

vSphere Installation and Setup

Update 1

Modified on 04 DEC 2017

VMware vSphere 6.5

VMware ESXi 6.5

vCenter Server 6.5



vmware®

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2009–2017 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About vSphere Installation and Setup	5
Updated Information for the vSphere 6.5 Installation and Setup	6
1 Introduction to vSphere Installation and Setup	7
Overview of the vSphere Installation and Setup Process	7
vCenter Server Components and Services	10
Overview of the vCenter Server Appliance	13
vCenter Server and Platform Services Controller Deployment Types	14
Understanding vSphere Domains, Domain Names, and Sites	17
Deployment Topologies with External Platform Services Controller Instances and High Availability	18
Enhanced Linked Mode Overview	21
About ESXi Evaluation and Licensed Modes	22
2 Installing and Setting Up ESXi	23
ESXi Requirements	23
Preparing for Installing ESXi	30
Installing ESXi	82
Setting Up ESXi	194
After You Install and Set Up ESXi	215
3 Deploying the vCenter Server Appliance and Platform Services Controller Appliance	219
System Requirements for the vCenter Server Appliance and Platform Services Controller Appliance	220
Preparing for Deployment of the vCenter Server Appliance and Platform Services Controller Appliance	230
Prerequisites for Deploying the vCenter Server Appliance or Platform Services Controller Appliance	232
GUI Deployment of the vCenter Server Appliance and Platform Services Controller Appliance	233
CLI Deployment of the vCenter Server Appliance and Platform Services Controller Appliance	257
4 Installing vCenter Server and Platform Services Controller on Windows	273
vCenter Server for Windows Requirements	274
Preparing for Installing vCenter Server and Platform Services Controller on Windows	285
Required Information for Installing vCenter Server or Platform Services Controller on Windows	307
Installing vCenter Server and Platform Services Controller on Windows	309

- 5 After You Install vCenter Server or Deploy the vCenter Server Appliance 318**
 - Log in to vCenter Server by Using the vSphere Web Client 318
 - Install the VMware Enhanced Authentication Plug-in 319
 - Collect vCenter Server Log Files 320
 - Repoint vCenter Server to Another External Platform Services Controller 321
 - Reconfigure a Standalone vCenter Server with an Embedded Platform Services Controller to a vCenter Server with an External Platform Services Controller 323

- 6 File-Based Backup and Restore of vCenter Server Appliance 327**
 - Considerations and Limitations for File-Based Backup and Restore 328
 - Back up a vCenter Server Appliance by Using the vCenter Server Appliance Management Interface 331
 - Restore a vCenter Server Appliance from a File-Based Backup 333

- 7 Image-Based Backup and Restore of a vCenter Server Environment 340**
 - Considerations and Limitations for Image-Based Backup and Restore 341
 - Use vSphere Data Protection to Back Up a vCenter Server Environment 344
 - Use vSphere Data Protection to Restore a vCenter Server Environment 349

- 8 Troubleshooting ESXi Booting 381**
 - Host Stops Unexpectedly at Bootup When Sharing a Boot Disk with Another Host 381
 - Host Fails to Boot After You Install ESXi in UEFI Mode 382

- 9 Troubleshooting vCenter Server Installation or Deployment 383**
 - Collecting Logs for Troubleshooting a vCenter Server Installation or Upgrade 383
 - Attempt to Install a Platform Services Controller After a Prior Installation Failure 385
 - Microsoft SQL Database Set to Unsupported Compatibility Mode Causes vCenter Server Installation or Upgrade to Fail 386

- 10 Decommissioning ESXi and vCenter Server 387**
 - Decommission an ESXi Host 387
 - Uninstall vCenter Server 387

About vSphere Installation and Setup

vSphere Installation and Setup describes how to install and configure VMware vCenter Server[®], deploy the VMware vCenter[®] Server Appliance™, and install and configure VMware ESXi™.

Intended Audience

vSphere Installation and Setup is intended for experienced administrators who want to install and configure vCenter Server, deploy and configure the vCenter Server Appliance, and install and configure ESXi.

This information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and data center operations. The information about using the Image Builder and VMware vSphere[®] Auto Deploy™ is written for administrators who have experience with Microsoft PowerShell and VMware vSphere[®] PowerCLI™.

vSphere Web Client and vSphere Client

Task instructions in this guide are based on the vSphere Web Client. You can also perform most of the tasks in this guide by using the new vSphere Client. The new vSphere Client user interface terminology, topology, and workflow are closely aligned with the same aspects and elements of the vSphere Web Client user interface. You can apply the vSphere Web Client instructions to the new vSphere Client unless otherwise instructed.

Note Not all functionality in the vSphere Web Client has been implemented for the vSphere Client in the vSphere 6.5 release. For an up-to-date list of unsupported functionality, see *Functionality Updates for the vSphere Client Guide* at <http://www.vmware.com/info?id=1413>.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

Updated Information for the vSphere 6.5 Installation and Setup

This *vSphere Installation and Setup* is updated with each release of the product or when necessary.

This table provides the update history of the *vSphere Installation and Setup* .

Revision	Description
04 DEC 2017	Added consideration item for SCP protocol in Considerations and Limitations for File-Based Backup and Restore .
EN-002587-01	Initial release.

Introduction to vSphere Installation and Setup

1

vSphere 6.5 provides various options for installation and setup. To ensure a successful vSphere deployment, understand the installation and setup options, and the sequence of tasks.

The two core components of vSphere are ESXi and vCenter Server. ESXi is the virtualization platform on which you can create and run virtual machines and virtual appliances. vCenter Server is a service that acts as a central administrator for ESXi hosts connected in a network. vCenter Server lets you pool and manage the resources of multiple hosts.

You can install vCenter Server on a Windows virtual machine or physical server, or deploy the vCenter Server Appliance. The vCenter Server Appliance is a preconfigured Linux-based virtual machine optimized for running vCenter Server and the vCenter Server components. You can deploy the vCenter Server Appliance on ESXi hosts 5.5 or later, or on vCenter Server instances 5.5 or later.

Starting with vSphere 6.0, all prerequisite services for running vCenter Server and the vCenter Server components are bundled in the VMware Platform Services Controller™. You can deploy vCenter Server with an embedded or external Platform Services Controller, but you must always install or deploy the Platform Services Controller before installing or deploying vCenter Server.

This section includes the following topics:

- [Overview of the vSphere Installation and Setup Process](#)
- [vCenter Server Components and Services](#)
- [Overview of the vCenter Server Appliance](#)
- [vCenter Server and Platform Services Controller Deployment Types](#)
- [Understanding vSphere Domains, Domain Names, and Sites](#)
- [Deployment Topologies with External Platform Services Controller Instances and High Availability](#)
- [Enhanced Linked Mode Overview](#)
- [About ESXi Evaluation and Licensed Modes](#)

Overview of the vSphere Installation and Setup Process

vSphere is a sophisticated product with multiple components to install and set up. To ensure a successful vSphere deployment, understand the sequence of tasks required.

Installing vSphere includes the following tasks:

Figure 1-1. vSphere Installation and Setup Workflow



- 1 Read the vSphere release notes.
- 2 Install ESXi.
 - a Verify that your system meets the minimum hardware requirements. See [ESXi Requirements](#).
 - b Determine the ESXi installation option to use. See [Options for Installing ESXi](#).

- c Determine where you want to locate and boot the ESXi installer. See [Media Options for Booting the ESXi Installer](#). If you are using PXE to boot the installer, verify that your network PXE infrastructure is properly set up. See [PXE Booting the ESXi Installer](#).
- d Create a worksheet with the information you will need when you install ESXi. See [Required Information for ESXi Installation](#).
- e Install ESXi.
 - [Installing ESXi Interactively](#)
 - [Installing or Upgrading Hosts by Using a Script](#)

Note You can also provision ESXi hosts by using vSphere Auto Deploy, but vSphere Auto Deploy is installed together with vCenter Server. To provision ESXi hosts by using Auto Deploy, you must deploy the vCenter Server Appliance or install vCenter Server.

- 3 Configure the ESXi boot and network settings, the direct console, and other settings. See [Setting Up ESXi](#) and [After You Install and Set Up ESXi](#).
- 4 Consider setting up a syslog server for remote logging, to ensure sufficient disk storage for log files. Setting up logging on a remote host is especially important for hosts with limited local storage. See [Required Free Space for System Logging](#) and [Configure Syslog on ESXi Hosts](#).
- 5 Determine the vCenter Server and Platform Services Controller deployment model that is suitable for your environment.

vCenter Server with an embedded Platform Services Controller deployment is suitable for small-scale environments. vCenter Server with an external Platform Services Controller deployment is suitable for environments with several vCenter Server instances. See [vCenter Server and Platform Services Controller Deployment Types](#).

- 6 Deploy or install vCenter Server and Platform Services Controller.

You can deploy the vCenter Server Appliance or Platform Services Controller appliance on an ESXi host or vCenter Server instance, or you can install vCenter Server and Platform Services Controller on a Windows virtual machine or physical server.

You can deploy or install multiple vCenter Server instances connected in Enhanced Linked Mode configuration by registering them to a common or different joined Platform Services Controller instances.

- Deploy the vCenter Server Appliance or Platform Services Controller appliance.
 - 1 Review the topics in [System Requirements for the vCenter Server Appliance and Platform Services Controller Appliance](#) and verify that your system meets the hardware and software requirements for deploying the appliance.
 - 2 Determine the deployment method to use.

You can use the GUI method to deploy the appliance interactively. You can use the CLI method to perform a silent deployment of the appliance. See [GUI Deployment of the vCenter Server Appliance and Platform Services Controller Appliance](#) and [CLI Deployment of the vCenter Server Appliance and Platform Services Controller Appliance](#).

- 3 Use the topic [Required Information for Deploying a vCenter Server Appliance or Platform Services Controller Appliance](#) to create a worksheet with the information you need for the GUI deployment, or use the topic [Prepare Your JSON Configuration File for CLI Deployment](#) to create your JSON templates for the CLI deployment.
 - 4 Deploy the appliance.
- Install vCenter Server or Platform Services Controller on a Windows virtual machine or physical server.
 - 1 Verify that your system meets the hardware and software requirements for installing vCenter Server. See [vCenter Server for Windows Requirements](#).
 - 2 (Optional) Set up an external vCenter Server database. See [Preparing vCenter Server Databases for Install](#).

For an environment with up to 20 hosts and 200 virtual machines, you can use the bundled PostgreSQL database. For production and large scale environments, set up an external database, because the migration from the embedded PostgreSQL database to an external database is not a trivial manual process.
 - 3 Create a worksheet with the information you need for installation. See [Required Information for Installing vCenter Server or Platform Services Controller on Windows](#).
 - 4 Install vCenter Server with an embedded Platform Services Controller, Platform Services Controller, or vCenter Server with an external Platform Services Controller.
- 7 Connect to vCenter Server from the vSphere Web Client. See [Chapter 5 After You Install vCenter Server or Deploy the vCenter Server Appliance](#).
 - 8 Configure the vCenter Server Appliance or vCenter Server instance. See *vCenter Server Appliance Configuration* and *vCenter Server and Host Management*.

vCenter Server Components and Services

vCenter Server provides a centralized platform for management, operation, resource provisioning, and performance evaluation of virtual machines and hosts.

When you install vCenter Server with an embedded Platform Services Controller, or deploy the vCenter Server Appliance with an embedded Platform Services Controller, vCenter Server, the vCenter Server components, and the services included in the Platform Services Controller are deployed on the same system.

When you install vCenter Server with an external Platform Services Controller, or deploy the vCenter Server Appliance with an external Platform Services Controller, vCenter Server and the vCenter Server components are deployed on one system, and the services included in the Platform Services Controller are deployed on another system.

The following components are included in the vCenter Server and vCenter Server Appliance installations:

- The VMware Platform Services Controller group of infrastructure services contains vCenter Single Sign-On, License service, Lookup Service, and VMware Certificate Authority.

- The vCenter Server group of services contains vCenter Server, vSphere Web Client, vSphere Auto Deploy, and vSphere ESXi Dump Collector. vCenter Server for Windows also contains the VMware vSphere Syslog Collector. The vCenter Server Appliance also contains the VMware vSphere Update Manager Extension service.

Note Starting with vSphere 6.5, all vCenter Server services and some Platform Services Controller services run as child processes of the VMware Service Lifecycle Manager service.

Services Installed with VMware Platform Services Controller

vCenter Single Sign-On The vCenter Single Sign-On authentication service provides secure authentication services to the vSphere software components. By using vCenter Single Sign-On, the vSphere components communicate with each other through a secure token exchange mechanism, instead of requiring each component to authenticate a user separately with a directory service like Active Directory. vCenter Single Sign-On constructs an internal security domain (for example, vsphere.local) where the vSphere solutions and components are registered during the installation or upgrade process, providing an infrastructure resource. vCenter Single Sign-On can authenticate users from its own internal users and groups, or it can connect to trusted external directory services such as Microsoft Active Directory. Authenticated users can then be assigned registered solution-based permissions or roles within a vSphere environment.

vCenter Single Sign-On is required with vCenter Server.

vSphere License Service The vSphere License service provides common license inventory and management capabilities to all vCenter Server systems that are connected to a Platform Services Controller or multiple linked Platform Services Controllers.

VMware Certificate Authority VMware Certificate Authority (VMCA) provisions each ESXi host with a signed certificate that has VMCA as the root certificate authority, by default. Provisioning occurs when the ESXi host is added to vCenter Server explicitly or as part of the ESXi host installation process. All ESXi certificates are stored locally on the host.

For information about all Platform Services Controller services and capabilities, see *Platform Services Controller Administration*.

Services Installed with vCenter Server

These additional components are installed silently when you install vCenter Server. The components cannot be installed separately as they do not have their own installers.

PostgreSQL	A bundled version of the VMware distribution of PostgreSQL database for vSphere and vCloud Hybrid Services.
vSphere Web Client	The vSphere Web Client lets you connect to vCenter Server instances by using a Web browser, so that you can manage your vSphere infrastructure.
vSphere Client	The new user interface that lets you connect to vCenter Server instances by using a Web browser. The terminology, topology, and workflow are closely aligned with the same aspects and elements of the vSphere Web Client user interface.

Note Not all functionality in the vSphere Web Client has been implemented for the vSphere Client in the vSphere 6.5 release. For an up-to-date list of unsupported functionality, see *Functionality Updates for the vSphere Client Guide* at <http://www.vmware.com/info?id=1413>.

vSphere ESXi Dump Collector	The vCenter Server support tool. You can configure ESXi to save the VMkernel memory to a network server, rather than to a disk, when the system encounters a critical failure. The vSphere ESXi Dump Collector collects such memory dumps over the network.
VMware vSphere Syslog Collector	<p>The vCenter Server on Windows support tool that enables network logging and combining of logs from multiple hosts. You can use the vSphere Syslog Collector to direct ESXi system logs to a server on the network, rather than to a local disk. The recommended maximum number of supported hosts to collect logs from is 30. For information about configuring vSphere Syslog Collector, see http://kb.vmware.com/kb/2021652.</p> <p>The vCenter Server Appliance uses the built-in Rsyslog service of the Linux OS. For information how to redirect the log files to another machine with the Appliance Management Interface, see <i>vCenter Server Appliance Configuration</i>.</p>

vSphere Auto Deploy	The vCenter Server support tool that can provision hundreds of physical hosts with ESXi software. You can specify the image to deploy and the hosts to provision with the image. Optionally, you can specify host profiles to apply to the hosts, and a vCenter Server location (folder or cluster) for each host.
VMware vSphere Update Manager Extension	Update Manager enables centralized, automated patch and version management for VMware vSphere and offers support for VMware ESXi hosts, virtual machines, and virtual appliances. The VMware vSphere Update Manager Extension is an optional service of only the vCenter Server Appliance 6.5.

Overview of the vCenter Server Appliance

The vCenter Server Appliance is a preconfigured Linux-based virtual machine that is optimized for running vCenter Server and the associated services.

The vCenter Server Appliance reduces the deployment time of vCenter Server and the associated services, and provides a low-cost alternative to the Windows-based vCenter Server installation.

The vCenter Server Appliance package contains the following software:

- Project Photon OS[®] 1.0
- The Platform Services Controller group of infrastructure services
- The vCenter Server group of services
- PostgreSQL
- VMware vSphere Update Manager Extension

Version 6.5 of the vCenter Server Appliance is deployed with virtual hardware version 10, which supports 64 virtual CPUs per virtual machine in ESXi.

The vCenter Server Appliance uses the embedded PostgreSQL database that has the scalability of up to 2,000 hosts and 35,000 virtual machines. During the deployment, you can choose the vCenter Server Appliance size for your vSphere environment size and the storage size for your database requirements.

Starting with vSphere 6.5, the vCenter Server uses the VMware vSphere Update Manager Extension service. An external VMware Update Manager instance on Windows is no longer required for vSphere centralized automated patch and version management. For information about the vCenter Server and Platform Services Controller services, see [vCenter Server Components and Services](#).

Starting with vSphere 6.5, the vCenter Server Appliance supports high availability. For information about configuring vCenter Server Appliance in a vCenter High Availability cluster, see *vSphere Availability*.

Starting with vSphere 6.5, the vCenter Server Appliance and Platform Services Controller appliance support file-based backup and restore. For information backing up and restoring, see [Chapter 6 File-Based Backup and Restore of vCenter Server Appliance](#).

For information about the vCenter Server Appliance maximums, see the *Configuration Maximums* documentation.

vCenter Server and Platform Services Controller Deployment Types

You can deploy the vCenter Server Appliance or install vCenter Server for Windows with an embedded or external Platform Services Controller. You can also deploy a Platform Services Controller as an appliance or install it on Windows. If necessary, you can use a mixed operating systems environment.

Before you deploy the vCenter Server Appliance or install vCenter Server for Windows, you must determine the deployment model that is suitable for your environment. For each deployment or installation, you must select one of the three deployment types.

Table 1-1. vCenter Server and Platform Services Controller Deployment Types

Deployment Type	Description
vCenter Server with an embedded Platform Services Controller	All services that are bundled with the Platform Services Controller are deployed together with the vCenter Server services on the same virtual machine or physical server.
Platform Services Controller	Only the services that are bundled with the Platform Services Controller are deployed on the virtual machine or physical server.
vCenter Server with an external Platform Services Controller (Requires external Platform Services Controller)	Only the vCenter Server services are deployed on the virtual machine or physical server. You must register such a vCenter Server instance with a Platform Services Controller instance that you previously deployed or installed.

vCenter Server with an Embedded Platform Services Controller

Using an embedded Platform Services Controller results in a standalone deployment that has its own vCenter Single Sign-On domain with a single site. vCenter Server with an embedded Platform Services Controller is suitable for small environments. You cannot join other vCenter Server or Platform Services Controller instances to this vCenter Single Sign-On domain.

Figure 1-2. vCenter Server with an Embedded Platform Services Controller



Installing vCenter Server with an embedded Platform Services Controller has the following advantages:

- The connection between vCenter Server and the Platform Services Controller is not over the network, and vCenter Server is not prone to outages caused by connectivity and name resolution issues between vCenter Server and the Platform Services Controller.
- If you install vCenter Server on Windows virtual machines or physical servers, you need fewer Windows licenses.
- You manage fewer virtual machines or physical servers.

Installing vCenter Server with an embedded Platform Services Controller has the following disadvantages:

- There is a Platform Services Controller for each product which might be more than required and which consumes more resources.
- The model is suitable only for small-scale environments.

You can configure the vCenter Server Appliance with an embedded Platform Services Controller in vCenter High Availability configuration. For information, see *vSphere Availability*.

Note After you deploy or install vCenter Server with an embedded Platform Services Controller, you can reconfigure the deployment type and switch to vCenter Server with an external Platform Services Controller.

See [Reconfigure a Standalone vCenter Server with an Embedded Platform Services Controller to a vCenter Server with an External Platform Services Controller](#).

Platform Services Controller and vCenter Server with an External Platform Services Controller

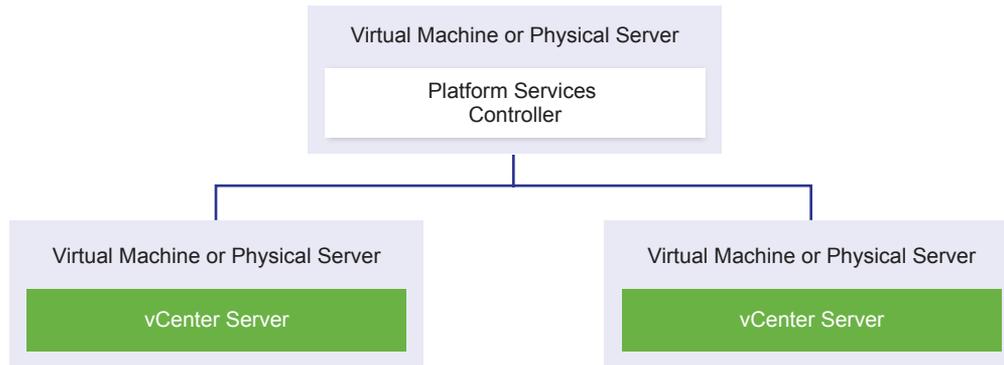
When you deploy or install a Platform Services Controller instance, you can create a vCenter Single Sign-On domain or join an existing vCenter Single Sign-On domain. Joined Platform Services Controller instances replicate their infrastructure data, such as authentication and licensing information, and can span multiple vCenter Single Sign-On sites. For information, see [Understanding vSphere Domains, Domain Names, and Sites](#).

For information about managing the Platform Services Controller services, see *Platform Services Controller Administration*.

You can register multiple vCenter Server instances with one common external Platform Services Controller instance. The vCenter Server instances assume the vCenter Single Sign-On site of the Platform Services Controller instance with which they are registered. All vCenter Server instances that are registered with one common or different joined Platform Services Controller instances are connected in Enhanced Linked Mode.

See [Enhanced Linked Mode Overview](#).

Figure 1-3. Example of Two vCenter Server Instances with a Common External Platform Services Controller



Installing vCenter Server with an external Platform Services Controller has the following advantages:

- Fewer resources consumed by the shared services in the Platform Services Controller instances.
- The model is suitable for large-scale environments.

Installing vCenter Server with an external Platform Services Controller has the following disadvantages:

- The connection between vCenter Server and Platform Services Controller might have connectivity and name resolution issues.
- If you install vCenter Server on Windows virtual machines or physical servers, you need more Microsoft Windows licenses.
- You must manage more virtual machines or physical servers.

For information about the Platform Services Controller and vCenter Server maximums, see the *Configuration Maximums* documentation.

For information about the deployment topologies and Platform Services Controller high availability, see [Deployment Topologies with External Platform Services Controller Instances and High Availability](#).

For information about configuring the vCenter Server Appliance with an external Platform Services Controller in vCenter High Availability configuration, see *vSphere Availability*.

Mixed Operating Systems Environment

A vCenter Server instance installed on Windows can be registered with either a Platform Services Controller installed on Windows or a Platform Services Controller appliance. A vCenter Server Appliance can be registered with either a Platform Services Controller installed on Windows or a Platform Services Controller appliance. Both vCenter Server and the vCenter Server Appliance can be registered with the same Platform Services Controller.

Figure 1-4. Example of a Mixed Operating Systems Environment with an External Platform Services Controller on Windows

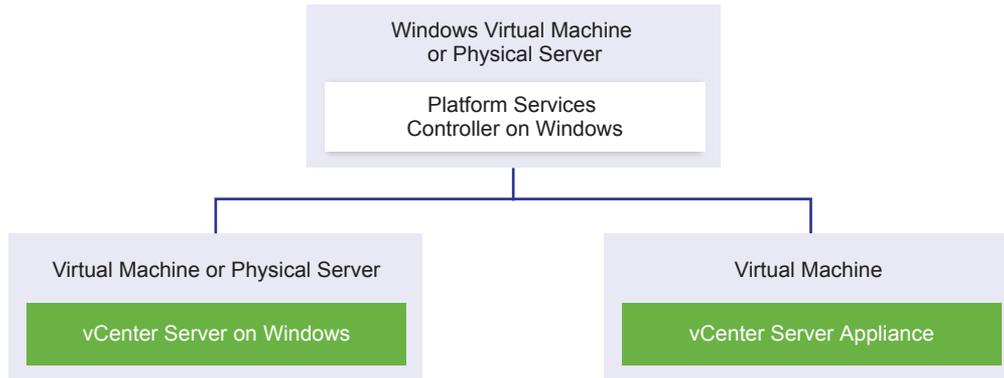
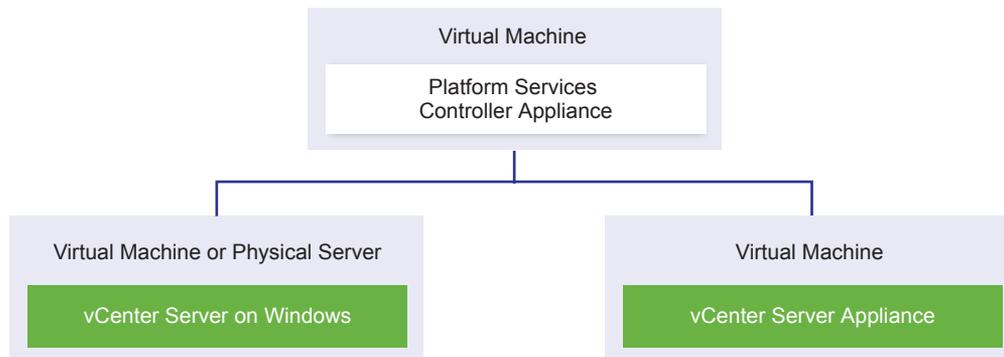


Figure 1-5. Example of a Mixed Operating Systems Environment with an External Platform Services Controller Appliance



Note To ensure easy manageability and maintenance, use only appliances or only Windows installations of vCenter Server and Platform Services Controller.

