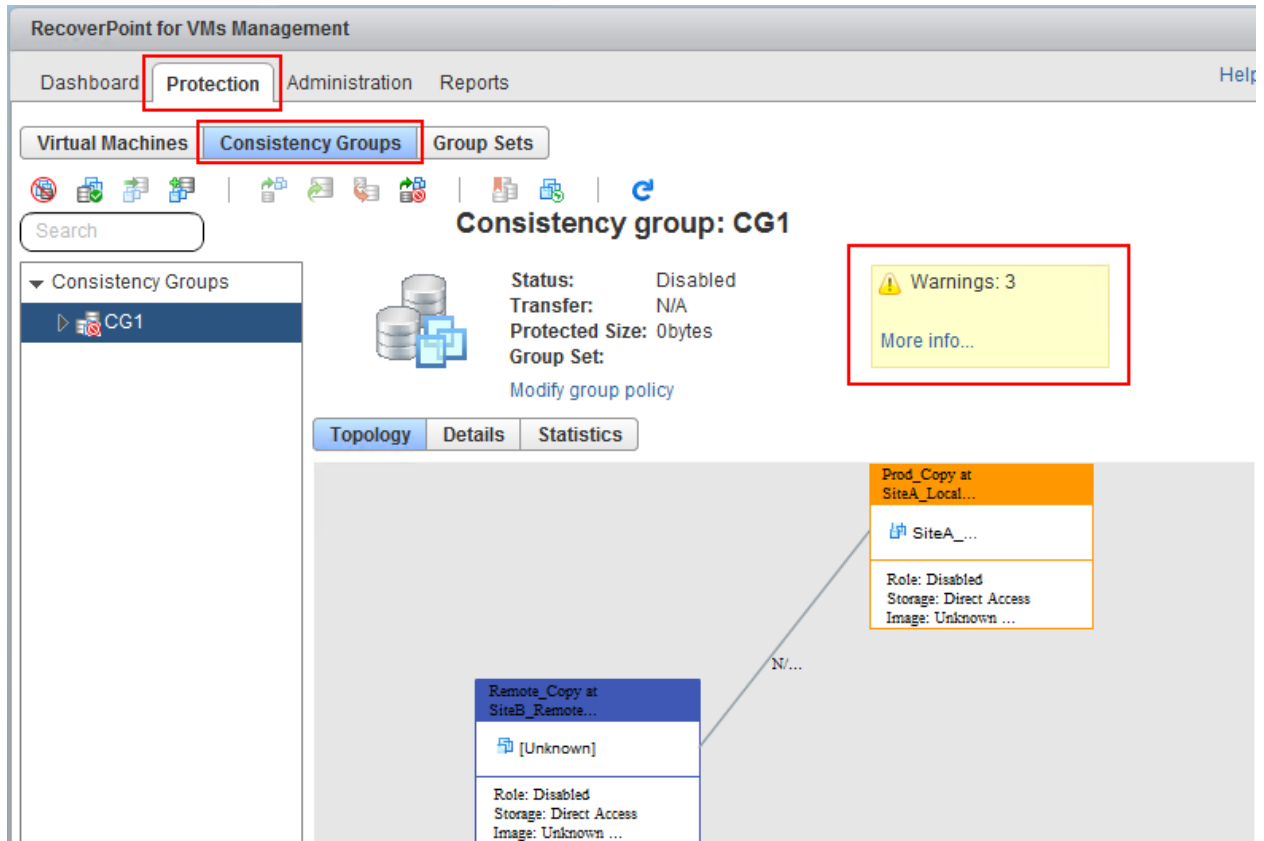
Verify Replication Status

From the RecoverPoint for VMs plug-in, verify the status of the replication.

Procedure

1. Select **RecoverPoint for VMs Management** > **Protection** > in the **Consistency Groups** tab, select the *CG Name*.