Understanding vSphere Domains, Domain Names, and Sites

Each Platform Services Controller is associated with a vCenter Single Sign-On domain. The domain name defaults to `vsphere.local`, but you can change it during installation of the first Platform Services Controller. The domain determines the local authentication space. You can split a domain into multiple sites, and assign each Platform Services Controller and vCenter Server instance to a site. Sites are logical constructs, but usually correspond to geographic location.

Platform Services Controller Domain

When you install a Platform Services Controller, you are prompted to create a vCenter Single Sign-On domain or join an existing domain.

The domain name is used by the VMware Directory Service (`vmdir`) for all Lightweight Directory Access Protocol (LDAP) internal structuring.

With vSphere 6.0 and later, you can give your vSphere domain a unique name. To prevent authentication conflicts, use a name that is not used by OpenLDAP, Microsoft Active Directory, and other directory services.

Note You cannot change the domain to which a Platform Services Controller or vCenter Server instance belongs.

If you are upgrading from vSphere 5.5, your vSphere domain name remains the default (vsphere.local). For all versions of vSphere, you cannot change the name of a domain.

After you specify the name of your domain, you can add users and groups. It usually makes more sense to add an Active Directory or LDAP identity source and allow the users and groups in that identity source to authenticate. You can also add vCenter Server or Platform Services Controller instances, or other VMware products, such as vRealize Operations, to the domain.

Platform Services Controller Sites

You can organize Platform Services Controller domains into logical sites. A site in the VMware Directory Service is a logical container for grouping Platform Services Controller instances within a vCenter Single Sign-On domain.

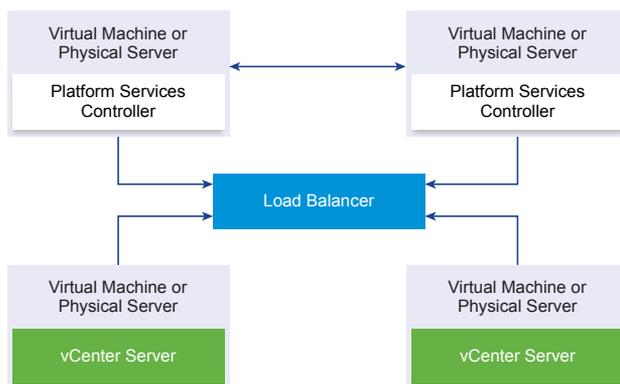
You are prompted for the site name when you install or upgrade a Platform Services Controller. See the *vSphere Installation and Setup* documentation.

Deployment Topologies with External Platform Services Controller Instances and High Availability

To ensure Platform Services Controller high availability in external deployments, you must install or deploy at least two joined Platform Services Controller instances in your vCenter Single Sign-On domain. When you use a third-party load balancer, you can ensure an automatic failover without downtime.

Platform Services Controller with a Load Balancer

Figure 1-6. Example of a Load Balanced Pair of Platform Services Controller Instances



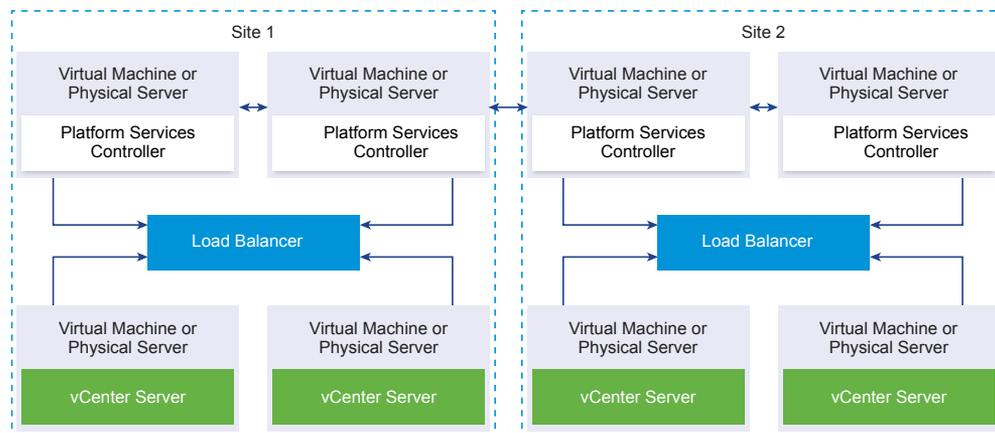
You can use a third-party load balancer per site to configure Platform Services Controller high availability with automatic failover for this site. For information about the maximum number of Platform Services Controller instances behind a load balancer, see the *Configuration Maximums* documentation.

Important To configure Platform Services Controller high availability behind a load balancer, the Platform Services Controller instances must be of the same operating system type. Mixed operating systems Platform Services Controller instances behind a load balancer are unsupported.

The vCenter Server instances are connected to the load balancer. When a Platform Services Controller instance stops responding, the load balancer automatically distributes the load among the other functional Platform Services Controller instances without downtime.

Platform Services Controller with Load Balancers Across vCenter Single Sign-On Sites

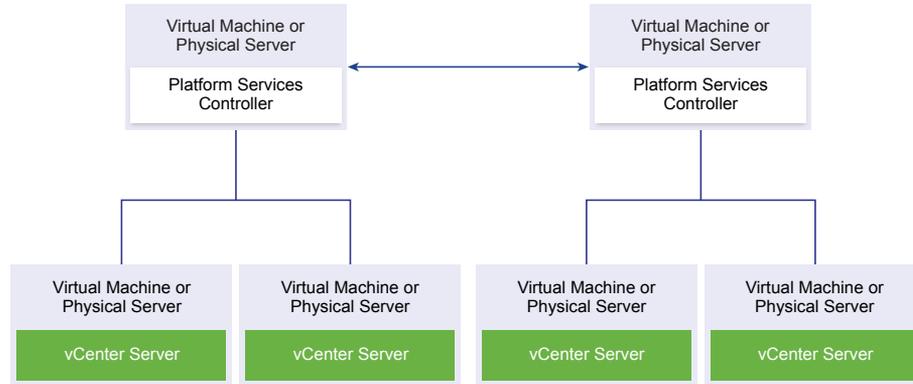
Figure 1-7. Example of Two Load Balanced Pairs of Platform Services Controller Instances Across Two Sites



Your vCenter Single Sign-On domain might span multiple sites. To ensure Platform Services Controller high availability with automatic failover throughout the domain, you must configure a separate load balancer in each site.

Platform Services Controller with No Load Balancer

Figure 1-8. Example of Two Joined Platform Services Controller Instances with No a Load Balancer



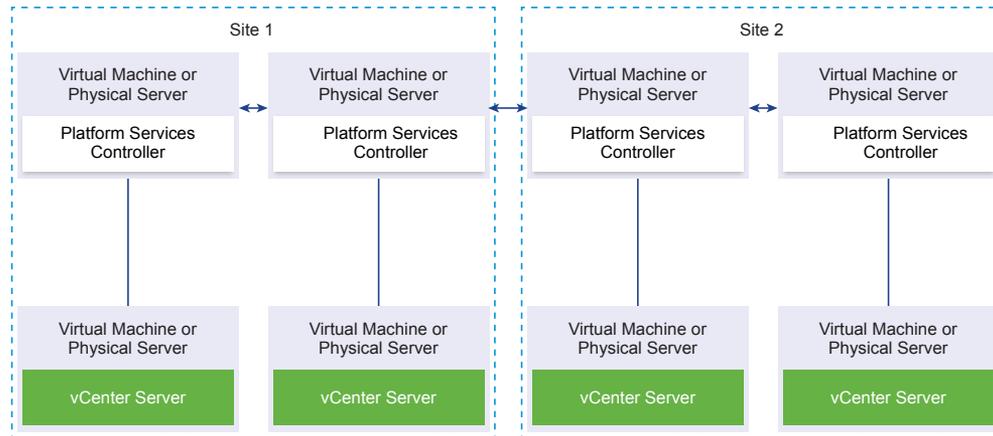
When you join two or more Platform Services Controller instances in the same site with no load balancer, you configure Platform Services Controller high availability with a manual failover for this site.

When a Platform Services Controller instance stops responding, you must manually fail over the vCenter Server instances that are registered to it. You fail over the instances by repointing them to other functional Platform Services Controller instances within the same site. See [Repoint vCenter Server to Another External Platform Services Controller](#).

Note If your vCenter Single Sign-On domain includes three or more Platform Services Controller instances, you can manually create a ring topology. A ring topology ensures Platform Services Controller reliability when one of the instances fails. To create a ring topology, run the `/usr/lib/vmware-vmdir/bin/vdcrepadmin -f createagreement` command against the first and last Platform Services Controller instance that you have deployed.

Platform Services Controller with No Load Balancer Across vCenter Single Sign-On Sites

Figure 1-9. Example of Two Joined Pairs of Platform Services Controller Instances Across Two Sites with No Load Balancer



Your vCenter Single Sign-On domain might span multiple sites. When no load balancer is available, you can manually repoint vCenter Server from a failed to a functional Platform Services Controller within the same site. See [Repoint vCenter Server to Another External Platform Services Controller](#).

Important Repointing vCenter Server between sites and domains is unsupported. If no functional Platform Services Controller instance is available in the site, you must deploy or install a new Platform Services Controller instance in this site. This new Platform Services Controller instance becomes the replication partner of the existing Platform Services Controller instance.

Enhanced Linked Mode Overview

Enhanced Linked Mode connects multiple vCenter Server systems together by using one or more Platform Services Controllers.

Enhanced Linked Mode lets you view and search across all linked vCenter Server systems and replicate roles, permissions, licenses, policies, and tags.

When you install vCenter Server or deploy the vCenter Server Appliance with an external Platform Services Controller, you must first install the Platform Services Controller. During installation of the Platform Services Controller, you can select whether to create a vCenter Single Sign-On domain or join an existing domain. You can select to join an existing vCenter Single Sign-On domain if you have already installed or deployed a Platform Services Controller instance and have created a vCenter Single Sign-On domain. When you join an existing vCenter Single Sign-On domain, the infrastructure data between the existing Platform Services Controller and the new Platform Services Controller is replicated.

With Enhanced Linked Mode, you can connect not only vCenter Server systems running on Windows but also many vCenter Server Appliances. You can also have an environment where multiple vCenter Server systems and vCenter Server Appliances are linked together.

If you install vCenter Server with an external Platform Services Controller, you first must deploy the Platform Services Controller on one virtual machines or physical server and then deploy vCenter Server on another virtual machine or physical server. While installing vCenter Server, you must select an existing external Platform Services Controller. You cannot select an existing Platform Services Controller that is a part of an embedded installation. For more information about the supported topologies, see [vCenter Server and Platform Services Controller Deployment Types](#).

About ESXi Evaluation and Licensed Modes

You can use evaluation mode to explore the entire set of features for ESXi hosts. The evaluation mode provides the set of features equal to a vSphere Enterprise Plus license. Before the evaluation mode expires, you must assign to your hosts a license that supports all the features in use.

For example, in evaluation mode, you can use vSphere vMotion technology, the vSphere HA feature, the vSphere DRS feature, and other features. If you want to continue using these features, you must assign a license that supports them.

The installable version of ESXi hosts is always installed in evaluation mode. ESXi Embedded is preinstalled on an internal storage device by your hardware vendor. It might be in evaluation mode or prelicensed.

The evaluation period is 60 days and begins when you turn on the ESXi host. At any time during the 60-day evaluation period, you can convert from licensed mode to evaluation mode. The time available in the evaluation period is decreased by the time already used.

For example, suppose that you use an ESXi host in evaluation mode for 20 days and then assign a vSphere Standard Edition license key to the host. If you set the host back in evaluation mode, you can explore the entire set of features for the host for the remaining evaluation period of 40 days.

For information about managing licensing for ESXi hosts, see the *vCenter Server and Host Management* documentation.

Installing and Setting Up ESXi

You can install and set up ESXi on your physical hardware so that it acts as a platform for virtual machines.

This section includes the following topics:

- [ESXi Requirements](#)
- [Preparing for Installing ESXi](#)
- [Installing ESXi](#)
- [Setting Up ESXi](#)
- [After You Install and Set Up ESXi](#)

ESXi Requirements

To install or upgrade ESXi, your system must meet specific hardware and software requirements.

ESXi Hardware Requirements

Make sure the host meets the minimum hardware configurations supported by ESXi6.5.

Hardware and System Resources

To install or upgrade ESXi, your hardware and system resources must meet the following requirements:

- Supported server platform. For a list of supported platforms, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.
- ESXi 6.5 requires a host machine with at least two CPU cores.
- ESXi 6.5 supports 64-bit x86 processors released after September 2006. This includes a broad range of multi-core processors. For a complete list of supported processors, see the VMware compatibility guide at <http://www.vmware.com/resources/compatibility>.
- ESXi 6.5 requires the NX/XD bit to be enabled for the CPU in the BIOS.
- ESXi 6.5 requires a minimum of 4 GB of physical RAM. It is recommended to provide at least 8 GB of RAM to run virtual machines in typical production environments.
- To support 64-bit virtual machines, support for hardware virtualization (Intel VT-x or AMD RVI) must be enabled on x64 CPUs.

- One or more Gigabit or faster Ethernet controllers. For a list of supported network adapter models, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.
- SCSI disk or a local, non-network, RAID LUN with unpartitioned space for the virtual machines.
- For Serial ATA (SATA), a disk connected through supported SAS controllers or supported on-board SATA controllers. SATA disks are considered remote, not local. These disks are not used as a scratch partition by default because they are seen as remote.

Note You cannot connect a SATA CD-ROM device to a virtual machine on an ESXi 6.5 host. To use the SATA CD-ROM device, you must use IDE emulation mode.

Storage Systems

For a list of supported storage systems, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>. For Software Fibre Channel over Ethernet (FCoE), see [Installing and Booting ESXi with Software FCoE](#).

ESXi Booting Requirements

vSphere 6.5 supports booting ESXi hosts from the Unified Extensible Firmware Interface (UEFI). With UEFI, you can boot systems from hard drives, CD-ROM drives, or USB media.

Starting with vSphere 6.5, VMware Auto Deploy supports network booting and provisioning of ESXi hosts with UEFI.

ESXi can boot from a disk larger than 2 TB if the system firmware and the firmware on any add-in card that you are using support it. See the vendor documentation.

Note Changing the boot type from legacy BIOS to UEFI after you install ESXi 6.5 might cause the host to fail to boot. In this case, the host displays an error message similar to Not a VMware boot bank. Changing the host boot type between legacy BIOS and UEFI is not supported after you install ESXi 6.5.

Storage Requirements for ESXi 6.5 Installation or Upgrade

Installing ESXi 6.5 or upgrading to ESXi 6.5 requires a boot device that is a minimum of 1 GB. When booting from a local disk, SAN or iSCSI LUN, a 5.2-GB disk is required to allow for the creation of the VMFS volume and a 4-GB scratch partition on the boot device. If a smaller disk or LUN is used, the installer attempts to allocate a scratch region on a separate local disk. If a local disk cannot be found the scratch partition, `/scratch`, is on the ESXi host ramdisk, linked to `/tmp/scratch`. You can reconfigure `/scratch` to use a separate disk or LUN. For best performance and memory optimization, do not leave `/scratch` on the ESXi host ramdisk.

To reconfigure `/scratch`, see [Set the Scratch Partition from the vSphere Web Client](#).

Due to the I/O sensitivity of USB and SD devices, the installer does not create a scratch partition on these devices. When installing or upgrading on USB or SD devices, the installer attempts to allocate a scratch region on an available local disk or datastore. If no local disk or datastore is found, `/scratch` is placed on the ramdisk. After the installation or upgrade, you should reconfigure `/scratch` to use a persistent

datastore. Although a 1GB USB or SD device suffices for a minimal installation, you should use a 4GB or larger device. The extra space is used for an expanded coredump partition on the USB/SD device. Use a high-quality USB flash drive of 16 GB or larger so that the extra flash cells can prolong the life of the boot media, but high-quality drives of 4 GB or larger are sufficient to hold the extended coredump partition. See Knowledge Base article <http://kb.vmware.com/kb/2004784>.

In Auto Deploy installations, the installer attempts to allocate a scratch region on an available local disk or datastore. If no local disk or datastore is found, /scratch is placed on ramdisk. You should reconfigure /scratch to use a persistent datastore following the installation.

For environments that boot from a SAN or use Auto Deploy, you need not allocate a separate LUN for each ESXi host. You can co-locate the scratch regions for many ESXi hosts onto a single LUN. The number of hosts assigned to any single LUN should be weighed against the LUN size and the I/O behavior of the virtual machines.

Supported Remote Management Server Models and Firmware Versions

You can use remote management applications to install or upgrade ESXi, or to manage hosts remotely.

Table 2-1. Supported Remote Management Server Models and Minimum Firmware Versions

Remote Management Server Model	Firmware Version	Java
Dell DRAC 7	1.30.30 (Build 43)	1.7.0_60-b19
Dell DRAC 6	1.54 (Build 15), 1.70 (Build 21)	1.6.0_24
Dell DRAC 5	1.0, 1.45, 1.51	1.6.0_20,1.6.0_203
Dell DRAC 4	1.75	1.6.0_23
HP ILO	1.81, 1.92	1.6.0_22, 1.6.0_23
HP ILO 2	1.8, 1.81	1.6.0_20, 1.6.0_23
HP ILO 3	1.28	1.7.0_60-b19
HP ILO 4	1.13	1.7.0_60-b19
IBM RSA 2	1.03, 1.2	1.6.0_22

Recommendations for Enhanced ESXi Performance

To enhance performance, install or upgrade ESXi on a robust system with more RAM than the minimum required and with multiple physical disks.

For ESXi system requirements, see [ESXi Hardware Requirements](#).

Table 2-2. Recommendations for Enhanced Performance

System Element	Recommendation
RAM	<p>ESXi hosts require more RAM than typical servers. Provide at least 8GB of RAM to take full advantage of ESXi features and run virtual machines in typical production environments. An ESXi host must have sufficient RAM to run concurrent virtual machines. The following examples are provided to help you calculate the RAM required by the virtual machines running on the ESXi host.</p> <p>Operating four virtual machines with Red Hat Enterprise Linux or Windows XP requires at least 3GB of RAM for baseline performance. This figure includes approximately 1024MB for the virtual machines, 256MB minimum for each operating system as recommended by vendors.</p> <p>Running these four virtual machines with 512MB RAM requires that the ESXi host have approximately 4GB RAM, which includes 2048MB for the virtual machines.</p> <p>These calculations do not take into account possible memory savings from using variable overhead memory for each virtual machine. See <i>vSphere Resource Management</i>.</p>
Dedicated Fast Ethernet adapters for virtual machines	<p>Place the management network and virtual machine networks on different physical network cards. Dedicated Gigabit Ethernet cards for virtual machines, such as Intel PRO 1000 adapters, improve throughput to virtual machines with high network traffic.</p>
Disk location	<p>Place all data that your virtual machines use on physical disks allocated specifically to virtual machines. Performance is better when you do not place your virtual machines on the disk containing the ESXi boot image. Use physical disks that are large enough to hold disk images that all the virtual machines use.</p>
VMFS5 partitioning	<p>The ESXi installer creates the initial VMFS volumes on the first blank local disk found. To add disks or modify the original configuration, use the vSphere Web Client. This practice ensures that the starting sectors of partitions are 64K-aligned, which improves storage performance.</p> <p>Note For SAS-only environments, the installer might not format the disks. For some SAS disks, it is not possible to identify whether the disks are local or remote. After the installation, you can use the vSphere Web Client to set up VMFS.</p>
Processors	<p>Faster processors improve ESXi performance. For certain workloads, larger caches improve ESXi performance.</p>
Hardware compatibility	<p>Use devices in your server that are supported by ESXi 6.5 drivers. See the <i>Hardware Compatibility Guide</i> at http://www.vmware.com/resources/compatibility.</p>

Incoming and Outgoing Firewall Ports for ESXi Hosts

The vSphere Web Client and the VMware Host Client allow you to open and close firewall ports for each service or to allow traffic from selected IP addresses.

The following table lists the firewalls for services that are installed by default. If you install other VIBs on your host, additional services and firewall ports might become available. The information is primarily for services that are visible in the vSphere Web Client but the table includes some other ports as well.

Table 2-3. Incoming Firewall Connections

Port	Protocol	Service	Description
5988	TCP	CIM Server	Server for CIM (Common Information Model).
5989	TCP	CIM Secure Server	Secure server for CIM.
427	TCP, UDP	CIM SLP	The CIM client uses the Service Location Protocol, version 2 (SLPv2) to find CIM servers.
546		DHCPv6	DHCP client for IPv6.
8301, 8302	UDP	DVSSync	DVSSync ports are used for synchronizing states of distributed virtual ports between hosts that have VMware FT record/replay enabled. Only hosts that run primary or backup virtual machines must have these ports open. On hosts that are not using VMware FT these ports do not have to be open.
902	TCP	NFC	Network File Copy (NFC) provides a file-type-aware FTP service for vSphere components. ESXi uses NFC for operations such as copying and moving data between datastores by default.
12345, 23451	UDP	vSAN Clustering Service	VMware vSAN Cluster Monitoring and Membership Directory Service. Uses UDP-based IP multicast to establish cluster members and distribute vSAN metadata to all cluster members. If disabled, vSAN does not work.
68	UDP	DHCP Client	DHCP client for IPv4.
53	UDP	DNS Client	DNS client.
8200, 8100, 8300	TCP, UDP	Fault Tolerance	Traffic between hosts for vSphere Fault Tolerance (FT).
6999	UDP	NSX Distributed Logical Router Service	NSX Virtual Distributed Router service. The firewall port associated with this service is opened when NSX VIBs are installed and the VDR module is created. If no VDR instances are associated with the host, the port does not have to be open. This service was called NSX Distributed Logical Router in earlier versions of the product.
2233	TCP	vSAN Transport	vSAN reliable datagram transport. Uses TCP and is used for vSAN storage IO. If disabled, vSAN does not work.
161	UDP	SNMP Server	Allows the host to connect to an SNMP server.
22	TCP	SSH Server	Required for SSH access.
8000	TCP	vMotion	Required for virtual machine migration with vMotion. ESXi hosts listen on port 8000 for TCP connections from remote ESXi hosts for vMotion traffic.
902, 443	TCP	vSphere Web Client	Client connections

Table 2-3. Incoming Firewall Connections (Continued)

Port	Protocol	Service	Description
8080	TCP	vsanvp	vSAN VASA Vendor Provider. Used by the Storage Management Service (SMS) that is part of vCenter to access information about vSAN storage profiles, capabilities, and compliance. If disabled, vSAN Storage Profile Based Management (SPBM) does not work.
80	TCP	vSphere Web Access	Welcome page, with download links for different interfaces.
5900 -5964	TCP	RFB protocol	
80, 9000	TCP	vSphere Update Manager	

Table 2-4. Outgoing Firewall Connections

Port	Protocol	Service	Description
427	TCP, UDP	CIM SLP	The CIM client uses the Service Location Protocol, version 2 (SLPv2) to find CIM servers.
547	TCP, UDP	DHCPv6	DHCP client for IPv6.
8301, 8302	UDP	DVSSync	DVSSync ports are used for synchronizing states of distributed virtual ports between hosts that have VMware FT record/replay enabled. Only hosts that run primary or backup virtual machines must have these ports open. On hosts that are not using VMware FT these ports do not have to be open.
44046, 31031	TCP	HBR	Used for ongoing replication traffic by vSphere Replication and VMware Site Recovery Manager.
902	TCP	NFC	Network File Copy (NFC) provides a file-type-aware FTP service for vSphere components. ESXi uses NFC for operations such as copying and moving data between datastores by default.
9	UDP	WOL	Used by Wake on LAN.
12345 23451	UDP	vSAN Clustering Service	Cluster Monitoring, Membership, and Directory Service used by vSAN.
68	UDP	DHCP Client	DHCP client.
53	TCP, UDP	DNS Client	DNS client.
80, 8200, 8100, 8300	TCP, UDP	Fault Tolerance	Supports VMware Fault Tolerance.
3260	TCP	Software iSCSI Client	Supports software iSCSI.
6999	UDP	NSX Distributed Logical Router Service	The firewall port associated with this service is opened when NSX VIBs are installed and the VDR module is created. If no VDR instances are associated with the host, the port does not have to be open.

Table 2-4. Outgoing Firewall Connections (Continued)

Port	Protocol	Service	Description
5671	TCP	rabbitmqproxy	A proxy running on the ESXi host. This proxy allows applications that are running inside virtual machines to communicate with the AMQP brokers that are running in the vCenter network domain. The virtual machine does not have to be on the network, that is, no NIC is required. Ensure that outgoing connection IP addresses include at least the brokers in use or future. You can add brokers later to scale up.
2233	TCP	vSAN Transport	Used for RDT traffic (Unicast peer to peer communication) between vSAN nodes.
8000	TCP	vMotion	Required for virtual machine migration with vMotion.
902	UDP	VMware vCenter Agent	vCenter Server agent.
8080	TCP	vsanvp	Used for vSAN Vendor Provider traffic.
9080	TCP	I/O Filter Service	Used by the I/O Filters storage feature

Table 2-5. Firewall Ports for Services That Are Not Visible in the UI by Default

Port	Protocol	Service	Comment
5900 -5964	TCP	RFB protocol	The RFB protocol is a simple protocol for remote access to graphical user interfaces.
8889	TCP	OpenWSMAN Daemon	Web Services Management (WS-Management is a DMTF open standard for the management of servers, devices, applications, and Web services.

Required Free Space for System Logging

If you used Auto Deploy to install your ESXi 6.5 host, or if you set up a log directory separate from the default location in a scratch directory on the VMFS volume, you might need to change your current log size and rotation settings to ensure that enough space is available for system logging .

All vSphere components use this infrastructure. The default values for log capacity in this infrastructure vary, depending on the amount of storage available and on how you have configured system logging. Hosts that are deployed with Auto Deploy store logs on a RAM disk, which means that the amount of space available for logs is small.

If your host is deployed with Auto Deploy, reconfigure your log storage in one of the following ways:

- Redirect logs over the network to a remote collector.
- Redirect logs to a NAS or NFS store.

If you redirect logs to non-default storage, such as a NAS or NFS store, you might also want to reconfigure log sizing and rotations for hosts that are installed to disk.

You do not need to reconfigure log storage for ESXi hosts that use the default configuration, which stores logs in a scratch directory on the VMFS volume. For these hosts, ESXi 6.5 configures logs to best suit your installation, and provides enough space to accommodate log messages.

Table 2-6. Recommended Minimum Size and Rotation Configuration for hostd, vpxa, and fdm Logs

Log	Maximum Log File Size	Number of Rotations to Preserve	Minimum Disk Space Required
Management Agent (hostd)	10 MB	10	100 MB
VirtualCenter Agent (vpxa)	5 MB	10	50 MB
vSphere HA agent (Fault Domain Manager, fdm)	5 MB	10	50 MB

For information about setting up a remote log server, see [Configure Syslog on ESXi Hosts](#).

VMware Host Client System Requirements

Make sure that your browser supports the VMware Host Client.

The following guest operating systems and Web browser versions are supported for the VMware Host Client.

Supported Browsers	Mac OS	Windows	Linux
Google Chrome	50+	50+	50+
Mozilla Firefox	45+	45+	45+
Microsoft Internet Explorer	N/A	11+	N/A
Microsoft Edge	N/A	38+	N/A
Safari	9.0+	N/A	N/A

Preparing for Installing ESXi

Before you install ESXi, determine the installation option that is suitable for your environment and prepare for the installation process.

Download the ESXi Installer

Download the installer for ESXi.

Prerequisites

Create a My VMware account at <https://my.vmware.com/web/vmware/>.

Procedure

- 1 Download the ESXi installer from the VMware Web site at <https://my.vmware.com/web/vmware/downloads>.

ESXi is listed under Datacenter & Cloud Infrastructure.

2 Confirm that the md5sum is correct.

See the VMware Web site topic Using MD5 Checksums at <http://www.vmware.com/download/md5.html>.

Options for Installing ESXi

ESXi can be installed in several ways. To ensure the best vSphere deployment, understand the options thoroughly before beginning the installation.

ESXi installations are designed to accommodate a range of deployment sizes.

Depending on the installation method you choose, different options are available for accessing the installation media and booting the installer.

Interactive ESXi Installation

Interactive installations are recommended for small deployments of fewer than five hosts.

You boot the installer from a CD or DVD, from a bootable USB device, or by PXE booting the installer from a location on the network. You follow the prompts in the installation wizard to install ESXi to disk. See [Installing ESXi Interactively](#).

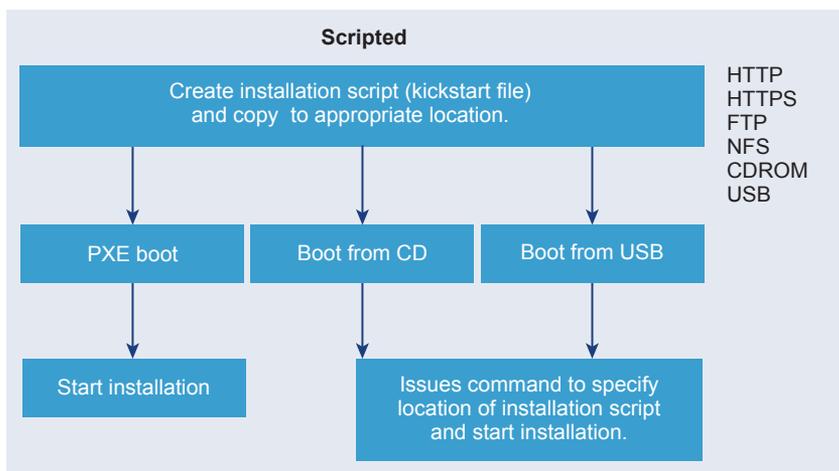
Scripted ESXi Installation

Running a script is an efficient way to deploy multiple ESXi hosts with an unattended installation.

The installation script contains the host configuration settings. You can use the script to configure multiple hosts with the same settings. See [Installing or Upgrading Hosts by Using a Script](#).

The installation script must be stored in a location that the host can access by HTTP, HTTPS, FTP, NFS, CDROM, or USB. You can PXE boot the ESXi installer or boot it from a CD/DVD or USB drive.

Figure 2-1. Scripted Installation



vSphere Auto Deploy ESXi Installation

vSphere 5.x and later provide several ways to install ESXi with vSphere Auto Deploy.

vSphere Auto Deploy can provision hundreds of physical hosts with ESXi software. You can specify the image to deploy and the hosts to provision with the image. Optionally, you can specify host profiles to apply to the hosts, a vCenter Server location (datacenter, folder, or cluster), and script bundle for each host.

vCenter Server makes ESXi updates and patches available for download in the form of an image profile. The host configuration is provided in the form of a host profile. You can create host profiles by using the vSphere Web Client. You can create custom image profiles by using vSphere ESXi Image Builder. See [Customizing Installations with vSphere ESXi Image Builder](#) and *vSphere Host Profiles*.

When you provision hosts by using vSphere Auto Deploy, vCenter Server loads the ESXi image directly into the host memory. vSphere Auto Deploy does not store the ESXi state on the host disk. The vSphere Auto Deploy server continues to provision this host every time the host boots.

You can also use vSphere Auto Deploy to install an ESXi host, and set up a host profile that causes the host to store the ESXi image and configuration on the local disk, a remote disk, or a USB drive. Subsequently, the ESXi host boots from this local image and vSphere Auto Deploy no longer provisions the host. This process is similar to performing a scripted installation. With a scripted installation, the script provisions a host and the host then boots from disk. For this case, vSphere Auto Deploy provisions a host and the host then boots from disk. For more information, see [Using vSphere Auto Deploy for Stateless Caching and Stateful Installs](#).

Media Options for Booting the ESXi Installer

The ESXi installer must be accessible to the system on which you are installing ESXi.

The following boot media are supported for the ESXi installer:

- Boot from a CD/DVD. See [Download and Burn the ESXi Installer ISO Image to a CD or DVD](#).
- Boot from a USB flash drive. See [Format a USB Flash Drive to Boot the ESXi Installation or Upgrade](#).
- PXE boot from the network. [PXE Booting the ESXi Installer](#)
- Boot from a remote location using a remote management application. See [Using Remote Management Applications](#)

Download and Burn the ESXi Installer ISO Image to a CD or DVD

If you do not have an ESXi installation CD/DVD, you can create one.

You can also create an installer ISO image that includes a custom installation script. See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#).

Procedure

- 1 Download the ESXi installer from the VMware Web site at <https://my.vmware.com/web/vmware/downloads>.

ESXi is listed under Datacenter & Cloud Infrastructure.

- 2 Confirm that the md5sum is correct.

See the VMware Web site topic Using MD5 Checksums at <http://www.vmware.com/download/md5.html>.

- 3 Burn the ISO image to a CD or DVD.

Format a USB Flash Drive to Boot the ESXi Installation or Upgrade

You can format a USB flash drive to boot the ESXi installation or upgrade.

The instructions in this procedure assume that the USB flash drive is detected as `/dev/sdb`.

Note The `ks.cfg` file that contains the installation script cannot be located on the same USB flash drive that you are using to boot the installation or upgrade.

Prerequisites

- Linux machine with superuser access to it
- USB flash drive that can be detected by the Linux machine
- The ESXi ISO image, `VMware-VMvisor-Installer-version_number-build_number.x86_64.iso`, which includes the `isolinux.cfg` file
- Syslinux package

Procedure

- 1 If your USB flash drive is not detected as `/dev/sdb`, or you are not sure how your USB flash drive is detected, determine how it is detected.

- a At the command line, run the command for displaying the current log messages.

```
tail -f /var/log/messages
```

- b Plug in your USB flash drive.

You see several messages that identify the USB flash drive in a format similar to the following message.

```
Oct 25 13:25:23 ubuntu kernel: [ 712.447080] sd 3:0:0:0: [sdb] Attached SCSI removable disk
```

In this example, `sdb` identifies the USB device. If your device is identified differently, use that identification, in place of `sdb`.

2 Create a partition table on the USB flash device.

```
/sbin/fdisk /dev/sdb
```

- a Enter `d` to delete partitions until they are all deleted.
- b Enter `n` to create a primary partition 1 that extends over the entire disk.
- c Enter `t` to set the type to an appropriate setting for the FAT32 file system, such as `c`.
- d Enter `a` to set the active flag on partition 1.
- e Enter `p` to print the partition table.

The result should be similar to the following message.

```
Disk /dev/sdb: 2004 MB, 2004877312 bytes
255 heads, 63 sectors/track, 243 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1            1           243     1951866   c   W95 FAT32 (LBA)
```

- f Enter `w` to write the partition table and exit the program.

3 Format the USB flash drive with the Fat32 file system.

```
/sbin/mkfs.vfat -F 32 -n USB /dev/sdb1
```

4 Install the Syslinux bootloader on the USB flash drive.

The locations of the Syslinux executable file and the `mbr.bin` file might vary for the different Syslinux versions. For example, if you downloaded Syslinux 6.02, run the following commands.

```
/usr/bin/syslinux /dev/sdb1
cat /usr/lib/syslinux/mbr/mbr.bin > /dev/sdb
```

5 Create a destination directory and mount the USB flash drive to it.

```
mkdir /usbdisk
mount /dev/sdb1 /usbdisk
```

6 Create a destination directory and mount the ESXi installer ISO image to it.

```
mkdir /esxi_cdrom
mount -o loop VMware-VMvisor-Installer-6.x.x-XXXXXX.x86_64.iso /esxi_cdrom
```

7 Copy the contents of the ISO image to the USB flash drive.

```
cp -r /esxi_cdrom/* /usbdisk
```

- 8 Rename the `isolinux.cfg` file to `syslinux.cfg`.

```
mv /usbdisk/isolinux.cfg /usbdisk/syslinux.cfg
```

- 9 In the `/usbdisk/syslinux.cfg` file, edit the `APPEND -c boot.cfg` line to `APPEND -c boot.cfg -p 1`.

- 10 Unmount the USB flash drive.

```
umount /usbdisk
```

- 11 Unmount the installer ISO image.

```
umount /esxi_cdrom
```

The USB flash drive can boot the ESXi installer.

Create a USB Flash Drive to Store the ESXi Installation Script or Upgrade Script

You can use a USB flash drive to store the ESXi installation script or upgrade script that is used during scripted installation or upgrade of ESXi.

When multiple USB flash drives are present on the installation machine, the installation software searches for the installation or upgrade script on all attached USB flash drives.

The instructions in this procedure assume that the USB flash drive is detected as `/dev/sdb`.

Note Do not store the `ks` file containing the installation or upgrade script on the same USB flash drive that you are using to boot the installation or upgrade.

Prerequisites

- Linux machine
- ESXi installation or upgrade script, the `ks.cfg` kickstart file
- USB flash drive

Procedure

- 1 Attach the USB flash drive to a Linux machine that has access to the installation or upgrade script.

- 2 Create a partition table.

```
/sbin/fdisk /dev/sdb
```

- a Type `d` to delete partitions until they are all deleted.
- b Type `n` to create primary partition 1 that extends over the entire disk.
- c Type `t` to set the type to an appropriate setting for the FAT32 file system, such as `c`.

- d Type `p` to print the partition table.

The result should be similar to the following text:

```
Disk /dev/sdb: 2004 MB, 2004877312 bytes
255 heads, 63 sectors/track, 243 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1             1           243     1951866   c   W95 FAT32 (LBA)
```

- e Type `w` to write the partition table and quit.

- 3 Format the USB flash drive with the Fat32 file system.

```
/sbin/mkfs.vfat -F 32 -n USB /dev/sdb1
```

- 4 Mount the USB flash drive.

```
mount /dev/sdb1 /usbdisk
```

- 5 Copy the ESXi installation script to the USB flash drive.

```
cp ks.cfg /usbdisk
```

- 6 Unmount the USB flash drive.

The USB flash drive contains the installation or upgrade script for ESXi.

What to do next

When you boot the ESXi installer, point to the location of the USB flash drive for the installation or upgrade script. See [Enter Boot Options to Start an Installation or Upgrade Script](#) and [PXELINUX Configuration Files](#).

Create an Installer ISO Image with a Custom Installation or Upgrade Script

You can customize the standard ESXi installer ISO image with your own installation or upgrade script. This customization enables you to perform a scripted, unattended installation or upgrade when you boot the resulting installer ISO image.

See also [About Installation and Upgrade Scripts](#) and [About the boot.cfg File](#).

Prerequisites

- Linux machine
- The ESXi ISO image `VMware-VMvisor-Installer-6.x.x-XXXXXX.x86_64.iso`, where `6.x.x` is the version of ESXi you are installing, and `XXXXXX` is the build number of the installer ISO image
- Your custom installation or upgrade script, the `ks_cust.cfg` kickstart file

Procedure

- 1 Download the ESXi ISO image from the VMware Web site.

- 2 Mount the ISO image in a folder:

```
mount -o loop VMware-VMvisor-Installer-6.x.x-XXXXXX.x86_64.iso /esxi_cdrom_mount
```

XXXXXX is the ESXi build number for the version that you are installing or upgrading to.

- 3 Copy the contents of cdrom to another folder:

```
cp -r /esxi_cdrom_mount /esxi_cdrom
```

- 4 Copy the kickstart file to /esxi_cdrom.

```
cp ks_cust.cfg /esxi_cdrom
```

- 5 (Optional) Modify the boot.cfg file to specify the location of the installation or upgrade script by using the kernelopt option.

You must use uppercase characters to provide the path of the script, for example,

```
kernelopt=runweasel ks=cdrom:/KS_CUST.CFG
```

The installation or upgrade becomes completely automatic, without the need to specify the kickstart file during the installation or upgrade.

- 6 Recreate the ISO image using the mkisofs or the genisoimage command.

Command	Syntax
mkisofs	mkisofs -relaxed-filenames -J -R -o custom_esxi.iso -b isolinux.bin -c boot.cat -no-emul-boot -boot-load-size 4 -boot-info-table -eltorito-alt-boot -eltorito-platform efi -b efiboot.img -no-emul-boot /esxi_cdrom
genisoimage	genisoimage -relaxed-filenames -J -R -o custom_esxi.iso -b isolinux.bin -c boot.cat -no-emul-boot -boot-load-size 4 -boot-info-table -eltorito-alt-boot -e efiboot.img -no-emul-boot /esxi_cdrom

You can use this ISO image for regular boot or UEFI secure boot.

The ISO image includes your custom installation or upgrade script.

What to do next

Install ESXi from the ISO image.

PXE Booting the ESXi Installer

You can use the preboot execution environment (PXE) to boot a host. Starting with vSphere 6.0, you can PXE boot the ESXi installer from a network interface on hosts with legacy BIOS or using UEFI.

ESXi is distributed in an ISO format that is designed to install to flash memory or to a local hard drive. You can extract the files and boot by using PXE.

PXE uses Dynamic Host Configuration Protocol (DHCP) and Trivial File Transfer Protocol (TFTP) to boot an operating system over a network.

PXE booting requires some network infrastructure and a machine with a PXE-capable network adapter. Most machines that can run ESXi have network adapters that can PXE boot.

Note PXE booting with legacy BIOS firmware is possible only over IPv4. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.

Sample DHCP Configurations

To PXE boot the ESXi installer, the DHCP server must send the address of the TFTP server and the filename of the initial boot loader to the ESXi host.

When the target machine first boots, it broadcasts a packet across the network requesting information to boot itself. The DHCP server responds. The DHCP server must be able to determine whether the target machine is allowed to boot and the location of the initial boot loader binary, typically a file on a TFTP server.

Caution Do not set up a second DHCP server if your network already has one. If multiple DHCP servers respond to DHCP requests, machines can obtain incorrect or conflicting IP addresses, or can fail to receive the proper boot information. Talk to a network administrator before setting up a DHCP server. For support on configuring DHCP, contact your DHCP server vendor.

Many DHCP servers can PXE boot hosts. If you are using a version of DHCP for Microsoft Windows, see the DHCP server documentation to determine how to pass the `next-server` and `filename` arguments to the target machine.

Example of Booting Using TFTP with IPv4

This example shows how to configure an ISC DHCP server to boot ESXi using a TFTP server at IPv4 address `xxx.xxx.xxx.xxx`.

```
#
# ISC DHCP server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option client-system-arch code 93 = unsigned integer 16;
class "pxeclients" {
    match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
    next-server xxx.xxx.xxx.xxx;
    if option client-system-arch = 00:07 or option client-system-arch = 00:09 {
        filename = "mboot.efi";
    } else {
        filename = "pxelinux.0";
    }
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `pxelinux.0` or `mboot.efi` binary file on the TFTP server.

Example of Booting Using TFTP with IPv6

This example shows how to configure an ISC DHCPv6 server to boot ESXi using a TFTP server at IPv6 address `xxxx:xxxx:xxxx:xxxx::xxxx`.

```
#
# ISC DHCPv6 server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option dhcp6.bootfile-url code 59 = string;
option dhcp6.bootfile-url "tftp://[xxxx:xxxx:xxxx:xxxx::xxxx]/mboot.efi";
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `mboot.efi` binary file on the TFTP server.

Example of Booting Using HTTP with IPv4

This example shows how to configure an ISC DHCP server to boot ESXi using a Web server at IPv4 address `xxx.xxx.xxx.xxx`. The example uses gPXE LINUX for legacy BIOS hosts and iPXE for UEFI hosts.

```
#
# ISC DHCPv6 server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option client-system-arch code 93 = unsigned integer 16;
class "pxeclients" {
    match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
    next-server xxx.xxx.xxx.xxx;
    if option client-system-arch = 00:07 or option client-system-arch = 00:09 {
        if exists user-class and option user-class = "iPXE" {
            # Instruct iPXE to load mboot.efi as secondary bootloader
            filename = "mboot.efi";
        } else {
            # Load the snponly.efi configuration of iPXE as initial bootloader
            filename = "snponly.efi";
        }
    } else {
        filename "gpxelinux.0";
    }
}
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `gpxelinux.0` or `snponly.efi` binary file on the TFTP server. In the UEFI case, iPXE then asks the DHCP server for the next file to load, and this time the server returns `mboot.efi` as the filename.

Example of Booting Using HTTP with IPv6

This example shows how to configure an ISC DHCPv6 server to boot ESXi using a TFTP server at IPv6 address `xxxx:xxxx:xxxx:xxxx::xxxx`.

```
#
# ISC DHCPv6 server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;

option dhcp6.bootfile-url code 59 = string;
if exists user-class and option user-class = "iPXE" {
    # Instruct iPXE to load mboot.efi as secondary bootloader
    option dhcp6.bootfile-url "tftp://[xxxx:xxxx:xxxx:xxxx::xxxx]/mboot.efi";
} else {
    # Load the snponly.efi configuration of iPXE as initial bootloader
    option dhcp6.bootfile-url "tftp://[xxxx:xxxx:xxxx:xxxx::xxxx]/snponly.efi";
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `snponly.efi` (iPXE) binary file on the TFTP server. iPXE then asks the DHCP server for the next file to load, and this time the server returns `mboot.efi` as the filename.

PXELINUX Configuration Files

You need a PXELINUX configuration file to boot the ESXi installer on a legacy BIOS system. The configuration file defines the menu displayed to the target ESXi host as it boots up and contacts the TFTP server for all SYSLINUX configurations, including PXELINUX and gPXELINUX.

This section gives general information about PXELINUX configuration files. For examples, see [Sample DHCP Configurations](#).

For syntax details, see the SYSLINUX web site at <http://www.syslinux.org/>.

Required Files

In the PXE configuration file, you must include paths to the following files:

- `mboot.c32` is the boot loader.
- `boot.cfg` is the boot loader configuration file.

See [About the boot.cfg File](#)

File Name for the PXE Configuration File

For the file name of the PXE configuration file, select one of the following options:

- `01-mac_address_of_target_ESXi_host`. For example, `01-23-45-67-89-0a-bc`
- The target ESXi host IP address in hexadecimal notation.
- `default`

The initial boot file, `pxelinux.0` or `gpxelinux.0`, tries to load a PXE configuration file in the following order:

- 1 It tries with the MAC address of the target ESXi host, prefixed with its ARP type code, which is 01 for Ethernet.
- 2 If that attempt fails, it tries with the hexadecimal notation of target ESXi system IP address.
- 3 Ultimately, it tries to load a file named `default`.

File Location for the PXE Configuration File

Save the file in `/tftpboot/pxelinux.cfg/` on the TFTP server.

For example, you might save the file on the TFTP server at `/tftpboot/pxelinux.cfg/01-00-21-5a-ce-40-f6`. The MAC address of the network adapter on the target ESXi host is 00-21-5a-ce-40-f6.

PXE Boot Background Information

Understanding the PXE boot process can help you during troubleshooting.

TFTP Server

Trivial File Transfer Protocol (TFTP) is similar to the FTP service, and is typically used only for network booting systems or loading firmware on network devices such as routers. TFTP is available on Linux and Windows.

- Most Linux distributions include a copy of the `tftp-hpa` server. If you require a supported solution, purchase a supported TFTP server from your vendor of choice. You can also acquire a TFTP server from one of the packaged appliances on the VMware Marketplace.
- If your TFTP server will run on a Microsoft Windows host, use `ttfpd32` version 2.11 or later. See <http://ttfpd32.jounin.net/>.

SYSLINUX, PXELINUX, and gPXELINUX

If you are using PXE in a legacy BIOS environment, you need to understand the different boot environments.

- SYSLINUX is an open source boot environment for machines that run legacy BIOS firmware. The ESXi boot loader for BIOS systems, `mboot.c32`, runs as a SYSLINUX plugin. You can configure SYSLINUX to boot from several types of media, including disk, ISO image, and network. You can find the SYSLINUX package at <http://www.kernel.org/pub/linux/utils/boot/syslinux/>.
- PXELINUX is a SYSLINUX configuration for booting from a TFTP server according to the PXE standard. If you use PXELINUX to boot the ESXi installer, the `pxelinux.0` binary file, `mboot.c32`, the configuration file, the kernel, and other files are all transferred by TFTP.

- gPXELINUX is a hybrid configuration that includes both PXELINUX and gPXE and supports booting from a Web server. gPXELINUX is part of the SYSLINUX package. If you use gPXELINUX to boot the ESXi installer, only the `gpxelinux.0` binary file, `mboot.c32`, and the configuration file are transferred via TFTP. The remaining files are transferred via HTTP. HTTP is typically faster and more reliable than TFTP, especially for transferring large amounts of data on a heavily loaded network.

Note VMware currently builds the `mboot.c32` plugin to work with SYSLINUX version 3.86 and tests PXE booting only with that version. Other versions are likely to be incompatible. This is not a statement of limited support. For support of third-party agents that you use to set up your PXE booting infrastructure, contact the vendor.

UEFI PXE and iPXE

Most UEFI firmware natively includes PXE support that allows booting from a TFTP server. The firmware can directly load the ESXi boot loader for UEFI systems, `mboot.efi`. Additional software such as PXELINUX is not required.

iPXE can also be useful for UEFI systems that do not include PXE in firmware and for older UEFI systems with bugs in their PXE support. For such cases you can try installing iPXE on a USB flash drive and booting from there.

Note Apple Macintosh products do not include PXE boot support. They include support for network booting via an Apple-specific protocol instead.

Alternative Approaches to PXE Booting

Alternative approaches to PXE booting different software on different hosts are also possible, for example:

- Configuring the DHCP server to provide different initial boot loader filenames to different hosts depending on MAC address or other criteria. See your DHCP server's documentation.
- Approaches using iPXE as the initial bootloader with an iPXE configuration file that selects the next bootloader based on the MAC address or other criteria.

Installing and Booting ESXi with Software FCoE

You can install and boot ESXi from an FCoE LUN using VMware software FCoE adapters and network adapters with FCoE offload capabilities. Your host does not require a dedicated FCoE HBA.

See the *vSphere Storage* documentation for information about installing and booting ESXi with software FCoE.

Using Remote Management Applications

Remote management applications allow you to install ESXi on servers that are in remote locations.

Remote management applications supported for installation include HP Integrated Lights-Out (iLO), Dell Remote Access Card (DRAC), IBM management module (MM), and Remote Supervisor Adapter II (RSA II). For a list of currently supported server models and remote management firmware versions, see [Supported Remote Management Server Models and Firmware Versions](#). For support on remote management applications, contact the vendor.

You can use remote management applications to do both interactive and scripted installations of ESXi remotely.

If you use remote management applications to install ESXi, the virtual CD might encounter corruption problems with systems or networks operating at peak capacity. If a remote installation from an ISO image fails, complete the installation from the physical CD media.

Customizing Installations with vSphere ESXi Image Builder

You can use VMware vSphere® ESXi™ Image Builder CLI to create ESXi installation images with a customized set of updates, patches, and drivers.

You can use vSphere ESXi Image Builder with the vSphere Web Client or with PowerCLI to create an ESXi installation image with a customized set of ESXi updates and patches. You can also include third-party network or storage drivers that are released between vSphere releases.

You can deploy an ESXi image created with vSphere ESXi Image Builder in either of the following ways:

- By burning it to an installation DVD.
- Through vCenter Server, using the Auto Deploy feature.

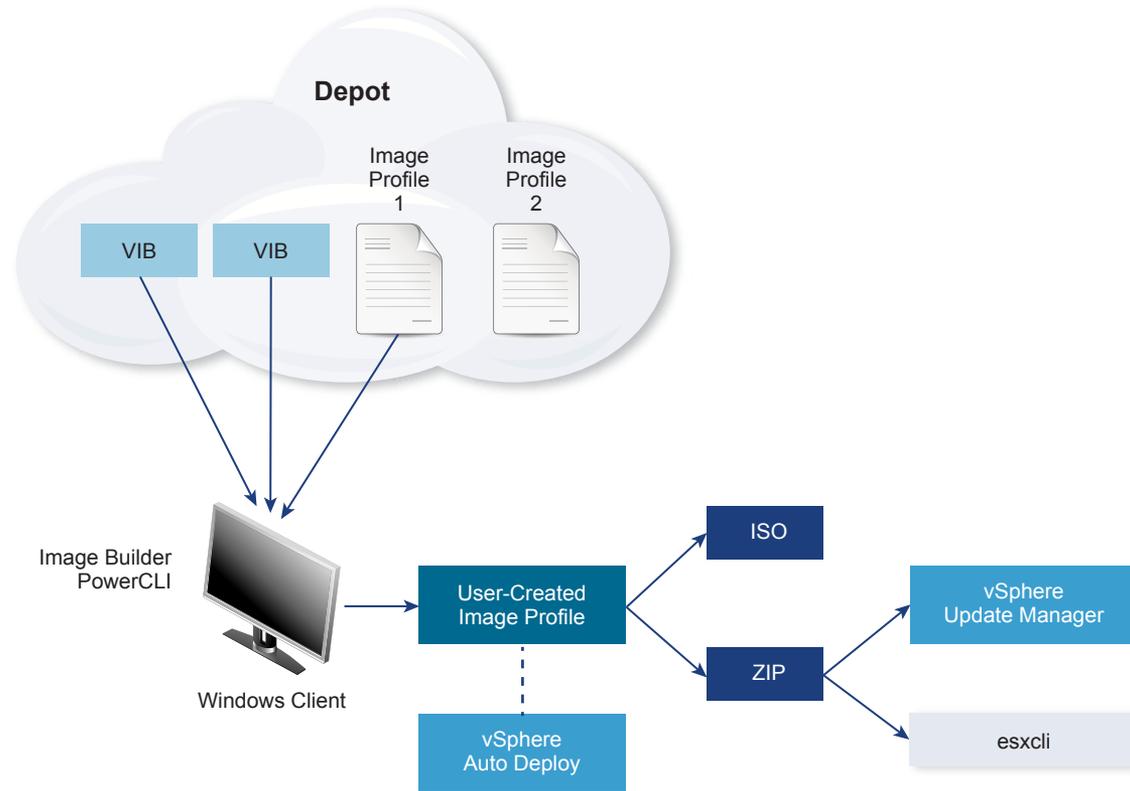
Understanding vSphere ESXi Image Builder

You can use the VMware vSphere® ESXi™ Image Builder CLI to manage software depots, image profiles, and software packages (VIBs). Image profiles and VIBs specify the software you want to use during installation or upgrade of an ESXi host.

vSphere ESXi Image Builder Overview

vSphere ESXi Image Builder lets you manage vSphere image profiles and VIBs.

VIBs are software packages, and image profiles contain a set of software packages. See [Software Depots and Their Components](#).

Figure 2-2. Image Builder Architecture

You use vSphere ESXi Image Builder cmdlets for managing the software to deploy to your ESXi hosts in several different situations.

Table 2-7. Cases Where You Can Use vSphere ESXi Image Builder

Use Case for vSphere ESXi Image Builder	Description
Create image profiles for use by vSphere Auto Deploy	Use vSphere ESXi Image Builder to create an image profile that defines the VIBs that vSphere Auto Deploy uses to provision hosts.
Add custom third-party drivers to existing image profile and export to ISO or bundle	When you add third-party driver or extension custom VIBs to your ESXi hosts, use vSphere ESXi Image Builder to clone the base image provided by VMware, add the custom VIBs, and export to ISO or to offline bundle ZIP file.
Perform upgrades	If you upgrade from a 4.0 or 4.1 system that includes custom extensions or drivers, you can use vSphere ESXi Image Builder to create an image profile that includes the vSphere 5 base VIB. You can create vSphere 5 VIBs for the custom extensions and add those VIBs to the base VIB. Export the custom image profile to an ISO you can install or to a ZIP that you can use with vSphere Update Manager.
Create custom images with reduced footprint	If you require a minimal footprint image, you can clone the ESXi base image profile and remove VIBs using vSphere ESXi Image Builder.

The vSphere ESXi Image Builder cmdlets take image profiles and VIBs as input and produce various outputs.

Table 2-8. Input and Output to the vSphere ESXi Image Builder Cmdlets

Parameter	Description
Input	Image profiles and VIBs that are located in a software depot are used as input to PowerCLI cmdlets running on a Windows client.
Output	PowerCLI cmdlets create custom image profiles that can be exported to an ISO image or an offline depot ZIP file. ISO images are used for installation. The ZIP depot can be used by Update Manager or by <code>esxcli</code> software commands to update or install images. Image profiles are also used in vSphere Auto Deploy rules to customize the software to provision ESXi hosts with.

Watch the video "Using Image Builder CLI" for information about vSphere ESXi Image Builder:



Using Image Builder CLI (http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_using_image_builder_cli)

Software Depots and Their Components

Understanding how depots, profiles, and VIBs are structured and where you can use them is a prerequisite for in-memory installation of a custom ESXi ISO, for provisioning ESXi hosts using vSphere Auto Deploy, and for certain custom upgrade operations.

The following technical terms are used throughout the vSphere documentation set in discussions of installation and upgrade tasks.

VIB

A VIB is an ESXi software package. VMware and its partners package solutions, drivers, CIM providers, and applications that extend the ESXi platform as VIBs. VIBs are available in software depots. You can use VIBs to create and customize ISO images or to upgrade ESXi hosts by installing VIBs asynchronously onto the hosts.

See [SoftwarePackage Object Properties](#).

Image Profile

An image profile defines an ESXi image and consists of VIBs. An image profile always includes a base VIB, and might include more VIBs. You examine and define an image profile by using vSphere ESXi Image Builder.

See [ImageProfile Object Properties](#).

Software Depot

A software depot is a collection of VIBs and image profiles. The software depot is a hierarchy of files and folders and can be available through an HTTP URL (online depot) or a ZIP file (offline depot). VMware and VMware partners make depots available. Companies with large VMware installations might create internal depots to provision ESXi hosts with vSphere Auto Deploy, or to export an ISO for ESXi installation.

vSphere ESXi Image Builder Cmdlets Overview

vSphere ESXi Image Builder cmdlets allow you to manage image profiles and VIBs.

vSphere ESXi Image Builder includes the following cmdlets.

Note When you run vSphere ESXi Image Builder cmdlets, provide all parameters on the command line when you invoke the cmdlet. Supplying parameters in interactive mode is not recommended.

Run `Get-Help cmdlet_name` at the PowerCLI prompt for detailed reference information.

Table 2-9. vSphere ESXi Image Builder Cmdlets

Cmdlet	Description
<code>Add-EsxSoftwareDepot</code>	Adds the software depot or ZIP file at the specified location to your current environment. Downloads metadata from the depot and analyzes VIBs for dependencies.
<code>Remove-EsxSoftwareDepot</code>	Disconnects from the specified software depot.
<code>Get-EsxSoftwareDepot</code>	Returns a list of software depots that are in the current environment. If you want to examine and manage image profiles and VIBs, you must first add the corresponding software depot to your environment.
<code>Get-EsxSoftwarePackage</code>	Returns a list of software package objects (VIBs). Use this cmdlet's options to filter the results.
<code>Get-EsxImageProfile</code>	Returns an array of <code>ImageProfile</code> objects from all currently added depots.
<code>New-EsxImageProfile</code>	Creates a new image profile. In most cases, creating a new profile by cloning an existing profile is recommended. See Clone an Image Profile .
<code>Set-EsxImageProfile</code>	Modifies a local <code>ImageProfile</code> object and performs validation tests on the modified profile. The cmdlet returns the modified object but does not persist it.
<code>Export-EsxImageProfile</code>	Exports an image profile as either an ESXi ISO image for ESXi installation, or as a ZIP file.
<code>Compare-EsxImageProfile</code>	Returns an <code>ImageProfileDiff</code> structure that shows whether the two profiles have the same VIB list and acceptance level. See Acceptance Levels .
<code>Remove-EsxImageProfile</code>	Removes the image profile from the software depot.
<code>Add-EsxSoftwarePackage</code>	Adds one or more new packages (VIBs) to an existing image profile.
<code>Remove-EsxSoftwarePackage</code>	Removes one or more packages (VIBs) from an image profile.

Image Profiles

Image profiles define the set of VIBs that an ESXi installation or update process uses. Image profiles apply to hosts provisioned with vSphere Auto Deploy and to other ESXi 5.x hosts. You define and manipulate image profiles with vSphere ESXi Image Builder.

Image Profile Requirements

You can create a custom image profile from scratch or clone an existing profile and add or remove VIBs. A profile must meet the following requirements to be valid.

- Each image profile must have a unique name and vendor combination.

- Each image profile has an acceptance level. When you add a VIB to an image profile with an vSphere ESXi Image Builder cmdlet, Image Builder checks that the VIB matches the acceptance level defined for the profile.
- You cannot remove VIBs that are required by other VIBs.
- You cannot include two versions of the same VIB in an image profile. When you add a new version of a VIB, the new version replaces the existing version of the VIB.

Image Profile Validation

An image profile and its VIBs must meet several criteria to be valid.

- Image profiles must contain at least one base VIB and one bootable kernel module.
- If any VIB in the image profile depends on another VIB, that other VIB must also be included in the image profile. VIB creators store that information in the SoftwarePackage object's Depends property.
- VIBs must not conflict with each other. VIB creators store conflict information in the SoftwarePackage object's Conflicts property.
- Two VIBs with the same name, but two different versions, cannot coexist. When you add a new version of a VIB, the new version replaces the existing version of the VIB.
- No acceptance level validation issues exist.

When you make a change to an image profile, vSphere ESXi Image Builder checks that the change does not invalidate the profile.

Dependency Validation

When you add or remove a VIB, vSphere ESXi Image Builder checks that package dependencies are met. Each SoftwarePackage object includes a Depends property that specifies a list of other VIBs that VIB depends on. See [Structure of ImageProfile, SoftwarePackage, and ImageProfileDiff Objects](#)

Acceptance Level Validation

vSphere ESXi Image Builder performs acceptance level validation each time an image profile is created or changed. vSphere ESXi Image Builder checks the acceptance level of VIBs in the image profile against the minimum allowed acceptance level of the profile. The acceptance level of the VIB is also validated each time the signature of a VIB is validated.

VIB Validation During Export

When you export an image profile to an ISO, vSphere ESXi Image Builder validates each VIB by performing the following actions.

- Checks that no conflicts exist by checking the Conflicts property of each SoftwarePackage object.
- Performs VIB signature validation. Signature validation prevents unauthorized modification of VIB packages. The signature is a cryptographic checksum that guarantees that a VIB was produced by its author. Signature validation also happens during installation of VIBs on an ESXi host and when the vSphere Auto Deploy server uses VIBs.

- Checks that VIBs follow file path usage rules. VMware tests `VMwareCertified` and `VMwareAccepted` VIBs to guarantee those VIBs always follow file path usage rules.

Acceptance Levels

Each VIB is released with an acceptance level that cannot be changed. The host acceptance level determines which VIBs can be installed to a host. You can change the host acceptance levels with `esxcli` commands.

VMware supports the following acceptance levels.

VMwareCertified	The <code>VMwareCertified</code> acceptance level has the most stringent requirements. VIBs with this level go through thorough testing fully equivalent to VMware in-house Quality Assurance testing for the same technology. Today, only I/O Vendor Program (IOVP) program drivers are published at this level. VMware takes support calls for VIBs with this acceptance level.
VMwareAccepted	VIBs with this acceptance level go through verification testing, but the tests do not fully test every function of the software. The partner runs the tests and VMware verifies the result. Today, CIM providers and PSA plug-ins are among the VIBs published at this level. VMware directs support calls for VIBs with this acceptance level to the partner's support organization.
PartnerSupported	VIBs with the <code>PartnerSupported</code> acceptance level are published by a partner that VMware trusts. The partner performs all testing. VMware does not verify the results. This level is used for a new or nonmainstream technology that partners want to enable for VMware systems. Today, driver VIB technologies such as Infiniband, ATAoE, and SSD are at this level with nonstandard hardware drivers. VMware directs support calls for VIBs with this acceptance level to the partner's support organization.
CommunitySupported	The <code>CommunitySupported</code> acceptance level is for VIBs created by individuals or companies outside of VMware partner programs. VIBs at this level have not gone through any VMware-approved testing program and are not supported by VMware Technical Support or by a VMware partner.

Structure of ImageProfile, SoftwarePackage, and ImageProfileDiff Objects

Knowing the structure of `ImageProfile`, `SoftwarePackage`, and `ImageProfileDiff` objects helps you manage deployment and upgrade processes.

ImageProfile Object Properties

The `ImageProfile` object, which is accessible with the `Get-EsxImageProfile` PowerCLI cmdlet, has the following properties.

Name	Type	Description
AcceptanceLevel	AcceptanceLevel	Determines which VIBs you can add to the profile. Levels are VMwareCertified, VMwareAccepted, PartnerSupported, and CommunitySupported. See Acceptance Levels .
Author	System.String	The person who created the profile. 60 characters or fewer.
CreationTime	System.DateTime	The timestamp of creation time.
Description	System.String	The full text description of profile. No length limit.
GUID	System.String	Globally unique ID of the image profile.
ModifiedTime	System.DateTime	The timestamp of last modification time.
Name	System.String	The name of the image profile. 80 characters or fewer.
ReadOnly	System.Boolean	When set to true, the profile cannot be edited. Use Set-EsxImageProfile -ReadOnly to make your custom image profiles read-only.
Rules	ImageProfileRule[]	Any OEM hardware requirements and restrictions that the image profile might have. vSphere Auto Deploy verifies the value of this property when deploying an image profile and deploys the profile if matching hardware is available.
Vendor	System.String	The organization that publishes the profile. 40 characters or fewer.
VibList	SoftwarePackage[]	The list of VIB IDs the image contains.

SoftwarePackage Object Properties

When preparing an image profile, you can examine software packages to decide which packages are suitable for inclusion. The SoftwarePackage object has the following properties.

Name	Type	Description
AcceptanceLevel	AcceptanceLevel	The acceptance level of this VIB.
Conflicts	SoftwareConstraint[]	A list of VIBs that cannot be installed at the same time as this VIB. Each constraint uses the following format: package-name[<< <= >= >> version]
Depends	SoftwareConstraint[]	A list of VIBs that must be installed at the same time as this VIB. Same constraint format as Conflicts property.
Description	System.String	The long description of the VIB.

Name	Type	Description
Guid	System.String	The unique ID for the VIB.
LiveInstallOk	System.Boolean	True if live installs of this VIB are supported.
LiveRemoveOk	System.Boolean	True if live removals of this VIB are supported.
MaintenanceMode	System.Boolean	True if hosts must be in maintenance mode for installation of this VIB.
Name	System.String	The name of the VIB. Usually uniquely describes the package on a running ESXi system.
Provides	SoftwareProvides	The list of virtual packages or interfaces this VIB provides. See SoftwareProvide Object Properties .
ReferenceURLs	SupportReference[]	The list of SupportReference objects with in-depth support information. The SupportReference object has two properties, Title and URL, both of type System.String.
Replaces	SoftwareConstraint[]	The list of SoftwareConstraint objects that identify VIBs that replace this VIB or make it obsolete. VIBs automatically replace VIBs with the same name but lower versions.
ReleaseDate	System.DateTime	Date and time of VIB publication or release.
SourceUrls	System.String[]	The list of source URLs from which this VIB can be downloaded.
StatelessReady	System.Boolean	True if the package supports host profiles or other technologies that make it suitable for use in conjunction with vSphere Auto Deploy.
Summary	System.String	A one-line summary of the VIB.
Tags	System.String[]	An array of string tags for this package defined by the vendor or publisher. Tags can be used to identify characteristics of a package.
Vendor	System.String	The VIB vendor or publisher.
Version	System.String	The VIB version.
VersionObject	Software.Version	The VersionObject property is of type SoftwareVersion. The SoftwareVersion class implements a static Compare method to compare two versions of strings. See SoftwareVersion Object Properties

ImageProfileDiff Object Properties

When you run the `Compare-EsxImageProfile` cmdlet, you pass in two parameters, first the reference profile, and then the comparison profile. The cmdlet returns an `ImageProfileDiff` object, which has the following properties.

Name	Type	Description
<code>CompAcceptanceLevel</code>	<code>System.String</code>	The acceptance level for the second profile that you passed to <code>Compare-EsxImageProfile</code> .
<code>DowngradeFromRef</code>	<code>System.String[]</code>	The list of VIBs in the second profile that are downgrades from VIBs in the first profile.
<code>Equal</code>	<code>System.Boolean</code>	True if the two image profiles have identical packages and acceptance levels.
<code>OnlyInComp</code>	<code>System.String</code>	The list of VIBs found only in the second profile that you passed to <code>Compare-EsxImageProfile</code> .
<code>OnlyInRef</code>	<code>System.String[]</code>	The list of VIBs found only in the first profile that you passed to <code>Compare-EsxImageProfile</code> .
<code>PackagesEqual</code>	<code>System.Boolean</code>	True if the image profiles have identical sets of VIB packages.
<code>RefAcceptanceLevel</code>	<code>System.String</code>	The acceptance level for the first profile that you passed to <code>Compare-EsxImageProfile</code> .
<code>UpgradeFromRef</code>	<code>System.String[]</code>	The list of VIBs in the second profile that are upgrades from VIBs in the first profile.

SoftwareVersion Object Properties

The `SoftwareVersion` object lets you compare two version strings. The object includes a `Compare` static method that accepts two strings as input and returns 1 if the first version string is a higher number than the second version string. `Compare` returns 0 if two versions strings are equal. `Compare` returns -1 if the second version string is a higher number than the first string. The object has the following properties.

Name	Type	Description
<code>Version</code>	<code>System.String</code>	The part of the version before the hyphen. This part indicates the primary version.
<code>Release</code>	<code>System.String</code>	The part of the version after the hyphen. This part indicates the release version.

SoftwareConstraint Object Properties

The `SoftwareConstraint` object implements a `MatchesProvide` method. The method accepts a `SoftwareProvides` or `SoftwarePackage` object as input and returns `True` if the constraint matches the `SoftwareProvide` or the `SoftwarePackage`, or returns `False` otherwise.

The `SoftwareConstraint` object includes the following properties.

Name	Type	Description
Name	System.String	The name of the constraint. This name should match a corresponding <code>SoftwareProvide Name</code> property.
Relation	System.String	An enum, or one of the following comparison indicators: <<, <=, = >=, >>. This property can be \$null if the constraint does not have a <code>Relation</code> and <code>Version</code> property.
Version	System.String	The version to match the constraint against. This property can be \$null if the constraint does not have a <code>Relation</code> and <code>Version</code> property.
VersionObject	SoftwareVersion	The version represented by a <code>SoftwareVersion</code> object.

SoftwareProvide Object Properties

The `SoftwareProvide` object includes the following properties.

Name	Type	Description
Name	System.String	The name of the provide.
Version	System.String	The version of the provide. Can be \$null if the provide does not specify a version.
Release	System.String	The version of the provide as represented by a <code>SoftwareVersion</code> object. See SoftwareVersion Object Properties .

vSphere ESXi Image Builder Installation and Usage

vSphere ESXi Image Builder consists of the vSphere ESXi Image Builder server and the vSphere ESXi Image Builder PowerShell cmdlets. The vSphere ESXi Image Builder server starts when you run the first vSphere ESXi Image Builder cmdlet.

Install vSphere ESXi Image Builder and Prerequisite Software

Before you can run vSphere ESXi Image Builder cmdlets, you must install PowerCLI and all prerequisite software. The vSphere ESXi Image Builder snap-in is included with the PowerCLI installation.

Prerequisites

If you want to manage vSphere ESXi Image Builder with PowerCLI cmdlets, verify that Microsoft .NET Framework 4.5 or 4.5.x and Windows PowerShell 3.0 or 4.0 are installed on a Microsoft Windows system. You can install PowerCLI on the Windows system on which vCenter Server is installed or on a different Windows system. See the *vSphere PowerCLI User's Guide*.

Procedure

- 1 Download the latest version of PowerCLI from the VMware Web site.
- 2 Navigate to the folder that contains the PowerCLI file you downloaded and double-click the executable file.

If the installation wizard detects an earlier version of PowerCLI on your system, it will attempt to upgrade your existing installation.

- 3 Follow the prompts in the wizard to complete the installation.

What to do next

Review [Using vSphere ESXi Image Builder Cmdlets](#). If you are new to PowerCLI, read the PowerCLI documentation.

Use vSphere ESXi Image Builder cmdlets and other PowerCLI cmdlets and PowerShell cmdlets to manage image profiles and VIBs. Use `Get-Help cmdlet_name` at any time for command-line help.

Configure the vSphere ESXi Image Builder Service Startup Type

Before you can use vSphere ESXi Image Builder with the vSphere Web Client, you must verify that the service is enabled and running.

Procedure

- 1 Log in to your vCenter Server system by using the vSphere Web Client.
- 2 On the vSphere Web Client Home page, click **Administration**.
- 3 Under **System Configuration** click **Services**.

4 Select **ImageBuilder Service**, click the **Actions** menu, and select **Edit Startup Type**.

- On Windows, the vSphere ESXi Image Builder service is disabled. In the **Edit Startup Type** window, select **Manual** or **Automatic** to enable Auto Deploy.
- On the vCenter Server Appliance, the vSphere ESXi Image Builder service by default is set to **Manual**. If you want the service to start automatically upon OS startup, select **Automatic**.

If you select the manual startup type, you must start the service manually upon OS startup every time you want to use the service.

5 (Optional) Click the **Start the service** icon.

6 (Optional) If you want to use vSphere ESXi Image Builder with the vSphere Web Client, log out of the vSphere Web Client and log in again.

The **Auto Deploy** icon is visible on the Home page of the vSphere Web Client.

What to do next

- [Add a Software Depot.](#)
- [Import a Software Depot.](#)
- [Clone an Image Profile.](#)
- [Create an Image Profile.](#)

Using vSphere ESXi Image Builder Cmdlets

vSphere ESXi Image Builder cmdlets are implemented as Microsoft PowerShell cmdlets and included in PowerCLI. Users of vSphere ESXi Image Builder cmdlets can take advantage of all PowerCLI features.

Experienced PowerShell users can use vSphere ESXi Image Builder cmdlets just like other PowerShell cmdlets. If you are new to PowerShell and PowerCLI, follow these tips.

You can type cmdlets, parameters, and parameter values in the PowerCLI shell.

- Get help for any cmdlet by running `Get-Help cmdlet_name`.
- Remember that PowerShell is not case sensitive.
- Use tab completion for cmdlet names and parameter names.
- Format any variable and cmdlet output by using `Format-List` or `Format-Table` or their short forms `fl` or `ft`. See `Get-Help Format-List`.
- Use wildcards for searching and filtering VIBs and image profiles. All wildcard expressions are supported.

Passing Parameters by Name

You can pass in parameters by name in most cases and surround parameter values that contain spaces or special characters with double quotes.

```
Add-EsxSoftwarePackage -ImageProfile profile42 -SoftwarePackage "partner package 35"
```

Passing Parameters as Objects

You can pass parameters as objects if you want to do scripting and automation. You can use the technique with cmdlets that return multiple objects or with cmdlets that return a single object.

- 1 Bind the output of a cmdlet that returns multiple objects to a variable.

```
$profs = Get-EsxImageProfile
```

- 2 When you run the cmdlet that needs the object as input, access the object by position, with the list starting with 0.

```
Add-EsxSoftwarePackage -ImageProfile $profs[4] -SoftwarePackage partner-pkg
```

The example adds the specified software package to the fifth image profile in the list returned by `Get-EsxImageProfile`.

Most of the examples in the *vSphere Installation and Setup* documentation pass in parameters by name. [vSphere ESXi Image Builder Workflows](#) includes examples that pass parameters as objects.

Using vSphere ESXi Image Builder with the vSphere Web Client

You can manage software packages (VIBs), image profiles, and software depots by using the vSphere ESXi Image Builder service in the vSphere Web Client.

- [Add a Software Depot](#)

Before you can work with software depots and customize image profiles, you must add one or more software depots to the vSphere ESXi Image Builder inventory. You can add a software depot by using the vSphere Web Client.

- [Import a Software Depot](#)

If an offline depot is located on your local file system, you can import the ZIP file to the vSphere ESXi Image Builder inventory by using the vSphere Web Client.

- [Clone an Image Profile](#)

You can use the vSphere Web Client to clone image profiles. You can clone an image profile when you want to make small changes to the VIB list in a profile, or if you want to use hosts from different vendors and want to use the same basic profile, but want to add vendor-specific VIBs.

- [Create an Image Profile](#)

You can create a new image profile by using the vSphere Web Client instead of cloning an existing one. You might consider creating a new image profile if it differs significantly from the image profiles in your inventory.

- [Edit an Image Profile](#)

You can edit image profiles by using the vSphere Web Client. You can change the name, details and VIB list of an image profile.

- [Compare Image Profiles](#)

You can compare two image profiles by using the vSphere Web Client, for example, to see if they have the same VIB list, version, or acceptance level.

- [Move an Image Profile to a Different Software Depot](#)

You can move image profiles between custom depots by using the vSphere Web Client. You can move an image profile to a custom depot to edit the image profile.

- [Export an Image Profile to ISO or Offline Bundle ZIP](#)

You can export an image profile to an ISO image or a ZIP file by using the vSphere Web Client. You can use the ISO image as an ESXi installer or to upgrade hosts with vSphere Upgrade Manager.

The ZIP file contains metadata and the VIBs of the image profile. You can use it for ESXi upgrades or as an offline depot.

Add a Software Depot

Before you can work with software depots and customize image profiles, you must add one or more software depots to the vSphere ESXi Image Builder inventory. You can add a software depot by using the vSphere Web Client.

Prerequisites

Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder Service Startup Type](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere ESXi Image Builder service.

- 2 On the **Software Depots** tab, click the **Add Software Depot** icon.
- 3 Select the type of depot that you want to create.

Option	Action
Online Depot	<ol style="list-style-type: none"> a Enter the name of the depot in the inventory. b Enter the URL of the online depot.
Custom Depot	Enter the name of the depot in the inventory.

- 4 Click **OK**.

What to do next

- You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
- You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
- [Edit the Image Profile Association of a Host](#).

Import a Software Depot

If an offline depot is located on your local file system, you can import the ZIP file to the vSphere ESXi Image Builder inventory by using the vSphere Web Client.

Prerequisites

Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder Service Startup Type](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.
By default, only the Administrator role has privileges to use the vSphere ESXi Image Builder service.
- 2 On the **Software Depots** tab, click the **Import Software Depot** icon.
- 3 Enter the name of the software depot in the inventory.
- 4 Click **Browse** and select a ZIP file from the local system, that contains the software depot you want to import.
- 5 Click **Upload**.

What to do next

- You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
- You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
- [Edit the Image Profile Association of a Host](#).

Clone an Image Profile

You can use the vSphere Web Client to clone image profiles. You can clone an image profile when you want to make small changes to the VIB list in a profile, or if you want to use hosts from different vendors and want to use the same basic profile, but want to add vendor-specific VIBs.

The acceptance level of the VIBs you add to the base image must be at least as high as the level of the base image. If you add a VIB with a lower acceptance level to the image profile, you must lower the image profile acceptance level. For more information, see [Working with Acceptance Levels](#).

Prerequisites

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder Service Startup Type](#).
- Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) and [Import a Software Depot](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere ESXi Image Builder service.

- 2 On the **Software Depots** tab, select the software depot that contains the image profile that you want to work with.
- 3 From the list of image profiles in the depot, select the image profile that you want to clone and click **Clone**.

- 4 Enter an image profile name, vendor, and description.

You must enter a unique image profile name.

- 5 From the **Software depot** drop-down list, select in which custom depot to add the new image profile and click **Next**.

- 6 (Optional) From the drop-down list, select an acceptance level for the image profile.

- 7 From the **Available** tab, select the VIBs that you want to add to the image profile and deselect the ones that you want to remove.

You can view the VIBs that will be added to the image profile from the **Selected** tab. You can filter the VIBs by software depot from the **Software depot** drop-down list on the **Available** tab.

Note The image profile must contain a bootable ESXi image to be valid.

- 8 Click **Next**.

vSphere ESXi Image Builder verifies that the change does not invalidate the profile. Some VIBs depend on other VIBs and become invalid if you include them in an image profile separately. When you add or remove a VIB, vSphere ESXi Image Builder checks whether the package dependencies are met.

- 9 On the Ready to complete page, review the summary information for the new image profile and click **Finish**.

What to do next

- You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
- You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
- [Edit the Image Profile Association of a Host](#).

Create an Image Profile

You can create a new image profile by using the vSphere Web Client instead of cloning an existing one. You might consider creating a new image profile if it differs significantly from the image profiles in your inventory.

The acceptance level of the VIBs you add to the base image must be at least as high as the level of the base image. If you add a VIB with a lower acceptance level to the image profile, you must lower the image profile acceptance level. For more information, see [Working with Acceptance Levels](#).

Prerequisites

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder Service Startup Type](#).
- Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) and [Import a Software Depot](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere ESXi Image Builder service.

- 2 On the **Software Depots** tab, select the custom depot in which you want to create a new image profile.
- 3 On the Image Profiles tab, click **New Image Profile**.
- 4 Enter an image profile name, vendor, and description.

You must enter a unique image profile name.

- 5 From the **Software depot** drop-down list, select in which custom depot to add the new image profile and click **Next**.
- 6 (Optional) From the drop-down list, select an acceptance level for the image profile.
- 7 From the **Available** tab, select the VIBs that you want to add to the image profile and deselect the ones that you want to remove.

You can view the VIBs that will be added to the image profile from the **Selected** tab. You can filter the VIBs by software depot from the **Software depot** drop-down list on the **Available** tab.

Note The image profile must contain a bootable ESXi image to be valid.

- 8 Click **Next**.

vSphere ESXi Image Builder verifies that the change does not invalidate the profile. Some VIBs are dependent on others and will not be valid if you include them in an image profile separately. When you add or remove a VIB, vSphere ESXi Image Builder checks that package dependencies are met.

- 9 On the Ready to complete page, review the summary information for the new image profile and click **Finish**.

What to do next

- You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
- You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
- [Edit the Image Profile Association of a Host](#).

Edit an Image Profile

You can edit image profiles by using the vSphere Web Client. You can change the name, details and VIB list of an image profile.

The acceptance level of the VIBs you add to the base image must be at least as high as the level of the base image. If you add a VIB with a lower acceptance level to the image profile, you must lower the image profile acceptance level. For more information, see [Working with Acceptance Levels](#).

Prerequisites

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder Service Startup Type](#).
- Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) and [Import a Software Depot](#).
- Verify that there is at least one custom depot in the vSphere ESXi Image Builder inventory.

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.
By default, only the Administrator role has privileges to use the vSphere ESXi Image Builder service.
- 2 On the **Software Depots** tab, select the software depot that contains the image profile that you want to work with.
- 3 On the **Image Profiles** tab, select the image profile that you want to edit and click **Edit**.
- 4 (Optional) Change the name, vendor and description information of the image profile.
- 5 Click **Next**.
- 6 From the **Available** tab, select the VIBs that you want to add to the image profile and deselect the ones that you want to remove.

You can view the VIBs that will be added to the image profile from the **Selected** tab. You can filter the VIBs by software depot from the **Software depot** drop-down list on the **Available** tab.

Note The image profile must contain a bootable ESXi image to be valid.

7 Click **Next**.

vSphere ESXi Image Builder verifies that the change does not invalidate the profile. Some VIBs depend on other VIBs and become invalid if you include them in an image profile separately. When you add or remove a VIB, vSphere ESXi Image Builder checks whether the package dependencies are met.

8 On the Ready to complete page, review the summary information for the edited image profile and click **Finish**.

What to do next

- You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
- You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
- [Edit the Image Profile Association of a Host](#).

Compare Image Profiles

You can compare two image profiles by using the vSphere Web Client, for example, to see if they have the same VIB list, version, or acceptance level.

Prerequisites

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder Service Startup Type](#).
- Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) and [Import a Software Depot](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.
By default, only the Administrator role has privileges to use the vSphere ESXi Image Builder service.
- 2 On the **Software Depots** tab, select the software depot that contains the image profile that you want to work with.
- 3 On the **Image Profiles** tab, select the image profile that you want to compare and click **Compare To**.
- 4 In the Compare Image Profile dialog box, from the **Software Depot** drop-down menu, select the software depot that contains the second image profile that you want to compare.
- 5 From the **Image Profile** drop-down menu, select the second image profile that you want to compare.
- 6 Under Software Packages, on the **All** tab, view the comparison of the two image profiles.

The left side of the list displays the names, versions, acceptance levels, and vendors of the VIBs that the first chosen image profile contains. The right part of the list provides information about the second image profile. The VIBs marked with no change are the same in both profiles. VIBs that are present in only one of the image profiles are marked **missing** in the image profile that they are not present in.

Move an Image Profile to a Different Software Depot

You can move image profiles between custom depots by using the vSphere Web Client. You can move an image profile to a custom depot to edit the image profile.

Prerequisites

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder Service Startup Type](#).
- Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) and [Import a Software Depot](#).
- Verify that there is at least one custom depot in the vSphere ESXi Image Builder inventory.

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.
By default, only the Administrator role has privileges to use the vSphere ESXi Image Builder service.
- 2 On the **Software Depots** tab, select the software depot that contains the image profile that you want to work with.
- 3 On the **Image Profiles** tab, select an image profile and click **Move to**.
- 4 From the drop-down list, select the custom depot in which you want to move the image profile.
- 5 Click **OK**.

Export an Image Profile to ISO or Offline Bundle ZIP

You can export an image profile to an ISO image or a ZIP file by using the vSphere Web Client. You can use the ISO image as an ESXi installer or to upgrade hosts with vSphere Upgrade Manager. The ZIP file contains metadata and the VIBs of the image profile. You can use it for ESXi upgrades or as an offline depot.

Prerequisites

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder Service Startup Type](#).
- Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) and [Import a Software Depot](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.
By default, only the Administrator role has privileges to use the vSphere ESXi Image Builder service.
- 2 On the **Software Depots** tab, select the software depot that contains the image profile that you want to work with.
- 3 On the **Image Profiles** tab, select the image profile that you want to export and click **Export Image Profile**.

- 4 Select the type of the exported file.

Option	Description
ISO	Exports the image profile to a bootable ISO image. If you want to create an ISO image that you can burn to a CD or DVD and use to boot up a stateless ESXi instance, select the Do not include an installer on the ISO check box.
ZIP	Exports the image profile to a ZIP file.

- 5 (Optional) If you want to bypass the acceptance level verification of the image profile, select **Skip acceptance level checking**.
- 6 Click the **Generate image** button.
- 7 When the image generates successfully, click **Download** to download the exported file.
- 8 Click **Close**.

Using vSphere ESXi Image Builder with PowerCLI Cmdlets

The vSphere ESXi Image Builder cmdlets allow you to manipulate software depots, image profiles, and VIBs.

Clone an Image Profile

Cloning a published profile is the easiest way to create a custom image profile. Cloning a profile is especially useful if you want to remove a few VIBs from a profile, or if you want to use hosts from different vendors and want to use the same basic profile, but want to add vendor-specific VIBs. VMware partners or large installations might consider creating a new profile.

Prerequisites

- Install the PowerCLI and all prerequisite software. See [vSphere ESXi Image Builder Installation and Usage](#).
- Verify that you have access to the software depot that contains the image profile you want to clone.

Procedure

- 1 In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file system. b Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- 2 (Optional) Run the `Get-EsxImageProfile` cmdlet to find the name of the profile that you want to clone.

You can use filtering options with `Get-EsxImageProfile`.

- 3 Run the `New-ExsImageProfile` cmdlet to create the new profile and use the `-CloneProfile` parameter to specify the profile you want to clone.

```
New-ExsImageProfile -CloneProfile My_Profile -Name "Test Profile 42"
```

This example clones the profile named *My_Profile* and assigns it the name Test Profile 42. You must specify a unique combination of name and vendor for the cloned profile.

What to do next

See [Examine Depot Contents](#) for some examples of filtering.

Customize the image profile by adding or removing VIBs. See [Add VIBs to an Image Profile](#).

Add VIBs to an Image Profile

You can add one or more VIBs to an image profile if that image profile is not set to read only. If the new VIB depends on other VIBs or conflicts with other VIBs in the profile, a message is displayed at the PowerShell prompt and the VIB is not added.

You can add VIBs from VMware or from VMware partners to an image profile. If you add VMware VIBs, vSphere ESXi Image Builder performs validation. If you add VIBs from two or more OEM partners simultaneously, no errors are reported but the resulting image profile might not work. Install VIBs from only one OEM vendor at a time.

If an error about acceptance level problems appears, change the acceptance level of the image profile and the acceptance level of the host. Consider carefully whether changing the host acceptance level is appropriate. VIB acceptance levels are set during VIB creation and cannot be changed.

You can add VIBs even if the resulting image profile is invalid.

Note VMware can support only environments and configurations that are proven to be stable and fully functional through rigorous and extensive testing. Use only those supported configurations. You can use custom VIBs if you lower your host acceptance level, and as a result, supportability. In that case, track the changes you made, so you can revert them if you want to remove custom VIBs and restore the host acceptance level to the default (Partner Supporter) later. See [Working with Acceptance Levels](#).

Prerequisites

Install the PowerCLI and all prerequisite software. See [Install vSphere ESXi Image Builder and Prerequisite Software](#)

Procedure

- 1 In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file system. b Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- 2 Run the `Get-EsxImageProfile` cmdlet to list all image profiles in all currently visible depots.

The cmdlet returns all available profiles. You can narrow your search by using the optional arguments to filter the output.

- 3 Clone the profile.

```
New-EsxImageProfile -CloneProfile My_Profile -Name "Test Profile 42" -Vendor "My Vendor"
```

Image profiles published by VMware and its partners are read only. To make changes, you must clone the image profile. The vendor parameter is required.

- 4 Run the `Add-EsxSoftwarePackage` cmdlet to add a new package to one of the image profiles.

```
Add-EsxSoftwarePackage -ImageProfile My_Profile -SoftwarePackage partner-package
```

The cmdlet runs the standard validation tests on the image profile. If validation succeeds, the cmdlet returns a modified, validated image profile. If the VIB that you want to add depends on a different VIB, the cmdlet displays that information and includes the VIB that would resolve the dependency. If the acceptance level of the VIB that you want to add is lower than the image profile acceptance level, an error occurs.

Export an Image Profile to ISO or Offline Bundle ZIP

You can export an image profile to an ISO image or a ZIP file of component files and folders. You cannot create both by running the cmdlet once. You can use the ISO image as an ESXi installer or upload the ISO into vSphere Update Manager for upgrades. You can use the ZIP file, which contains metadata and the VIBs specified in the image profile, for upgrades to ESXi 5.0 and later.

Prerequisites

Install the PowerCLI and all prerequisite software. See [Install vSphere ESXi Image Builder and Prerequisite Software](#).

Procedure

- 1 In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file system. b Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- 2 Run `Export-EsxImageProfile` to export the image profile.

Export Format	Cmdlet
ISO images	<code>Export-EsxImageProfile</code> with the <code>-ExportToIso</code> parameter
Offline depot ZIP files	<code>Export-EsxImageProfile</code> with the <code>-ExportToBundle</code> parameter

For the ISO image, vSphere ESXi Image Builder validates VIB signatures, adds VIB binaries to the image, and downloads the image to the specified location. For the ZIP file, vSphere ESXi Image Builder validates VIB signatures and downloads the VIB binaries to the specified location.

Example: Exporting an Image Profile

Follow these steps to export an image profile to an ISO image.

- 1 Add the software depot.

```
Add-EsxSoftwareDepot -DepotUrl url_or_file
```

- 2 View all available image profiles to find the name of the image profile to export.

```
Get-EsxImageProfile
```

- 3 Export the image profile.

```
Export-EsxImageProfile -ImageProfile "myprofile" -ExportToIso -FilePath iso_name
```

Follow these steps to export an image profile to a ZIP file of component files and folders.

- 1 Add the software depot.

```
Add-EsxSoftwareDepot -DepotUrl url_or_file
```

- 2 View all available image profiles to find the name of the image profile to export.

```
Get-EsxImageProfile
```

3 Export the image profile.

```
Export-ESXImageProfile -ImageProfile "myprofile" -ExportToBundle -FilePath C:\my_bundle.zip
```

What to do next

Use the ISO image in an ESXi installation or upload the ISO image into vSphereUpdate Manager to perform upgrades.

Use the ZIP file to upgrade an ESXi installation.

- Import the ZIP file into vSphere Update Manager for use with patch baselines.
- Download the ZIP file to an ESXi host or a datastore and run `esxcli software vib` commands to import the VIBs in the ZIP file.

See the *vSphere Upgrade* documentation.

Preserve Image Profiles Across Sessions

When you create an image profile and exit the PowerCLI session, the image profile is no longer available when you start a new session. You can export the image profile to a ZIP file software depot, and add that depot in the next session.

Prerequisites

Install the PowerCLI and all prerequisite software. See [Install vSphere ESXi Image Builder and Prerequisite Software](#).

Procedure

- 1 In a PowerCLI session, create an image profile, for example by cloning an existing image profile and adding a VIB.
- 2 Export the image profile to a ZIP file by calling `Export-ESXImageProfile` with the `ExportToBundle` parameter.

```
Export-ESXImageProfile -ImageProfile "my_profile" -ExportToBundle -FilePath
"C:\isos\temp-base-plus-vib25.zip"
```

- 3 Exit the PowerCLI session.
- 4 When you start a new PowerCLI session, add the depot that contains your image profile to access it.

```
Add-ESXSoftwareDepot "C:\isos\temp-base-plus-vib25.zip"
```

Compare Image Profiles

You can compare two image profiles by using the `Compare-EsxImageProfile` cmdlet, for example, to see if they have the same VIB list or acceptance level. Comparing image profiles or their properties is also possible by using the PowerShell comparison operators.

Prerequisites

Install the PowerCLI and all prerequisite software. See [Install vSphere ESXi Image Builder and Prerequisite Software](#).

Procedure

- 1 In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file system. b Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- 2 (Optional) Run the `Get-EsxImageProfile` cmdlet to view a list of all image profiles in all available depots.

In the list, you can locate the names of the image profiles you want to compare.

- 3 Before comparing the image profiles, assign them to variables.

For example, you can create variables `$imageProfile1` and `$imageProfile2` to hold the names of the compared images profiles.

```
$imageProfile1
= Get-EsxImageProfile -Name "ImageProfile1"
$imageProfile2
= Get-EsxImageProfile -Name "ImageProfile2"
```

- 4 Compare the two image profiles by using the `Compare-EsxImageProfile` cmdlet or the `-eq` comparison operator, which returns a Boolean value.
 - Compare the two image profiles to get a full description of the differences by using the `Compare-EsxImageProfile` cmdlet.

```
Compare-EsxImageProfile -ReferenceProfile
                        $imageProfile1 -ComparisonProfile $imageProfile2
```

- Compare the two image profiles by VIB list and acceptance level using the `-eq` comparison operator.

```
if ($imageProfile1 -eq $imageProfile2) {
    Write-host "Successfully verified that both image profiles are equal."
} else {
    Write-host "Failed to verify that the image profiles are equal."
}
```

- Compare the two image profiles by a specific property using the `-eq` comparison operator.

```
if ($imageProfile1.vendor -eq $imageProfile2.vendor) {
    Write-host "Successfully verified that both image profiles are equal."
} else {
    Write-host "Failed to verify that the image profiles are equal."
}
```

Compare VIBs

You can compare two VIBs or their properties by using the PowerShell comparison operators.

Prerequisites

Install the PowerCLI and all prerequisite software. See [Install vSphere ESXi Image Builder and Prerequisite Software](#).

Procedure

- 1 In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file system. b Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- 2 (Optional) Run the `Get-EsxSoftwarePackage` cmdlet to view all available VIBs.

In the list, you can locate the names of the VIBs you want to compare.

3 Before comparing the VIBs, assign them to variables.

For example, you can create variables `$vib1` and `$vib2` to hold the names of the compared VIBs.

```
$vib1 = Get-ExxSoftwarePackage -Name "ReferenceVIB"  
$vib2 = Get-ExxSoftwarePackage -Name "ComparisonVIB"
```

4 Use a comparison operator to compare the VIBs by contents and acceptance level or by a specific property.

- Compare the two VIBs by their contents and acceptance level.

```
if ($vib1 -eq $vib2) {  
    Write-host "Successfully verified that both VIBs are equal."  
} else {  
    Write-host "Failed to verify that the VIBs are equal."  
}
```

- Compare a specific property of the VIBs by using a comparison operator such as `-eq`, `-lt`, `-le`, `-gt` or `-ge`.

```
if ($vib1.VersionObject -lt $vib2.VersionObject) {  
    Write-host "Successfully verified that both the VIBs are equal."  
} else {  
    Write-host "Failed to verify that the VIBs are equal."  
}
```

Working with Acceptance Levels

Hosts, image profiles, and individual VIBs have acceptance levels. VIB acceptance levels show how the VIB was tested. Understanding what each acceptance level implies, how to change levels, and what a change implies is an important part of installation and update procedures.

Acceptance levels are set for hosts, image profiles, and individual VIBs. The default acceptance level for an ESXi image or image profile is PartnerSupported.

Host acceptance levels The host acceptance level determines which VIBs you can install on a host. You can change a host's acceptance level with ESXCLI commands. By default, ESXi hosts have an acceptance level of PartnerSupported to allow for easy updates with PartnerSupported VIBs.

Note VMware supports hosts at the PartnerSupported acceptance level. For problems with individual VIBs with PartnerSupported acceptance level, contact your partner's support organization.

Image profile acceptance levels

The image profile acceptance level is set to the lowest VIB acceptance level in the image profile. If you want to add a VIB with a low acceptance level to an image profile, you can change the image profile acceptance level with the `Set-EsxImageProfile` cmdlet. See [Set the Image Profile Acceptance Level](#).

The vSphere Update Manager does not display the actual acceptance level. Use vSphere ESXi Image Builder cmdlets to retrieve the acceptance level information for VIBs and image profiles.

VIB acceptance levels A VIB's acceptance level is set when the VIB is created. Only the VIB creator can set the acceptance level.

If you attempt to provision a host with an image profile or VIB that has a lower acceptance level than the host, an error occurs. Change the acceptance level of the host to install the image profile or VIB. See [Change the Host Acceptance Level](#). Changing the acceptance level of the host changes the support level for that host.

The acceptance level of a host, image profile, or VIB lets you determine who tested the VIB and who supports the VIB. VMware supports the following acceptance levels .

VMwareCertified	The VMwareCertified acceptance level has the most stringent requirements. VIBs with this level go through thorough testing fully equivalent to VMware in-house Quality Assurance testing for the same technology. Today, only I/O Vendor Program (IOVP) program drivers are published at this level. VMware takes support calls for VIBs with this acceptance level.
VMwareAccepted	VIBs with this acceptance level go through verification testing, but the tests do not fully test every function of the software. The partner runs the tests and VMware verifies the result. Today, CIM providers and PSA plug-ins are among the VIBs published at this level. VMware directs support calls for VIBs with this acceptance level to the partner's support organization.
PartnerSupported	VIBs with the PartnerSupported acceptance level are published by a partner that VMware trusts. The partner performs all testing. VMware does not verify the results. This level is used for a new or nonmainstream technology that partners want to enable for VMware systems. Today, driver VIB technologies such as Infiniband, ATAOE, and SSD are at this level with nonstandard hardware drivers. VMware directs support calls for VIBs with this acceptance level to the partner's support organization.
CommunitySupported	The CommunitySupported acceptance level is for VIBs created by individuals or companies outside of VMware partner programs. VIBs at this level have not gone through any VMware-approved testing program and are not supported by VMware Technical Support or by a VMware partner.

Change the Host Acceptance Level

You can lower the host acceptance level to match the acceptance level for a VIB or image profile you want to install.

The acceptance level of each VIB on a host must be at least as high as the acceptance level of the host. For example, you cannot install a VIB with PartnerSupported acceptance level on a host with VMwareAccepted acceptance level. You must first lower the acceptance level of the host. For more information on acceptance levels, see [Acceptance Levels](#).

Changing the host acceptance level to CommunitySupported affects the supportability of your host and might affect the security of your host.

Prerequisites

Install vCLI or deploy the vSphere Management Assistant (vMA) virtual machine. See *Getting Started with vSphere Command-Line Interfaces*. For troubleshooting, run `esxcli` commands in the ESXi Shell.

Procedure

- 1 Retrieve the acceptance level for the VIB or image profile.

Option	Description
View information for all VIBs	<code>esxcli --server=<i>server_name</i> software sources vib list --depot=<i>depot_URL</i></code>
View information for a specified VIB	<code>esxcli --server=<i>server_name</i> software sources vib list --viburl=<i>vib_URL</i></code>
View information for all image profiles	<code>esxcli --server=<i>server_name</i> software sources profile list --depot=<i>depot_URL</i></code>
View information for a specified image profile	<code>esxcli --server=<i>server_name</i> software sources profile get --depot=<i>depot_URL</i> --profile=<i>profile_name</i></code>

- 2 View the host acceptance level.

```
esxcli --server=server_name software acceptance get
```

- 3 Change the acceptance level of the host.

```
esxcli
    --server=server_name software acceptance set --level=acceptance_level
```

The value for *acceptance_level* can be `VMwareCertified`, `VMwareAccepted`, `PartnerSupported`, or `CommunitySupported`. The values for *acceptance_level* are case-sensitive.

Note If the host has a higher acceptance level than the VIB or image profile you want to add, you can run commands in the `esxcli software vib` or `esxcli software profile` namespace with the `--force` option. When you use the `--force` option, a warning appears because you enforce a VIB or image profile with lower acceptance level than the acceptance level of the host and your setup is no longer consistent. The warning is repeated when you install VIBs, remove VIBs, or perform certain other operations on the host that has inconsistent acceptance levels.

Set the Image Profile Acceptance Level

If you want to add a VIB to an image profile, and the acceptance level of the VIB is lower than that of the image profile, you can clone the image profile with a lower acceptance level or change the image profile's acceptance level.

You can specify `VMwareCertified`, `VMwareAccepted`, `PartnerSupported`, or `CommunitySupported` as an acceptance level of an image profile. If you lower the acceptance level, the level of support for the image profile and hosts that you provision with it changes. For more information, see [Acceptance Levels](#).

Prerequisites

Install PowerCLI and all prerequisite software. See [Install vSphere ESXi Image Builder and Prerequisite Software](#).

Procedure

- 1 In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file system. b Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- 2 Get the acceptance level for the image profile.

```
Get-EsxImageProfile -Name string
```

- 3 Set the acceptance level of the image profile.

```
Set-EsxImageProfile -Name string -AcceptanceLevel level
```

vSphere ESXi Image Builder Workflows

vSphere ESXi Image Builder workflows are examples for cmdlet usage. Workflows do not represent actual tasks, but illustrate how you might explore different ways of using a cmdlet. Administrators trying out the workflows benefit from some experience with PowerCLI, Microsoft PowerShell, or both.

Examine Depot Contents

You can examine software depots and VIBs with vSphere ESXi Image Builder cmdlets. You can use wildcards to examine depot contents. All wildcard expressions are supported.

The workflow itself passes parameters by name. However, you can pass parameters as objects by accessing variables.

You can use filtering options and wildcard expressions to examine depot contents.

Prerequisites

Verify that PowerCLI and prerequisite software is installed. See [Install vSphere ESXi Image Builder and Prerequisite Software](#).

Procedure

- 1 In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file system. b Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- 2 Retrieve image profiles.

You can filter by vendor, name, and acceptance level.

- `Get-EsxImageProfiles`

Returns an array of `ImageProfile` objects from all depots you added to the session.

- `Get-EsxImageProfile -Vendor "C*"`

Returns all image profiles created by a vendor with a name that starts with the letter C.

- 3 Retrieve software packages by using the `Get-EsxSoftwarePackage` cmdlet.

You can filter, for example by vendor or version, and you can use the standard PowerShell wildcard characters.

- `Get-EsxSoftwarePackage -Vendor "V*"`

Returns all software packages from a vendor with a name that starts with the letter V.

- `Get-EsxSoftwarePackage -Vendor "V*" -Name "*scsi*"`

Returns all software packages with a name that contains the string `scsi` in it from a vendor with a name that starts with the letter V.

- `Get-EsxSoftwarePackage -Version "2.0*"`

Returns all software packages with a version string that starts with 2.0.

4 Use `-Newest` to find the latest package.

```
Get-ExSoftwarePackage -Vendor "V*" -Newest
```

Returns the newest package for the vendors with a name that starts with the letter V, and displays the information as a table.

```
Get-ExSoftwarePackage -Vendor "V*" -Newest | format-list
```

Returns detailed information about each software package by using a pipeline to link the output of the request for software packages to the PowerShell `format-list` cmdlet.

5 View the list of VIBs in the image profile.

```
(Get-ExImageProfile -Name "Robin's Profile").VibList
```

`VibList` is a property of the `ImageProfile` object.

6 Retrieve software packages released before or after a certain date by using the `CreatedBefore` or `CreatedAfter` parameter.

```
Get-ExSoftwarePackage -CreatedAfter 7/1/2010
```

Example: Depot Content Examination Using Variables

This workflow example examines depot contents by passing in parameters as objects accessed by position in a variable, instead of passing in parameters by name. You can run the following commands in sequence from the PowerCLI prompt. Replace names with names that are appropriate in your installation.

```
Get-ExSoftwarePackage -Vendor "V*"
Get-ExSoftwarePackage -Vendor "V*" -Name "r*"
Get-ExSoftwarePackage -Version "2.0*"
$ip1 = Get-ExImageProfile -name ESX-5.0.0-123456-full
$ip1.VibList
Get-ExSoftwarePackage -CreatedAfter 7/1/2010
```

Create Image Profiles by Cloning Workflow

You can use vSphere ESXi Image Builder cmdlets to check which depots are available, to add a depot, to view image profile information, and to create a new image profile by cloning one of the available image profiles.

Published profiles are usually read-only and cannot be modified. Even if a published profile is not read-only, cloning instead of modifying the profile is a best practice, because modifying the original profile erases the original. You cannot revert to the original, unmodified profile except by reconnecting to a depot.

A profile cloning workflow might include checking the current state of the system, adding a software depot, and cloning the profile.

Prerequisites

Verify that PowerCLI and prerequisite software is installed. See [Install vSphere ESXi Image Builder and Prerequisite Software](#).

Procedure

- 1 In a PowerShell window, check whether any software depots are defined for the current session.

```
$DefaultSoftwareDepots
```

PowerShell returns the currently defined depots, or nothing if you just started PowerShell.

- 2 If the depot containing the profile that you want to clone does not appear in the results, add it to the current session.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file path. b Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

PowerShell adds the specified depot to your current session and lists all current depots.

- 3 (Optional) Check the `$DefaultSoftwareDepots` variable, which now returns the newly added depot.
- 4 View all available image profiles.

```
Get-EsxImageProfile
```

- 5 To clone an image profile, enter its name, a new name for the new profile, and a name of the vendor.

```
$ip = New-EsxImageProfile -CloneProfile base-tbd-v1 -Name "Test Profile 42" -Vendor "Vendor20"
```

- 6 (Optional) View the newly created image profile, `$ip`.

PowerShell returns the information about the image profile in tabular format.

Name	Vendor	Last Modified	Acceptance Level
----	-----	-----	-----
Test Profile 42	Vendor20	9/15/2010 5:45:43...	PartnerSupported

Example: Creating Image Profile by Cloning Using Variables

This workflow example repeats the steps of this workflow by passing in parameters as objects accessed by position in a variable, instead of passing in parameters by name. You can run the following cmdlets in sequence from the PowerCLI prompt.

```
$DefaultSoftwareDepots
Add-EsxSoftwareDepot -DepotUrl depot_url
$DefaultSoftwareDepots
$profs = Get-EsxImageProfile
$profs
$ip = New-EsxImageProfile -CloneProfile $profs[2] -Name "new_profile_name" -Vendor "my_vendor"
$ip
```

Create New Image Profiles Workflow

In most situations, you create an image profile by cloning an existing profile. Some VMware customers or partners might need to create a new image profile. Pay careful attention to dependencies and acceptance levels if you create an image profile from scratch.

The system expects that the acceptance level of the VIBs you add to the base image is at least as high as the level of the base image. If you have to add a VIB with a lower acceptance level to the image profile, you must lower the image profile acceptance level. For more information, see [Set the Image Profile Acceptance Level](#).

As an alternative to specifying the parameters on the command line, you can use the PowerShell prompting mechanism to specify string parameters. Prompting does not work for other parameters such as objects.

Prerequisites

- PowerCLI and prerequisite software is installed. See [Install vSphere ESXi Image Builder and Prerequisite Software](#).
- You have access to a depot that includes a base image and one or more VIBs. VMware and VMware partners have public depots, accessible by a URL. VMware or VMware partners can create a ZIP file that you can unzip to your local environment and access by using a file path.

Procedure

- 1 In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file system. b Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- 2 Run the `Get-ExsImageProfile` cmdlet to list all image profiles in all currently visible depots. You can narrow your search by using the optional arguments to filter the output.

```
Get-ExsSoftwarePackage -CreatedAfter 7/1/2010
```

- 3 Create a new profile, assign it a name and vendor, and add a base package.

```
New-ExsImageProfile -NewProfile -Name "Test #2" -vendor "Vendor42" -SoftwarePackage esx-base[0],esx-xlibs[0]
```

The example uses the `esx-base` package. In most cases, you include the `esx-base` package when you create a new image profile. Names that contain spaces are surrounded by quotes.

- 4 Use a pipeline to pass the new image profile to `format-list` for detailed information about the new package.

```
(Get-ExsImageProfile -Name "Test #2").VibList | format-list
```

Example: Creating Image Profiles from Scratch Using Variables

This command sequence repeats the steps of the workflow, but passes parameters as objects, accessed by position in a variable, instead of passing parameters by name. You can run the following commands in sequence at the `PowerCLI` prompt.

```
Add-ExsSoftwareDepot depoturl
$pkgs = Get-ExsSoftwarePackage -CreatedAfter 7/1/2010
$ip2 = New-ExsImageProfile -NewProfile -Name "Test #2" -vendor "Vendor42" -SoftwarePackage $pkgs[0]
$ip2.VibList | format-list
```

Edit Image Profiles Workflow

You can create a custom image by cloning and editing an image profile by using `PowerCLI`. You can add or remove one or more VIBs in the existing profile. If adding or removing VIBs prevents the image profile from working correctly, an error occurs.

Prerequisites

- `PowerCLI` and prerequisite software is installed. See [Install vSphere ESXi Image Builder and Prerequisite Software](#).
- You have access to a depot that includes a base image and one or more VIBs. VMware and VMware partners make public depots, accessible by a URL, available. VMware or VMware partners can create a ZIP file that you can download to your local environment and access by using a file path.

Procedure

- 1 In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file system. b Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- 2 Use a pipeline to pass the image profile you intend to edit to `format-list` to see detailed information.

In this example, the image profile created in [Create New Image Profiles Workflow](#) contains only the base image. A newly created image profile is not included in the depot. Instead, you access the image profile by name or by binding it to a variable.

```
Get-EsxImageProfile "Test #2" | format-list
```

PowerShell returns the information.

```
Name           : Test #2
Vendor         : Vendor42
...
VibList        : {esx-base 5.0.0.-...,}
```

- 3 (Optional) If you are adding a VIB with a lower acceptance level than that of the image profile, change the acceptance level of the image profile.

```
Set-EsxImageProfile -ImageProfile "Test #2" -AcceptanceLevel VMwareAccepted
```

PowerShell returns the information about the changed profile in tabular format.

```
Name           Vendor           Last Modified           Acceptance Level
----           -
Test #2        Vendor42          9/22/2010 12:05:...    VMwareAccepted
```

- 4 Add a software package (VIB) to the image profile. You can add the package by name.

```
Add-EsxSoftwarePackage -ImageProfile "Test #2"
                        -SoftwarePackage NewPack3
```

PowerShell returns the information about the image profile in tabular format.

Name	Vendor	Last Modified	Acceptance Level
Test #2	Vendor42	9/22/2010 12:05:...	VMwareAccepted

Note If an error occurs when you add the software package, you might have a problem with acceptance levels, see [Working with Acceptance Levels](#)

- 5 View the image profile again.

```
Get-EsxImageProfile "Test #2" | format-list
```

The VIB list is updated to include the new software package and the information is displayed.

```
Name           : Test #2
Vendor         : Vendor42
...
VibList        : {esx-base 5.0.0.-..., NewPack3}
```

Example: Editing Image Profiles by Using Variables

This cmdlet sequence repeats the steps of the workflow but passes parameters as objects, accessed by position in a variable, instead of passing parameters by name. You can run the following cmdlets in sequence from the PowerCLI prompt.

```
Add-EsxSoftwareDepot -DepotUrl depot_url
$ip2 = Get-EsxImageProfile -name "Test #2"
$ip2 | format-list
Set-EsxImageProfile -ImageProfile $ip2 -AcceptanceLevel VMwareAccepted
Add-EsxImageSoftwarePackage -ImageProfile $ip2 -SoftwarePackage NewPack3
$ip2 | format-list
```

Required Information for ESXi Installation

In an interactive installation, the system prompts you for the required system information. In a scripted installation, you must supply this information in the installation script.

For future use, note the values you use during the installation. These notes are useful if you must reinstall ESXi and reenter the values that you originally chose.

Table 2-10. Required Information for ESXi Installation

Information	Required or Optional	Default	Comments
Keyboard layout	Required	U.S. English	
VLAN ID	Optional	None	Range: 0 through 4094
IP address	Optional	DHCP	You can allow DHCP to configure the network during installation. After installation, you can change the network settings.
Subnet mask	Optional	Calculated based on the IP address	
Gateway	Optional	Based on the configured IP address and subnet mask	
Primary DNS	Optional	Based on the configured IP address and subnet mask	
Secondary DNS	Optional	None	
Host name	Required for static IP settings	None	The vSphere Web Client can use either the host name or the IP address to access the ESXi host.
Install location	Required	None	Must be at least 5 GB if you install the components on a single disk.
Migrate existing ESXi settings. Preserve existing VMFS datastore.	Required if you are installing ESXi on a drive with an existing ESXi installation.	None	If you have an existing ESXi 5.x installation, the ESXi installer offers a choice between preserving or overwriting the VMFS datastore during installation
Root password	Required	None	The root password must contain between 8 and 40 characters. For information about passwords see the <i>vSphere Security</i> documentation.

Installing ESXi

You can install ESXi interactively, with a scripted installation, or with vSphere Auto Deploy.

Installing ESXi Interactively

Use the interactive installation option for small deployments of fewer than five hosts.

In a typical interactive installation, you boot the ESXi installer and respond to the installer prompts to install ESXi to the local host disk. The installer reformats and partitions the target disk and installs the ESXi boot image. If you have not installed ESXi on the target disk before, all data on the drive is overwritten, including hardware vendor partitions, operating system partitions, and associated data.

Note To ensure that you do not lose any data, migrate the data to another machine before you install ESXi.

If you are installing ESXi on a disk that contains a previous installation of ESXi or ESX, or a VMFS datastore, the installer provides you with options for upgrading. See the *vSphere Upgrade* documentation.

Install ESXi Interactively

You use the ESXi CD/DVD or a USB flash drive to install the ESXi software onto a SAS, SATA, SCSI hard drive, or USB drive.

Prerequisites

- You must have the ESXi installer ISO in one of the following locations:
 - On CD or DVD. If you do not have the installation CD/DVD, you can create one. See [Download and Burn the ESXi Installer ISO Image to a CD or DVD](#)
 - On a USB flash drive. See [Format a USB Flash Drive to Boot the ESXi Installation or Upgrade](#).

Note You can also PXE boot the ESXi installer to run an interactive installation or a scripted installation. See [PXE Booting the ESXi Installer](#).

- Verify that the server hardware clock is set to UTC. This setting is in the system BIOS.
- Verify that a keyboard and monitor are attached to the machine on which the ESXi software is installed. Alternatively, use a remote management application. See [Using Remote Management Applications](#).
- Consider disconnecting your network storage. This action decreases the time it takes the installer to search for available disk drives. When you disconnect network storage, any files on the disconnected disks are unavailable at installation.

Do not disconnect a LUN that contains an existing ESX or ESXi installation. Do not disconnect a VMFS datastore that contains the Service Console of an existing ESX installation. These actions can affect the outcome of the installation.

- Gather the information required by the ESXi installation wizard. See [Required Information for ESXi Installation](#).
- Verify that ESXi Embedded is not present on the host machine. ESXi Installable and ESXi Embedded cannot exist on the same host.

Procedure

- 1 Insert the ESXi installer CD/DVD into the CD/DVD-ROM drive, or attach the Installer USB flash drive and restart the machine.
- 2 Set the BIOS to boot from the CD-ROM device or the USB flash drive.

See your hardware vendor documentation for information on changing boot order.

- 3 On the Select a Disk page, select the drive on which to install ESXi, and press Enter.

Press F1 for information about the selected disk.

Note Do not rely on the disk order in the list to select a disk. The disk order is determined by the BIOS and might be out of order. This might occur on systems where drives are continuously being added and removed.

If you select a disk that contains data, the Confirm Disk Selection page appears.

If you are installing on a disc with a previous ESXi or ESX installation or VMFS datastore, the installer provides several choices.

Important If you are upgrading or migrating an existing ESX/ESXi installation, see the *vSphere Upgrade* documentation. The instructions in this *vSphere Installation and Setup* documentation are for a fresh installation of ESXi.

If you select a disk that is in vSAN disk group, the resulting installation depends on the type of disk and the group size:

- If you select an SSD, the SSD and all underlying HDDs in the same disk group are wiped.
- If you select an HDD, and the disk group size is greater than two, only the selected HDD is wiped.
- If you select an HDD disk, and the disk group size is two or less, the SSD and the selected HDD is wiped.

For more information about managing vSAN disk groups, see the *vSphere Storage* documentation.

- 4 Select the keyboard type for the host.

You can change the keyboard type after installation in the direct console.

- 5 Enter the root password for the host.

You can change the password after installation in the direct console.

- 6 Press Enter to start the installation.

- 7 When the installation is complete, remove the installation CD, DVD, or USB flash drive.

- 8 Press Enter to reboot the host.

If you are performing a new installation, or you chose to overwrite an existing VMFS datastore, during the reboot operation, VFAT scratch, and VMFS partitions are created on the host disk.

- 9 Set the first boot device to be the drive on which you installed ESXi in [Step 3](#).

For information about changing boot order, see your hardware vendor documentation.

Note UEFI systems might require additional steps to set the boot device. See [Host Fails to Boot After You Install ESXi in UEFI Mode](#)

After the installation is complete, you can migrate existing VMFS data to the ESXi host.

You can boot a single machine from each ESXi image. Booting multiple devices from a single shared ESXi image is not supported.

What to do next

Set up basic administration and network configuration for ESXi. See [After You Install and Set Up ESXi](#).

Install ESXi on a Software iSCSI Disk

When you install ESXi to a software iSCSI disk, you must configure the target iSCSI qualified name (IQN).

During system boot, the system performs a Power-On Self Test (POST), and begins booting the adapters in the order specified in the system BIOS. When the boot order comes to the iSCSI Boot Firmware Table (iBFT) adapter, the adapter attempts to connect to the target, but does not boot from it. See Prerequisites.

If the connection to the iSCSI target is successful, the iSCSI boot firmware saves the iSCSI boot configuration in the iBFT. The next adapter to boot must be the ESXi installation media, either a mounted ISO image or a physical CD-ROM.

Prerequisites

- Verify that the target IQN is configured in the iBFT BIOS target parameter setting. This setting is in the option ROM of the network interface card (NIC) to be used for the iSCSI LUN. See the vendor documentation for your system.
- Disable the iBFT adapter option to boot to the iSCSI target. This action is necessary to make sure that the ESXi installer boots, rather than the iSCSI target. When you start your system, follow the prompt to log in to your iBFT adapter and disable the option to boot to the iSCSI target. See the vendor documentation for your system and iBFT adapter. After you finish the ESXi installation, you can reenble the option to boot from the LUN you install ESXi on.

Procedure

- 1 Start an interactive installation from the ESXi installation CD/DVD or mounted ISO image.
- 2 On the Select a Disk screen, select the iSCSI target you specified in the iBFT BIOS target parameter setting.

If the target does not appear in this menu, make sure that the TCP/IP and initiator iSCSI IQN settings are correct. Check the network Access Control List (ACL) and confirm that the adapter has adequate permissions to access the target.

- 3 Follow the prompts to complete the installation.
- 4 Reboot the host.
- 5 In the host BIOS settings, enter the iBFT adapter BIOS configuration, and change the adapter parameter to boot from the iSCSI target.

See the vendor documentation for your system.

What to do next

On your iBFT adapter, reenabling the option to boot to the iSCSI target, so the system will boot from the LUN you installed ESXi on.

Installing or Upgrading Hosts by Using a Script

You can quickly deploy ESXi hosts by using scripted, unattended installations or upgrades. Scripted installations or upgrades provide an efficient way to deploy multiple hosts.

The installation or upgrade script contains the installation settings for ESXi. You can apply the script to all hosts that you want to have a similar configuration.

For a scripted installation or upgrade, you must use the supported commands to create a script. You can edit the script to change settings that are unique for each host.

The installation or upgrade script can reside in one of the following locations:

- FTP server
- HTTP/HTTPS server
- NFS server
- USB flash drive
- CD-ROM drive

Approaches for Scripted Installation

You can install ESXi on multiple machines using a single script for all of them or a separate script for each machine.

For example, because disk names vary from machine to machine, one of the settings that you might want to configure in a script is the selection for the disk to install ESXi on.

Table 2-11. Scripted Installation Choices

Option	Action
Always install on the first disk on multiple machines.	Create one script.
Install ESXi on a different disk for each machine.	Create multiple scripts.

For information about the commands required to specify the disk to install on, see [Installation and Upgrade Script Commands](#).

Enter Boot Options to Start an Installation or Upgrade Script

You can start an installation or upgrade script by typing boot options at the ESXi installer boot command line.

At boot time you might need to specify options to access the kickstart file. You can enter boot options by pressing Shift+O in the boot loader. For a PXE boot installation, you can pass options through the `kerneLOPTS` line of the `boot.cfg` file. See [About the boot.cfg File](#) and [PXE Booting the ESXi Installer](#).

To specify the location of the installation script, set the `ks=filepath` option, where *filepath* indicates the location of your Kickstart file. Otherwise, a scripted installation or upgrade cannot start. If `ks=filepath` is omitted, the text installer is run.

Supported boot options are listed in [Boot Options](#).

Procedure

- 1 Start the host.
- 2 When the ESXi installer window appears, press Shift+O to edit boot options.



- 3 At the `runweasel` command prompt, type `ks=location of installation script plus boot command-line options`.

Example: Boot Option

You type the following boot options:

```
ks=http://00.00.00.00/kickstart/ks-osdc-pdp101.cfg nameserver=00.00.0.0 ip=00.00.00.000
netmask=255.255.255.0 gateway=00.00.00.000
```

Boot Options

When you perform a scripted installation, you might need to specify options at boot time to access the kickstart file.

Supported Boot Options

Table 2-12. Boot Options for ESXi Installation

Boot Option	Description
<code>BOOTIF=<i>hwtype-MAC address</i></code>	Similar to the <code>netdevice</code> option, except in the PXELINUX format as described in the <code>IPAPPEND</code> option under SYSLINUX at the syslinux.zytor.com site.
<code>gateway=<i>ip address</i></code>	Sets this network gateway as the default gateway to be used for downloading the installation script and installation media.
<code>ip=<i>ip address</i></code>	Sets up a static IP address to be used for downloading the installation script and the installation media. Note: the PXELINUX format for this option is also supported. See the <code>IPAPPEND</code> option under SYSLINUX at the syslinux.zytor.com site.

Table 2-12. Boot Options for ESXi Installation (Continued)

Boot Option	Description
<code>ks=cdrom:/path</code>	<p>Performs a scripted installation with the script at <i>path</i>, which resides on the CD in the CD-ROM drive. Each CDROM is mounted and checked until the file that matches the path is found.</p> <p>Important If you have created an installer ISO image with a custom installation or upgrade script, you must use uppercase characters to provide the path of the script, for example, <code>ks=cdrom:/KS_CUST.CFG</code>.</p>
<code>ks=file://path</code>	Performs a scripted installation with the script at <i>path</i> .
<code>ks=protocol://serverpath</code>	<p>Performs a scripted installation with a script located on the network at the given URL. <i>protocol</i> can be <code>http</code>, <code>https</code>, <code>ftp</code>, or <code>nfs</code>. An example using <code>nfs</code> protocol is <code>ks=nfs://host/porturl-path</code>. The format of an NFS URL is specified in RFC 2224.</p>
<code>ks=usb</code>	Performs a scripted installation, accessing the script from an attached USB drive. Searches for a file named <code>ks.cfg</code> . The file must be located in the root directory of the drive. If multiple USB flash drives are attached, they are searched until the <code>ks.cfg</code> file is found. Only FAT16 and FAT32 file systems are supported.
<code>ks=usb:/path</code>	Performs a scripted installation with the script file at the specified path, which resides on USB.
<code>ksdevice=device</code>	Tries to use a network adapter <i>device</i> when looking for an installation script and installation media. Specify as a MAC address, for example, <code>00:50:56:C0:00:01</code> . This location can also be a <code>vmnicNN</code> name. If not specified and files need to be retrieved over the network, the installer defaults to the first discovered network adapter that is plugged in.
<code>nameserver=ip address</code>	Specifies a domain name server to be used for downloading the installation script and installation media.
<code>netdevice=device</code>	Tries to use a network adapter <i>device</i> when looking for an installation script and installation media. Specify as a MAC address, for example, <code>00:50:56:C0:00:01</code> . This location can also be a <code>vmnicNN</code> name. If not specified and files need to be retrieved over the network, the installer defaults to the first discovered network adapter that is plugged in.
<code>netmask=subnet mask</code>	Specifies subnet mask for the network interface that downloads the installation script and the installation media.
<code>vlanid=vlanid</code>	Configure the network card to be on the specified VLAN.

About Installation and Upgrade Scripts

The installation/upgrade script is a text file, for example `ks.cfg`, that contains supported commands.

The command section of the script contains the ESXi installation options. This section is required and must appear first in the script.

About the Default ks.cfg Installation Script

The ESXi installer includes a default installation script that performs a standard installation to the first detected disk.

The default `ks.cfg` installation script is located in the initial RAM disk at `/etc/vmware/weasel/ks.cfg`. You can specify the location of the default `ks.cfg` file with the `ks=file:///etc/vmware/weasel/ks.cfg` boot option. See [Enter Boot Options to Start an Installation or Upgrade Script](#).

When you install ESXi using the `ks.cfg` script, the default root password is `mypassword`.

You cannot modify the default script on the installation media. After the installation, you can use the vSphere Web Client to log in to the vCenter Server that manages the ESXi host and modify the default settings.

The default script contains the following commands:

```
#
# Sample scripted installation file
#

# Accept the VMware End User License Agreement
vmaccepteula

# Set the root password for the DCUI and Tech Support Mode
rootpw mypassword

# Install on the first local disk available on machine
install --firstdisk --overwritevmfs

# Set the network to DHCP on the first network adapter
network --bootproto=dhcp --device=vmnic0

# A sample post-install script
%post --interpreter=python --ignorefailure=true
import time
stampFile = open('/finished.stamp', mode='w')
stampFile.write( time.asctime() )
```

Locations Supported for Installation or Upgrade Scripts

In scripted installations and upgrades, the ESXi installer can access the installation or upgrade script, also called the kickstart file, from several locations.

The following locations are supported for the installation or upgrade script:

- CD/DVD. See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#).
- USB Flash drive. See [Create a USB Flash Drive to Store the ESXi Installation Script or Upgrade Script](#).
- A network location accessible through the following protocols: NFS, HTTP, HTTPS, FTP

Path to the Installation or Upgrade Script

You can specify the path to an installation or upgrade script.

`ks=http://XXX.XXX.XXX.XXX/kickstart/KS.CFG` is the path to the ESXi installation script, where `XXX.XXX.XXX.XXX` is the IP address of the machine where the script resides. See [About Installation and Upgrade Scripts](#).

To start an installation script from an interactive installation, you enter the `ks=` option manually. See [Enter Boot Options to Start an Installation or Upgrade Script](#).

Installation and Upgrade Script Commands

To modify the default installation or upgrade script or to create your own script, use supported commands. Use supported commands in the installation script, which you specify with a boot command when you boot the installer.

To determine which disk to install or upgrade ESXi on, the installation script requires one of the following commands: `install`, `upgrade`, or `installorupgrade`. The `install` command creates the default partitions, including a VMFS datastore that occupies all available space after the other partitions are created.

accepteula or vmaccepteula (required)

Accepts the ESXi license agreement.

clearpart (optional)

Clears any existing partitions on the disk. Requires the `install` command to be specified. Carefully edit the `clearpart` command in your existing scripts.

<code>--drives=</code>	Remove partitions on the specified drives.
<code>--alldrives</code>	Ignores the <code>--drives=</code> requirement and allows clearing of partitions on every drive.
<code>--ignoredrives=</code>	Removes partitions on all drives except those specified. Required unless the <code>--drives=</code> or <code>--alldrives</code> flag is specified.
<code>--overwritevmfs</code>	Allows overwriting of VMFS partitions on the specified drives. By default, overwriting VMFS partitions is not allowed.
<code>--firstdisk= disk-type1 [disk-type2, ...]</code>	Partitions the first eligible disk found. By default, the eligible disks are set to the following order: <ol style="list-style-type: none"> 1 Locally attached storage (<code>local</code>) 2 Network storage (<code>remote</code>) 3 USB disks (<code>usb</code>)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including `esx` for the first disk with ESXi installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name `ST3120814A` and any disk that uses the `mptsas` driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use `localesx` for local storage that contains ESXi image or `remoteesx` for remote storage that contains ESXi image.

dryrun (optional)

Parses and checks the installation script. Does not perform the installation.

install

Specifies that this is a fresh installation. Replaces the deprecated `autopart` command used for ESXi 4.1 scripted installations. Either the `install`, `upgrade`, or `installorupgrade` command is required to determine which disk to install or upgrade ESXi on.

`--disk=` or `--drive=` Specifies the disk to partition. In the command `--disk=diskname`, the `diskname` can be in any of the forms shown in the following examples:

- Path: `--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0`
- MPX name: `--disk=mpx.vmhba1:C0:T0:L0`
- VML name: `--disk=vmL.000000034211234`
- vmkLUN UID: `--disk=vmkLUN_UID`

For accepted disk name formats, see [Disk Device Names](#).

`--firstdisk=`
`disk-type1,`
`[disk-type2,...]` Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

- 1 Locally attached storage (`local`)
- 2 Network storage (`remote`)
- 3 USB disks (`usb`)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including `esx` for the first disk with ESX installed on it, model and vendor information, or the name of the vmkernel device driver. For example, to prefer a disk with the model name `ST3120814A` and any disk that uses the `mptsas`

driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use `local` for local storage that contains ESXi image or `remoteesx` for remote storage that contains ESXi image.

--ignoressd Excludes solid-state disks from eligibility for partitioning. This option can be used with the `install` command and the `--firstdisk` option. This option takes precedence over the `--firstdisk` option. This option is invalid with the `--drive` or `--disk` options and with the `upgrade` and `installorupgrade` commands. See the *vSphere Storage* documentation for more information about preventing SSD formatting during auto-partitioning.

--overwritevsan You must use the `--overwritevsan` option when you install ESXi on a disk, either SSD or HDD (magnetic), that is in a vSAN disk group. If you use this option and no vSAN partition is on the selected disk, the installation will fail. When you install ESXi on a disk that is in vSAN disk group, the result depends on the disk that you select:

- If you select an SSD, the SSD and all underlying HDDs in the same disk group will be wiped.
- If you select an HDD, and the disk group size is greater than two, only the selected HDD will be wiped.
- If you select an HDD disk, and the disk group size is two or less, the SSD and the selected HDD will be wiped.

For more information about managing vSAN disk groups, see the *vSphere Storage* documentation.

--overwritevmfs Required to overwrite an existing VMFS datastore on the disk before installation.

--preservevmfs Preserves an existing VMFS datastore on the disk during installation.

--novmfsdisk Prevents a VMFS partition from being created on this disk. Must be used with `--overwritevmfs` if a VMFS partition already exists on the disk.

installorupgrade

Either the `install`, `upgrade`, or `installorupgrade` command is required to determine which disk to install or upgrade ESXi on.

--disk= or --drive= Specifies the disk to partition. In the command `--disk=diskname`, the *diskname* can be in any of the forms shown in the following examples:

- Path: `--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0`
- MPX name: `--disk=mpx.vmhba1:C0:T0:L0`

- VML name: `--disk=vm1.000000034211234`
- vmkLUN UID: `--disk=vmkLUN_UID`

For accepted disk name formats, see [Disk Device Names](#).

`--firstdisk=
disk-type1,
[disk-type2,...]`

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

- 1 Locally attached storage (`local`)
- 2 Network storage (`remote`)
- 3 USB disks (`usb`)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including `esx` for the first disk with ESX installed on it, model and vendor information, or the name of the vmkernel device driver. For example, to prefer a disk with the model name `ST3120814A` and any disk that uses the `mptsas` driver rather than a normal local disk, the argument is

`--firstdisk=ST3120814A,mptsas,local`. You can use `localeSX` for local storage that contains ESXi image or `remoteesx` for remote storage that contains ESXi image.

`--overwritevsan`

You must use the `--overwritevsan` option when you install ESXi on a disk, either SSD or HDD (magnetic), that is in a vSAN disk group. If you use this option and no vSAN partition is on the selected disk, the installation will fail. When you install ESXi on a disk that is in a vSAN disk group, the result depends on the disk that you select:

- If you select an SSD, the SSD and all underlying HDDs in the same disk group will be wiped.
- If you select an HDD, and the disk group size is greater than two, only the selected HDD will be wiped.
- If you select an HDD disk, and the disk group size is two or less, the SSD and the selected HDD will be wiped.

For more information about managing vSAN disk groups, see the *vSphere Storage* documentation.

`--overwritevmfs`

Install ESXi if a VMFS partition exists on the disk, but no ESX or ESXi installation exists. Unless this option is present, the installer will fail if a VMFS partition exists on the disk, but no ESX or ESXi installation exists.

keyboard (optional)

Sets the keyboard type for the system.

keyboardType Specifies the keyboard map for the selected keyboard type. *keyboardType* must be one of the following types.

- Belgian
- Brazilian
- Croatian
- Czechoslovakian
- Danish
- Estonian
- Finnish
- French
- German
- Greek
- Icelandic
- Italian
- Japanese
- Latin American
- Norwegian
- Polish
- Portuguese
- Russian
- Slovenian
- Spanish
- Swedish
- Swiss French
- Swiss German
- Turkish
- Ukrainian
- United Kingdom
- US Default

- US Dvorak

serialnum or vmserialnum (optional)

Deprecated in ESXi 5.0.x. Supported in ESXi 5.1 and later. Configures licensing. If not included, ESXi installs in evaluation mode.

`--esx=<license-key>` Specifies the vSphere license key to use. The format is 5 five-character groups (XXXXX-XXXXX-XXXXX-XXXXX-XXXXX).

network (optional)

Specifies a network address for the system.

`--bootproto=[dhcp|static]` Specifies whether to obtain the network settings from DHCP or set them manually.

`--device=` Specifies either the MAC address of the network card or the device name, in the form `vmnicNN`, as in `vmnic0`. This options refers to the uplink device for the virtual switch.

`--ip=` Sets an IP address for the machine to be installed, in the form `xxx.xxx.xxx.xxx`. Required with the `--bootproto=static` option and ignored otherwise.

`--gateway=` Designates the default gateway as an IP address, in the form `xxx.xxx.xxx.xxx`. Used with the `--bootproto=static` option.

`--nameserver=` Designates the primary name server as an IP address. Used with the `--bootproto=static` option. Omit this option if you do not intend to use DNS.

The `--nameserver` option can accept two IP addresses. For example: `--nameserver="10.126.87.104[,10.126.87.120]"`

`--netmask=` Specifies the subnet mask for the installed system, in the form `255.xxx.xxx.xxx`. Used with the `--bootproto=static` option.

`--hostname=` Specifies the host name for the installed system.

`--vlanid= vlanid` Specifies which VLAN the system is on. Used with either the `--bootproto=dhcp` or `--bootproto=static` option. Set to an integer from 1 to 4096.

`--addvmportgroup=(0|1)` Specifies whether to add the VM Network port group, which is used by virtual machines. The default value is 1.

paranoid (optional)

Causes warning messages to interrupt the installation. If you omit this command, warning messages are logged.

part or partition (optional)

Creates an additional VMFS datastore on the system. Only one datastore per disk can be created. Cannot be used on the same disk as the `install` command. Only one partition can be specified per disk and it can only be a VMFS partition.

<i>datastore name</i>	Specifies where the partition is to be mounted.
<code>--ondisk=</code> or <code>--ondrive=</code>	Specifies the disk or drive where the partition is created.
<code>--firstdisk=</code> <i>disk-type1,</i> <i>[disk-type2, ...]</i>	Partitions the first eligible disk found. By default, the eligible disks are set to the following order: <ol style="list-style-type: none"> 1 Locally attached storage (<code>local</code>) 2 Network storage (<code>remote</code>) 3 USB disks (<code>usb</code>) <p>You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including <code>esx</code> for the first disk with ESX installed on it, model and vendor information, or the name of the vmkernel device driver. For example, to prefer a disk with the model name <code>ST3120814A</code> and any disk that uses the <code>mptsas</code> driver rather than a normal local disk, the argument is <code>--firstdisk=ST3120814A,mptsas,local</code>. You can use <code>localesex</code> for local storage that contains ESXi image or <code>remoteesx</code> for remote storage that contains ESXi image.</p>

reboot (optional)

Reboots the machine after the scripted installation is complete.

`<--noeject>` The CD is not ejected after the installation.

rootpw (required)

Sets the root password for the system.

`--iscrypted` Specifies that the password is encrypted.

password Specifies the password value.

upgrade

Either the `install`, `upgrade`, or `installorupgrade` command is required to determine which disk to install or upgrade ESXi on.

`--disk=` or `--drive=` Specifies the disk to partition. In the command `--disk=diskname`, the *diskname* can be in any of the forms shown in the following examples:

- Path: `--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0`
- MPX name: `--disk=mpx.vmhba1:C0:T0:L0`
- VML name: `--disk=vm1.000000034211234`
- vmkLUN UID: `--disk=vmkLUN_UID`

For accepted disk name formats, see [Disk Device Names](#).

`--firstdisk=`
`disk-type1,`
`[disk-type2,...]` Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

- 1 Locally attached storage (`local`)
- 2 Network storage (`remote`)
- 3 USB disks (`usb`)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including `esx` for the first disk with ESX installed on it, model and vendor information, or the name of the vmkernel device driver. For example, to prefer a disk with the model name `ST3120814A` and any disk that uses the `mptsas` driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use `localesx` for local storage that contains ESXi image or `remoteesx` for remote storage that contains ESXi image.

%include or include (optional)

Specifies another installation script to parse. This command is treated similarly to a multiline command, but takes only one argument.

`filename` For example: `%include part.cfg`

%pre (optional)

Specifies a script to run before the kickstart configuration is evaluated. For example, you can use it to generate files for the kickstart file to include.

`--interpreter` Specifies an interpreter to use. The default is `busybox`.

`=[python|busybox]`

%post (optional)

Runs the specified script after package installation is complete. If you specify multiple %post sections, they run in the order that they appear in the installation script.

<code>--interpreter</code>	Specifies an interpreter to use. The default is busybox.
<code>=[python busybox]</code>	
<code>--timeout=secs</code>	Specifies a timeout for running the script. If the script is not finished when the timeout expires, the script is forcefully terminated.
<code>--ignorefailure</code>	If true, the installation is considered a success even if the %post script terminated with an error.
<code>=[true false]</code>	

%firstboot

Creates an `init` script that runs only during the first boot. The script has no effect on subsequent boots. If multiple %firstboot sections are specified, they run in the order that they appear in the kickstart file.

Note You cannot check the semantics of %firstboot scripts until the system is booting for the first time. A %firstboot script might contain potentially catastrophic errors that are not exposed until after the installation is complete.

<code>--interpreter</code>	Specifies an interpreter to use. The default is busybox.
<code>=[python busybox]</code>	

Note You cannot check the semantics of the %firstboot script until the system boots for the first time. If the script contains errors, they are not exposed until after the installation is complete.

Disk Device Names

The `install`, `upgrade`, and `installorupgrade` installation script commands require the use of disk device names.

Table 2-13. Disk Device Names

Format	Example	Description
VML	vml.00025261	The device name as reported by the VMkernel
MPX	mpx.vmhba0:C0:T0:L0	The device name

About the boot.cfg File

The boot loader configuration file `boot.cfg` specifies the kernel, the kernel options, and the boot modules that the `mboot.c32` or `mboot.efi` boot loader uses in an ESXi installation.

The `boot.cfg` file is provided in the ESXi installer. You can modify the `kernelopt` line of the `boot.cfg` file to specify the location of an installation script or to pass other boot options.

The `boot.cfg` file has the following syntax:

```
# boot.cfg -- mboot configuration file
#
# Any line preceded with '#' is a comment.

title=STRING
prefix=DIRPATH
kernel=FILEPATH
kernelopt=STRING
modules=FILEPATH1 --- FILEPATH2... --- FILEPATHn

# Any other line must remain unchanged.
```

The commands in `boot.cfg` configure the boot loader.

Table 2-14. Commands in `boot.cfg` .

Command	Description
<code>title=STRING</code>	Sets the boot loader title to <i>STRING</i> .
<code>prefix=STRING</code>	(Optional) Adds <i>DIRPATH</i> / in front of every <i>FILEPATH</i> in the <code>kernel=</code> and <code>modules=</code> commands that do not already start with / or with <code>http://</code> .
<code>kernel=FILEPATH</code>	Sets the kernel path to <i>FILEPATH</i> .
<code>kernelopt=STRING</code>	Appends <i>STRING</i> to the kernel boot options.
<code>modules=FILEPATH1 --- FILEPATH2... --- FILEPATHn</code>	Lists the modules to be loaded, separated by three hyphens (---).

See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#) and [PXE Booting the ESXi Installer](#).

Install or Upgrade ESXi from a CD or DVD by Using a Script

You can install or upgrade ESXi from a CD-ROM or DVD-ROM drive by using a script that specifies the installation or upgrade options.

You can start the installation or upgrade script by entering a boot option when you start the host. You can also create an installer ISO image that includes the installation script. With an installer ISO image, you can perform a scripted, unattended installation when you boot the resulting installer ISO image. See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#).

Prerequisites

Before you run the scripted installation or upgrade, verify that the following prerequisites are met:

- The system on which you are installing or upgrading meets the hardware requirements. See [ESXi Hardware Requirements](#).
- You have the ESXi installer ISO on an installation CD or DVD . See [Download and Burn the ESXi Installer ISO Image to a CD or DVD](#).

- The default installation or upgrade script (`ks.cfg`) or a custom installation or upgrade script is accessible to the system. See [About Installation and Upgrade Scripts](#).
- You have selected a boot command to run the scripted installation or upgrade. See [Enter Boot Options to Start an Installation or Upgrade Script](#). For a complete list of boot commands, see [Boot Options](#).

Procedure

- 1 Boot the ESXi installer from the local CD-ROM or DVD-ROM drive.
- 2 When the ESXi installer window appears, press Shift+O to edit boot options.



- 3 Type a boot option that calls the default installation or upgrade script or an installation or upgrade script file that you created.

The boot option has the form `ks=`.

- 4 Press Enter.

The installation, upgrade, or migration runs, using the options that you specified.

Install or Upgrade ESXi from a USB Flash Drive by Using a Script

You can install or upgrade ESXi from a USB flash drive by using a script that specifies the installation or upgrade options.

Supported boot options are listed in [Boot Options](#).

Prerequisites

Before running the scripted installation or upgrade, verify that the following prerequisites are met:

- The system that you are installing or upgrading to ESXi meets the hardware requirements for the installation or upgrade. See [ESXi Hardware Requirements](#).
- You have the ESXi installer ISO on a bootable USB flash drive. See [Format a USB Flash Drive to Boot the ESXi Installation or Upgrade](#).
- The default installation or upgrade script (`ks.cfg`) or a custom installation or upgrade script is accessible to the system. See [About Installation and Upgrade Scripts](#).

- You have selected a boot option to run the scripted installation, upgrade, or migration. See [Enter Boot Options to Start an Installation or Upgrade Script](#).

Procedure

- 1 Boot the ESXi installer from the USB flash drive.
- 2 When the ESXi installer window appears, press Shift+O to edit boot options.



- 3 Type a boot option that calls the default installation or upgrade script or an installation or upgrade script file that you created.

The boot option has the form ks=.

- 4 Press Enter.

The installation, upgrade, or migration runs, using the options that you specified.

Performing a Scripted Installation or Upgrade of ESXi by Using PXE to Boot the Installer

ESXi 6.5 provides many options for using PXE to boot the installer and using an installation or upgrade script.

- For information about setting up a PXE infrastructure, see [PXE Booting the ESXi Installer](#).
- For information about creating and locating an installation script, see [About Installation and Upgrade Scripts](#).
- For specific procedures to use PXE to boot the ESXi installer and use an installation script, see one of the following topics:
 - [PXE Boot the ESXi Installer Using a Web Server](#)
 - [PXE Boot the ESXi Installer Using TFTP](#)
- For information about using vSphere Auto Deploy to perform a scripted installation by using PXE to boot, see [Installing ESXi Using vSphere Auto Deploy](#).

PXE Booting the ESXi Installer

You can use the preboot execution environment (PXE) to boot a host. Starting with vSphere 6.0, you can PXE boot the ESXi installer from a network interface on hosts with legacy BIOS or using UEFI.

ESXi is distributed in an ISO format that is designed to install to flash memory or to a local hard drive. You can extract the files and boot by using PXE.

PXE uses Dynamic Host Configuration Protocol (DHCP) and Trivial File Transfer Protocol (TFTP) to boot an operating system over a network.

PXE booting requires some network infrastructure and a machine with a PXE-capable network adapter. Most machines that can run ESXi have network adapters that can PXE boot.

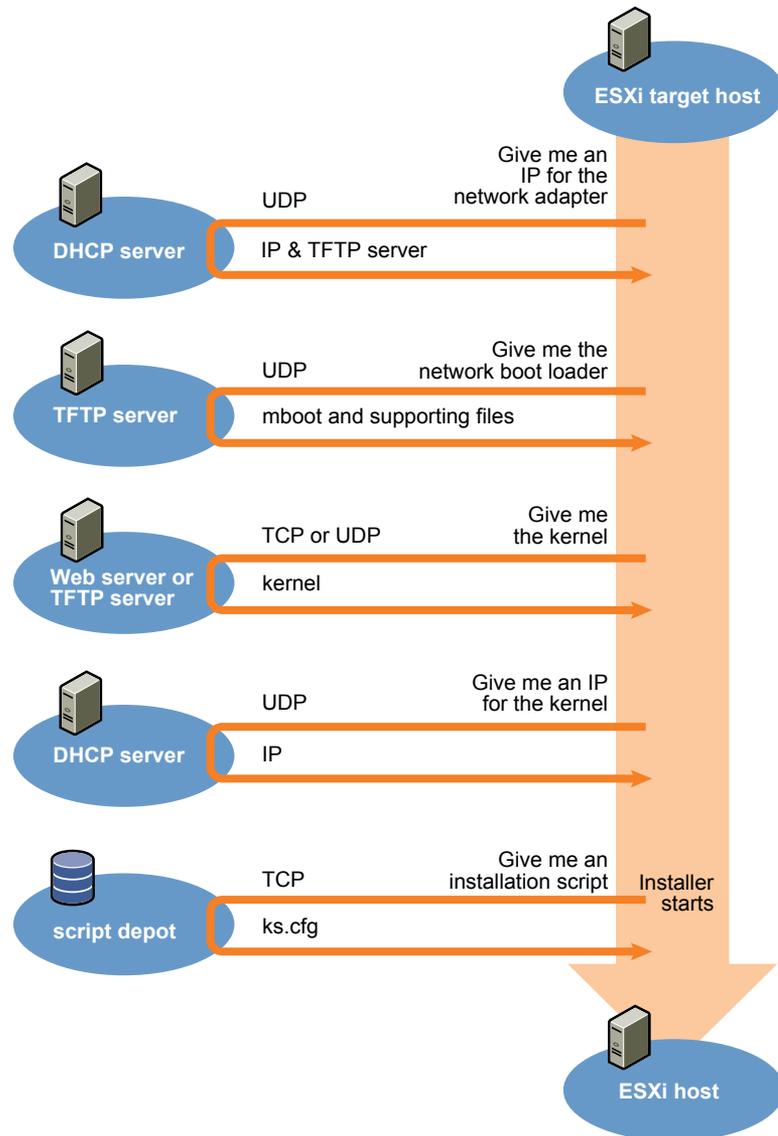
Note PXE booting with legacy BIOS firmware is possible only over IPv4. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.

Overview of the PXE Boot Installation Process

Some of the details of the PXE boot process vary depending on whether the target host is using legacy BIOS or UEFI firmware, and whether the boot process uses TFTP only or TFTP plus HTTP.

When you boot the target host, it interacts with the different servers in the environment to get the network adapter, boot loader, kernel, IP address for the kernel, and finally the installation script. When all components are in place, installation starts, as shown in the following illustration.

Figure 2-3. Overview of PXE Boot Installation Process



The interaction between the ESXi host and other servers proceeds as follows:

- 1 The user boots the target ESXi host.
- 2 The target ESXi host makes a DHCP request.
- 3 The DHCP server responds with the IP information and the location of the TFTP server.
- 4 The ESXi host contacts the TFTP server and requests the file that the DHCP server specified.
- 5 The TFTP server sends the network boot loader, and the ESXi host executes it. The initial boot loader might load additional boot loader components from the TFTP server.
- 6 The boot loader searches for a configuration file on the TFTP server, downloads the kernel and other ESXi components from the HTTP server or the TFTP server and boots the kernel on the ESXi host.
- 7 The installer runs interactively or using a kickstart script, as specified in the configuration file.

PXE Boot the ESXi Installer Using TFTP

You can use a TFTP server to PXE boot the ESXi installer. The process differs slightly depending on whether you use UEFI or boot from a legacy BIOS. Because most environments include ESXi hosts that support UEFI boot and hosts that support only legacy BIOS, this topic discusses prerequisites and steps for both types of hosts.

- For legacy BIOS machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `pxelinux.0` or `gpxelinux.0` initial boot loader for all target machines, but potentially different PXELINUX configuration files depending on the target machine's MAC address.
- For UEFI machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `mboot.efi` initial boot loader for all target machines, but potentially different `boot.cfg` files depending on the target machine's MAC address.

Prerequisites

Verify that your environment meets the following prerequisites.

- ESXi installer ISO image, downloaded from the VMware Web site.
- Target host with a hardware configuration that is supported for your version of ESXi. See the *VMware Compatibility Guide*.
- Network adapter with PXE support on the target ESXi host.
- DHCP server configured for PXE booting. See [Sample DHCP Configurations](#).
- TFTP server.
- Network security policies to allow TFTP traffic (UDP port 69).
- For legacy BIOS, you can use only IPv4 networking. For UEFI PXE boot, you can use IPv4 or IPv6 networking.
- (Optional) Installation script (kickstart file).
- Use a native VLAN in most cases. If you want to specify the VLAN ID to be used with PXE booting, check that your NIC supports VLAN ID specification.

For legacy BIOS systems, version 3.86 of the SYSLINUX package, available from <https://www.kernel.org/pub/linux/utils/boot/syslinux/>.

Procedure

- 1 Configure the DHCP server for TFTP boot.

2 (Legacy BIOS only) Obtain and configure PXELINUX:

- a Obtain SYSLINUX version 3.86, unpack it, and copy the `pxelinux.0` file to the top-level `/tftpboot` directory on your TFTP server.
- b Create a PXELINUX configuration file using the following code model.

ESXi-6.x.x-XXXXXX is the name of the TFTP subdirectory that contains the ESXi installer files.

```
DEFAULT install
NOHALT 1
LABEL install
  KERNEL ESXi-6.x.x-XXXXXX/mboot.c32
  APPEND -c ESXi-6.x.x-XXXXXX/boot.cfg
  IPAPPEND 2
```

- c Save the PXELINUX file in the `/tftpboot/pxelinux.cfg` directory on your TFTP server with a filename that will determine whether all hosts boot this installer by default:

Option	Description
Same installer	Name the file <code>default</code> if you want for all host to boot this ESXi installer by default.
Different installers	Name the file with the MAC address of the target host machine (01- <i>mac_address_of_target_ESXi_host</i>) if you want only a specific host to boot with this file, for example, 01-23-45-67-89-0a-bc.

- 3 (UEFI only) Copy the file `efi/boot/bootx64.efi` from the ESXi installer ISO image to `/tftpboot/mboot.efi` on your TFTP server.

Note Newer versions of `mboot.efi` can generally boot older versions of ESXi, but older versions of `mboot.efi` might be unable to boot newer versions of ESXi. If you plan to configure different hosts to boot different versions of the ESXi installer, use the `mboot.efi` from the newest version.

- 4 Create a subdirectory of your TFTP server's top-level `/tftpboot` directory and name it after the version of ESXi it will hold, for example, `/tftpboot/ESXi-6.x.x-xxxxx`.
- 5 Copy the contents of the ESXi installer image to the directory you just created.
- 6 Modify the `boot.cfg` file
 - a Add the following line:

```
prefix=ESXi-6.x.x-xxxxxx
```

Here, `ESXi-6.x.x-xxxxxx` is the pathname of the installer files relative to the TFTP server's root directory.

- b If the filenames in the `kernel=` and `modules=` lines begin with a forward slash (`/`) character, delete that character.

- 7 (Optional) For a scripted installation, in the `boot.cfg` file, add the `kernelopt` option to the line after the kernel command, to specify the location of the installation script.

Use the following code as a model, where `XXX.XXX.XXX.XXX` is the IP address of the server where the installation script resides, and `esxi_ksFiles` is the directory that contains the `ks.cfg` file.

```
kernelopt=ks=http://XXX.XXX.XXX.XXX/esxi_ksFiles/ks.cfg
```

- 8 (UEFI only) Specify whether you want for all UEFI hosts to boot the same installer.

Option	Description
Same installer	Copy or link the <code>boot.cfg</code> file to <code>/tftpboot/boot.cfg</code>
Different installers	<ul style="list-style-type: none"> a Create a subdirectory of <code>/tftpboot</code> named after the MAC address of the target host machine (<code>01-mac_address_of_target_ESXi_host</code>), for example, <code>01-23-45-67-89-0a-bc</code>. b Place a copy of (or a link to) the host's <code>boot.cfg</code> file in that directory, for example, <code>/tftpboot/01-23-45-67-89-0a-bc/boot.cfg</code>.

PXE Boot the ESXi Installer Using a Web Server

You can use a Web server to PXE boot the ESXi installer. Because most environments include ESXi hosts that support UEFI boot and hosts that support only legacy BIOS, this topic discusses prerequisites and steps for both types of hosts.

- For legacy BIOS machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `pxelinux.0` or `gpxelinux.0` initial boot loader for all target machines, but potentially different PXELINUX configuration files depending on the target machine's MAC address.
- For UEFI machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `mboot.efi` initial boot loader for all target machines, but potentially different `boot.cfg` files depending on the target machine's MAC address.

Prerequisites

Verify that your environment has the following components:

- ESXi installer ISO image, downloaded from the VMware Web site.
- Target host with a hardware configuration that is supported for your version of ESXi. See the *VMware Compatibility Guide*.
- Network adapter with PXE support on the target ESXi host.
- DHCP server configured for PXE booting. See [Sample DHCP Configurations](#).
- TFTP server.
- Network security policies to allow TFTP traffic (UDP port 69).
- For legacy BIOS, you can use only IPv4 networking. For UEFI PXE boot, you can use IPv4 or IPv6 networking.

- (Optional) Installation script (kickstart file).
- Use a native VLAN in most cases. If you want to specify the VLAN ID to be used with PXE booting, check that your NIC supports VLAN ID specification.

Verify that your environment also meets the following prerequisites required for PXE boot using a Web Server:

- Verify that the HTTP Web server is accessible by your target ESXi hosts.
- (UEFI) Obtain iPXE, available at <http://ipxe.org>.
- (Legacy BIOS) Obtain version 3.86 of the SYSLINUX package, available from <https://www.kernel.org/pub/linux/utils/boot/syslinux/>.

Procedure

- 1 Configure the DHCP server for HTTP boot.
- 2 (UEFI only) Obtain and configure iPXE:
 - a Obtain the iPXE source code, as described at <http://ipxe.org/download>.
 - b Follow the instructions on that page, but use the following make command:

```
make bin-x86_64-efi/snponly.efi
```
 - c Copy the resulting file `snponly.efi` to `/tftpboot` directory on your TFTP server.
- 3 (UEFI only) Copy the file `efi/boot/bootx64.efi` from the ESXi installer ISO image to `/tftpboot/mboot.efi` on your TFTP server.

Note Newer versions of `mboot.efi` can generally boot older versions of ESXi, but older versions of `mboot.efi` might be unable to boot newer versions of ESXi. If you plan to configure different hosts to boot different versions of the ESXi installer, use the `mboot.efi` from the newest version.

4 (Legacy BIOS only) Obtain and configure PXELINUX:

- a Obtain SYSLINUX version 3.86, unpack it, and copy the `gpxelinux.0` file to the top-level `/tftpboot` directory on your TFTP server.
- b Create a PXELINUX configuration file using the following code model.

ESXi-6.x.x-XXXXXX is the name of the TFTP subdirectory that contains the ESXi installer files.

```
DEFAULT install
NOHALT 1
LABEL install
  KERNEL ESXi-6.x.x-XXXXXX/mboot.c32
  APPEND -c ESXi-6.x.x-XXXXXX/boot.cfg
  IPAPPEND 2
```

- c Save the PXELINUX file in the `/tftpboot/pxelinux.cfg` directory on your TFTP server with a filename that will determine whether all hosts boot this installer by default:

Option	Description
Same installer	Name the file <code>default</code> if you want for all host to boot this ESXi installer by default.
Different installers	Name the file with the MAC address of the target host machine (01- <i>mac_address_of_target_ESXi_host</i>) if you want only a specific host to boot with this file, for example, 01-23-45-67-89-0a-bc.

- 5 Create a directory on your HTTP server named for the version of ESXi it will hold, for example, `/var/www/html/ESXi-6.x.x-XXXXXX`.

- 6 Copy the contents of the ESXi installer image to the directory you just created.

- 7 Modify the `boot.cfg` file

- a Add the following line:

```
prefix=http://XXX.XXX.XXX.XXX/ESXi-6.x.x-XXXXXX
```

where `http://XXX.XXX.XXX.XXX/ESXi-6.x.x-XXXXXX` is the location of the installer files on the HTTP server.

- b If the filenames in the `kernel=` and `modules=` lines begin with a forward slash (`/`) character, delete that character.

- 8 (Optional) For a scripted installation, in the `boot.cfg` file, add the `kernelopt` option to the line after the kernel command, to specify the location of the installation script.

Use the following code as a model, where `XXX.XXX.XXX.XXX` is the IP address of the server where the installation script resides, and `esxi_ksFiles` is the directory that contains the `ks.cfg` file.

```
kernelopt=ks=http://XXX.XXX.XXX.XXX/esxi_ksFiles/ks.cfg
```

9 (UEFI only) Specify whether you want for all UEFI hosts to boot the same installer.

Option	Description
Same installer	Copy or link the <code>boot.cfg</code> file to <code>/tftpboot/boot.cfg</code>
Different installers	<ol style="list-style-type: none"> Create a subdirectory of <code>/tftpboot</code> named after the MAC address of the target host machine (<code>01-mac_address_of_target_ESXi_host</code>), for example, <code>01-23-45-67-89-0a-bc</code>. Place a copy of (or a link to) the host's <code>boot.cfg</code> file in that directory, for example, <code>/tftpboot/01-23-45-67-89-0a-bc/boot.cfg</code>.

Installing ESXi Using vSphere Auto Deploy

vSphere Auto Deploy lets you provision hundreds of physical hosts with ESXi software.

Using Auto Deploy, experienced system administrators can manage large deployments efficiently. Hosts are network-booted from a central Auto Deploy server. Optionally, hosts are configured with a host profile of a reference host. The host profile can be set up to prompt the user for input. After boot up and configuration complete, the hosts are managed by vCenter Server just like other ESXi hosts.

Auto Deploy can also be used for stateless caching or stateful installs.

Important Auto Deploy requires a secure separation between the production network and the management or deployment networks as discussed in [vSphere Auto Deploy Security Considerations](#). Using Auto Deploy without this separation is insecure.

Stateless caching By default, Auto Deploy does not store ESXi configuration or state on the host disk. Instead, an image profile defines the image that the host is provisioned with, and other host attributes are managed through host profiles. A host that uses Auto Deploy for stateless caching still needs to connect to the Auto Deploy server and the vCenter Server.

Stateful installs You can provision a host with Auto Deploy and set up the host to store the image to disk. On subsequent boots, the host boots from disk.

Understanding vSphere Auto Deploy

vSphere Auto Deploy can provision hundreds of physical hosts with ESXi software. You can specify the image to deploy and the hosts to provision with the image. Optionally, you can specify host profiles to apply to the hosts, a vCenter Server location (datacenter, folder or cluster), and assign a script bundle for each host.

Introduction to vSphere Auto Deploy

When you start a physical host that is set up for vSphere Auto Deploy, vSphere Auto Deploy uses PXE boot infrastructure in conjunction with vSphere host profiles to provision and customize that host. No state is stored on the host itself. Instead, the vSphere Auto Deploy server manages state information for each host.

State Information for ESXi Hosts

vSphere Auto Deploy stores the information for the ESXi hosts to be provisioned in different locations. Information about the location of image profiles and host profiles is initially specified in the rules that map machines to image profiles and host profiles.

Table 2-15. vSphere Auto Deploy Stores Information for Deployment

Information Type	Description	Source of Information
Image state	The executable software to run on an ESXi host.	Image profile, created with vSphere ESXi Image Builder.
Configuration state	The configurable settings that determine how the host is configured, for example, virtual switches and their settings, driver settings, boot parameters, and so on.	Host profile, created by using the host profile UI. Often comes from a template host.
Dynamic state	The runtime state that is generated by the running software, for example, generated private keys or runtime databases.	Host memory, lost during reboot.
Virtual machine state	The virtual machines stored on a host and virtual machine autostart information (subsequent boots only).	Virtual machine information sent by vCenter Server to vSphere Auto Deploy must be available to supply virtual machine information to vSphere Auto Deploy.
User input	State that is based on user input, for example, an IP address that the user provides when the system starts up, cannot automatically be included in the host profile.	Host customization information, stored by vCenter Server during first boot. You can create a host profile that requires user input for certain values. When vSphere Auto Deploy applies a host profile that requires user provided information, the host is placed in maintenance mode. Use the host profile UI to check the host profile compliance, and respond to the prompt to customize the host.

vSphere Auto Deploy Architecture

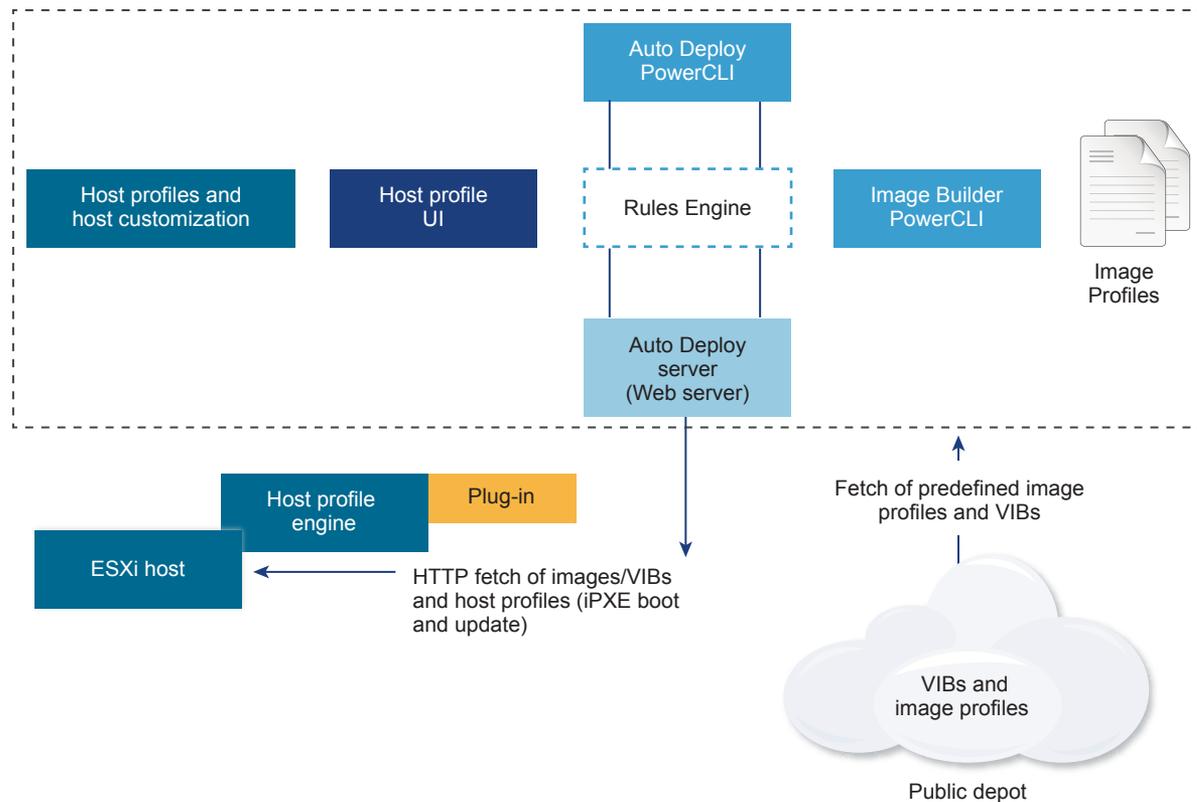
The vSphere Auto Deploy infrastructure consists of several components.

For more information, watch the video "Auto Deploy Architecture":



Auto Deploy Architecture (http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_auto_deploy_architecture)

Figure 2-4. vSphere Auto Deploy Architecture



vSphere Auto Deploy server

Serves images and host profiles to ESXi hosts.

vSphere Auto Deploy rules engine

Sends information to the vSphere Auto Deploy server which image profile and which host profile to serve to which host. Administrators use vSphere Auto Deploy to define the rules that assign image profiles and host profiles to hosts. For more information on vSphere Auto Deploy rules and rule sets, see [Rules and Rule Sets](#).

Image profiles

Define the set of VIBs to boot ESXi hosts with.

- VMware and VMware partners make image profiles and VIBs available in public depots. Use vSphere ESXi Image Builder to examine the depot and use the vSphere Auto Deploy rules engine to specify which image profile to assign to which host.
- VMware customers can create a custom image profile based on the public image profiles and VIBs in the depot and apply that image profile to the host. See [Customizing Installations with vSphere ESXi Image Builder](#).

Host profiles

Define machine-specific configuration such as networking or storage setup. Use the host profile UI to create host profiles. You can create a host profile for a reference host and apply that host profile to other hosts in your environment for a consistent configuration. For more information, see the *vSphere Host Profiles* documentation or the [Setting Up a vSphere Auto Deploy Reference Host](#) section.

Host customization

Stores information that the user provides when host profiles are applied to the host. Host customization might contain an IP address or other information that the user supplied for that host. For more information about host customizations, see the *vSphere Host Profiles* documentation.

Host customization was called answer file in earlier releases of vSphere Auto Deploy.

Rules and Rule Sets

You specify the behavior of the vSphere Auto Deploy server by using a set of rules. The vSphere Auto Deploy rules engine checks the rule set for matching host patterns to decide which items (image profile, host profile, vCenter Server location, or script object) to provision each host with.

The rules engine maps software and configuration settings to hosts based on the attributes of the host. For example, you can deploy image profiles or host profiles to two clusters of hosts by writing two rules, each matching on the network address of one cluster.

For hosts that have not yet been added to a vCenter Server system, the vSphere Auto Deploy server checks with the rules engine before serving image profiles, host profiles, and inventory location information to hosts. For hosts that are managed by a vCenter Server system, the image profile, host profile, and inventory location that vCenter Server has stored in the host object is used. If you make changes to rules, you can use the vSphere Web Client or vSphere Auto Deploy cmdlets in a PowerCLI session to test and repair rule compliance. When you repair rule compliance for a host, that host's image profile and host profile assignments are updated.

The rules engine includes rules and rule sets.

Rules

Rules can assign image profiles and host profiles to a set of hosts, or specify the location (folder or cluster) of a host on the target vCenter Server system. A rule can identify target hosts by boot MAC address, SMBIOS information, BIOS UUID, Vendor, Model, or fixed DHCP IP address. In most cases, rules apply to multiple hosts. You create rules by using the vSphere Web Client or vSphere Auto Deploy cmdlets in a PowerCLI session. After you create a rule, you must add it to a rule set. Only two rule sets, the active rule set and the working rule set, are supported. A rule can belong to both sets, the default, or only to the working rule set. After you add a rule to a rule set, you can no longer change the rule. Instead, you copy the rule and replace items or patterns in the copy. If you are managing vSphere Auto Deploy with the vSphere Web Client, you can edit a rule if it is in inactive state.

You can specify the following parameters in a rule.

Parameter	Description																						
Name	Name of the rule, specified with the <code>-Name</code> parameter.																						
Item	One or more items, specified with the <code>-Item</code> parameter. An item can be an image profile, a host profile, a vCenter Server inventory location (datacenter, folder, cluster) for the target host, or a custom script. You can specify multiple items separated by commas.																						
Pattern	The pattern specifies the host or group of hosts to which the rule applies. <table border="0" style="margin-left: 20px;"> <tr> <td>vendor</td> <td>Machine vendor name.</td> </tr> <tr> <td>model</td> <td>Machine model name.</td> </tr> <tr> <td>serial</td> <td>Machine serial number.</td> </tr> <tr> <td>hostname</td> <td>Machine hostname.</td> </tr> <tr> <td>domain</td> <td>Domain name.</td> </tr> <tr> <td>ipv4</td> <td>IPv4 address of the machine.</td> </tr> <tr> <td>ipv6</td> <td>IPv6 address of the machine.</td> </tr> <tr> <td></td> <td>PXE booting with BIOS firmware is possible only with IPv4, PXE booting with UEFI firmware is possible with either IPv4 or IPv6.</td> </tr> <tr> <td>mac</td> <td>Boot NIC MAC address.</td> </tr> <tr> <td>asset</td> <td>Machine asset tag.</td> </tr> <tr> <td>oemstring</td> <td>OEM-specific strings in the SMBIOS.</td> </tr> </table> <p>You can specify <code>-AllHosts</code> to apply the item or items to all hosts.</p>	vendor	Machine vendor name.	model	Machine model name.	serial	Machine serial number.	hostname	Machine hostname.	domain	Domain name.	ipv4	IPv4 address of the machine.	ipv6	IPv6 address of the machine.		PXE booting with BIOS firmware is possible only with IPv4, PXE booting with UEFI firmware is possible with either IPv4 or IPv6.	mac	Boot NIC MAC address.	asset	Machine asset tag.	oemstring	OEM-specific strings in the SMBIOS.
vendor	Machine vendor name.																						
model	Machine model name.																						
serial	Machine serial number.																						
hostname	Machine hostname.																						
domain	Domain name.																						
ipv4	IPv4 address of the machine.																						
ipv6	IPv6 address of the machine.																						
	PXE booting with BIOS firmware is possible only with IPv4, PXE booting with UEFI firmware is possible with either IPv4 or IPv6.																						
mac	Boot NIC MAC address.																						
asset	Machine asset tag.																						
oemstring	OEM-specific strings in the SMBIOS.																						

Active Rule Set

When a newly started host contacts the vSphere Auto Deploy server with a request for an image profile, the vSphere Auto Deploy server checks the active rule set for matching rules. The image profile, host profile, vCenter Server inventory location, and script object that are mapped by matching rules are then used to boot the host. If more than one item of the same type is mapped by the rules, the vSphere Auto Deploy server uses the item that is first in the rule set.

Working Rule Set

The working rule set allows you to test changes to rules before making the changes active. For example, you can use vSphere Auto Deploy cmdlets for testing compliance with the working rule set. The test verifies that hosts managed by a vCenter Server system are following the rules in the working rule set. By default, cmdlets add the rule to the working rule set and activate the rules. Use the `NoActivate` parameter to add a rule only to the working rule set.

You use the following workflow with rules and rule sets.

- 1 Make changes to the working rule set.

- 2 Test the working rule set rules against a host to make sure that everything is working correctly.
- 3 Refine and retest the rules in the working rule set.
- 4 Activate the rules in the working rule set.

If you add a rule in a PowerCLI session and do not specify the `NoActivate` parameter, all rules that are currently in the working rule set are activated. You cannot activate individual rules.

See the PowerCLI command-line help and [Managing vSphere Auto Deploy with PowerCLI Cmdlets](#) for more information on using vSphere Auto Deploy with PowerCLI cmdlets. See [Managing vSphere Auto Deploy with the vSphere Web Client](#) for more information on using vSphere Auto Deploy with the vSphere Web Client.

vSphere Auto Deploy Boot Process

When you boot a host that you want to provision or reprovision with vSphere Auto Deploy, the vSphere Auto Deploy infrastructure supplies the image profile and, optionally, a host profile, a vCenter Server location, and script bundle for that host.

The boot process is different for hosts that have not yet been provisioned with vSphere Auto Deploy (first boot) and for hosts that have been provisioned with vSphere Auto Deploy and added to a vCenter Server system (subsequent boot).

First Boot Prerequisites

Before a first boot process, you must set up your system. Setup includes the following tasks, which are discussed in more detail in [Preparing for vSphere Auto Deploy](#).

- Set up a DHCP server that assigns an IP address to each host upon startup and that points the host to the TFTP server to download the iPXE boot loader from.
- If the hosts that you plan to provision with vSphere Auto Deploy are with legacy BIOS, verify that the vSphere Auto Deploy server has an IPv4 address. PXE booting with legacy BIOS firmware is possible only over IPv4. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.
- Identify an image profile to be used in one of the following ways.
 - Choose an ESXi image profile in a public depot.
 - (Optional) Create a custom image profile by using vSphere ESXi Image Builder, and place