You might encounter temporary warnings and alerts as shown in this screen and the alert messages that follow:



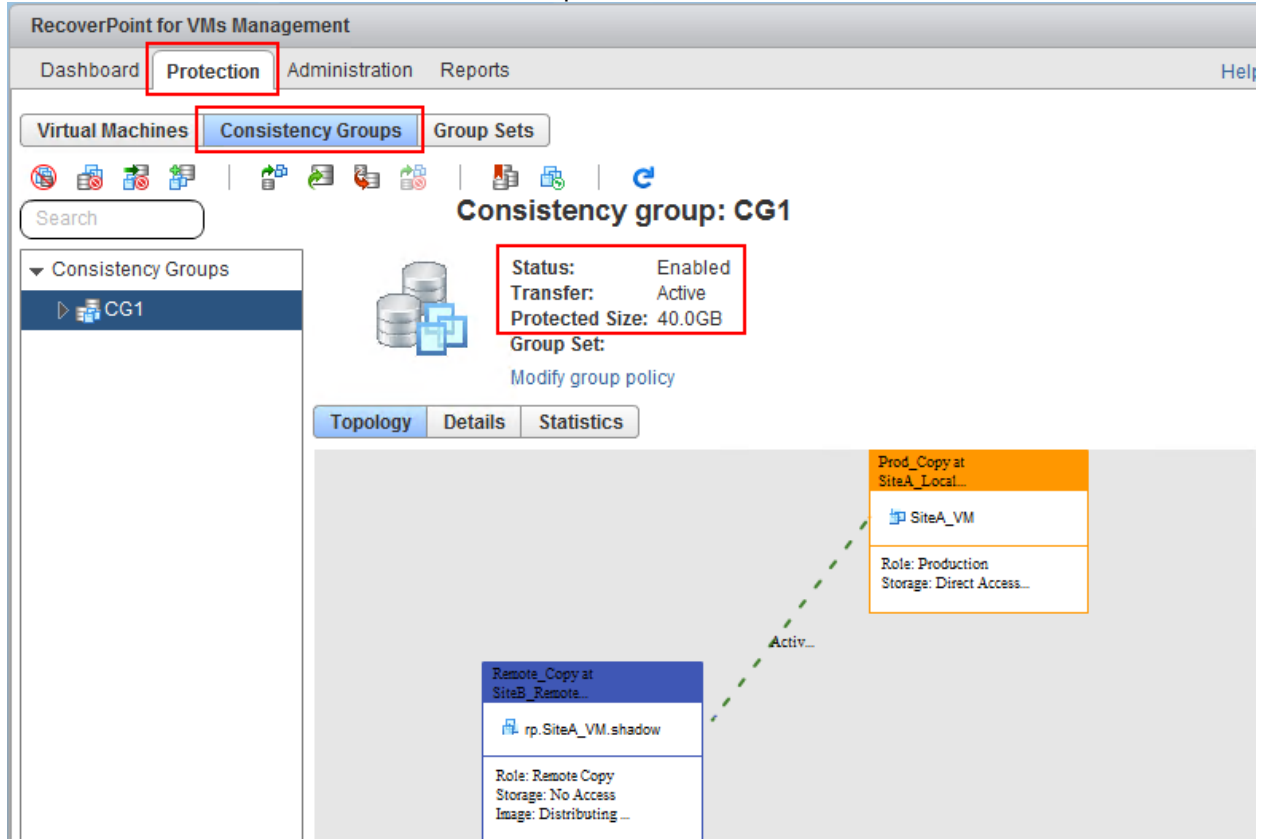
Alerts:

One or more volumes are currently being provisioned for copy [CG1, Remote_copy].

One or more volumes are currently being provisioned for copy [CG1, prod_Copy]

Accessibility of volume [CG1, prod_Copy], CG_RSET_SiteA_VM copy_0_0] by splitter Local-Cluster is unknown.

After a few minutes, the final status shows that protection is enabled and the transfer is active:

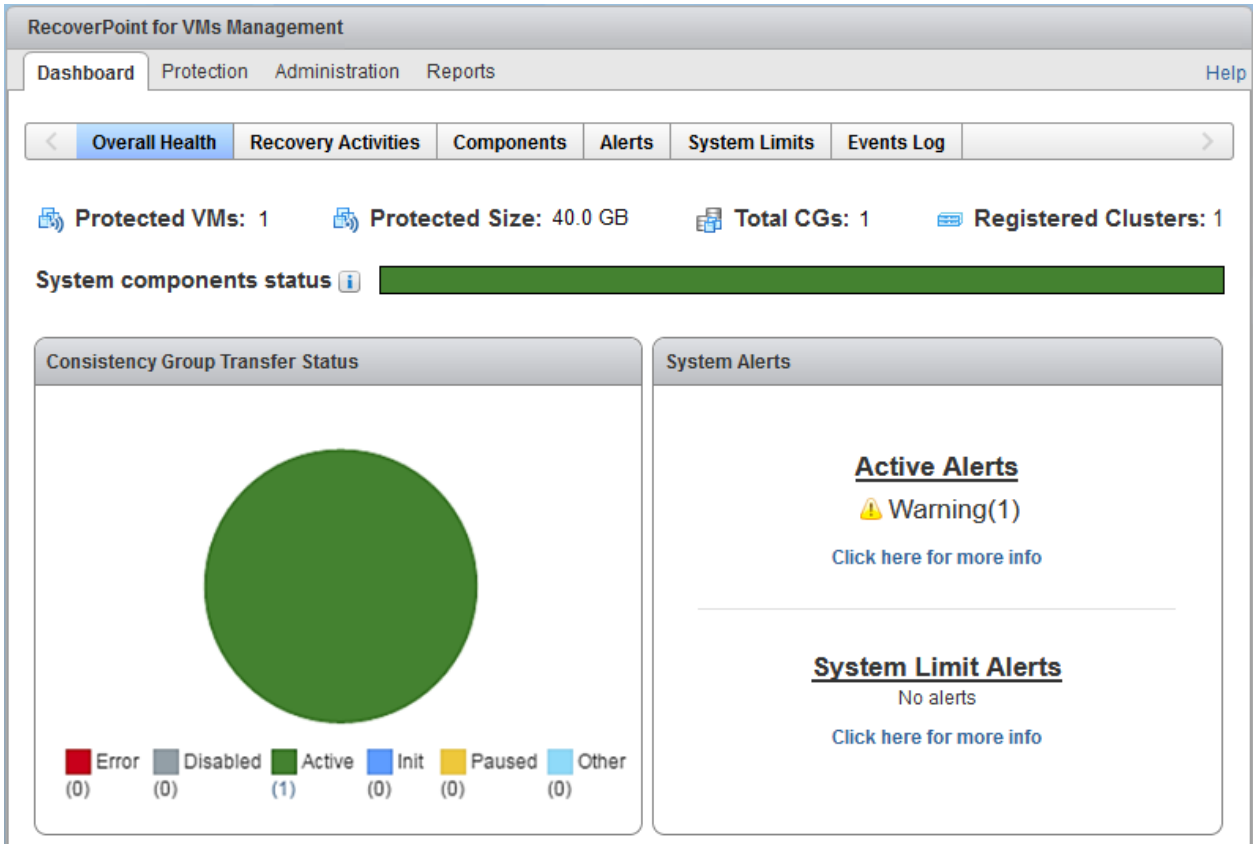


If warnings and alerts do not disappear after 10-15 minutes, contact Customer Support.

Note

- If the environment is configured in a way that causes one or both of the vRPAs to reside on the same ESXi with the replicated virtual machine, one warning remains as a notification:
At least one virtual RPA is running on the same ESX as the VM it is replicating. It is recommended not to have the virtual RPA that replicates a VM running on the same ESX as the replicated VM. Use vMotion to move one of them to another ESX.
- The replication process might require several minutes and initialization might require up to several hours depending on the size of the virtual disks and the speed of the data transfer between the two sites.

2. Monitor the health of the system from the RecoverPoint for VMs plug-in **Dashboard** tab.



Results

You have finished verifying replication status and are ready to perform the final steps.

Final RecoverPoint for VMs installation steps

About this task

Licensed users should activate their entitlements, add licenses, register their RecoverPoint for VMs system, and establish communication between their system and the System Reports database.

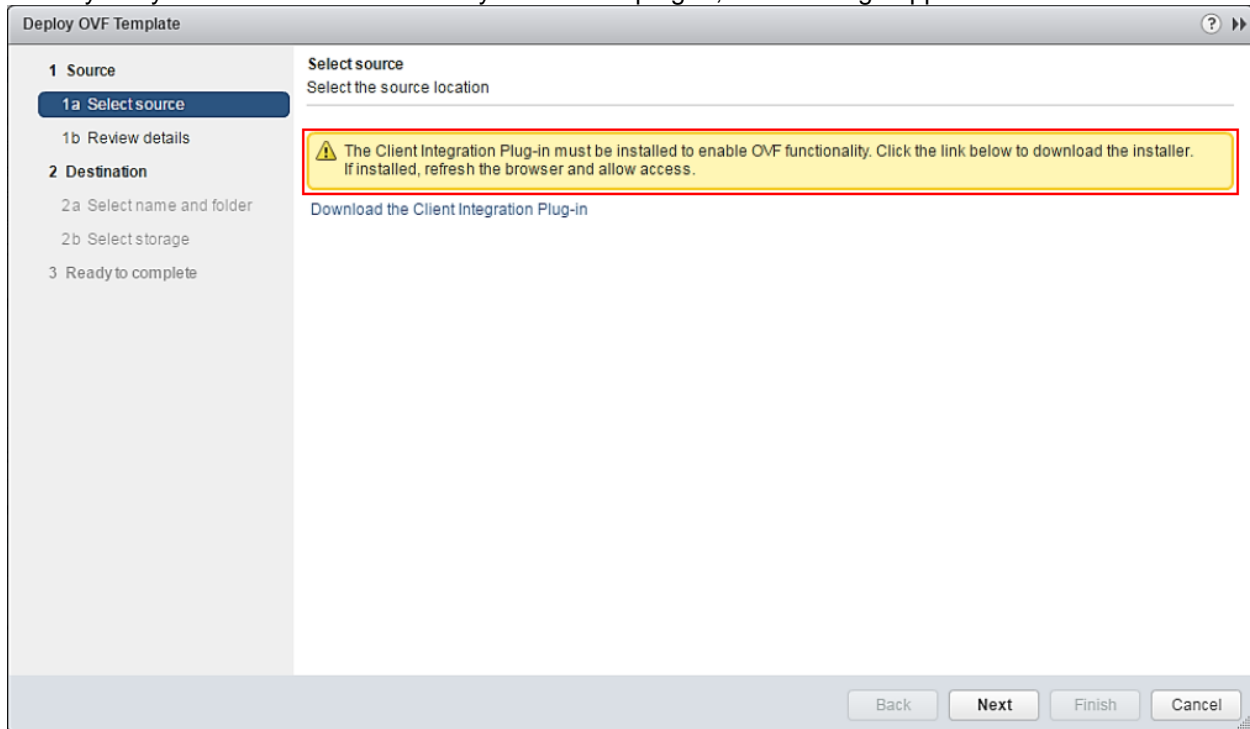
Steps to complete these tasks are in the "Getting started" section of the *RecoverPoint for Virtual Machines Administrator's Guide*, which is available at <https://support.emc.com>.

Appendix A: Install vSphere Client Integration Plug-in

Ensure that you have installed the Client Integration Plug-in for the vSphere web client; this plug-in is required for deploying OVF functionality.


About this task

When you try to enable OVF functionality without the plug-in, this message appears:



Procedure

1. If the Client Integration Plug-in is not installed, notice this option at the bottom of the vSphere web client login page (`https://vc ip or fqdn:9443`).



2. To install the Client Integration Plug-in, close all web browsers, click **Download Client Integration Plugin**, and select the default installation
3. After successful plug-in installation, ensure that the browser detects and enables the plug-in. If the plug-in is browser-enabled, you notice this option at the login screen:

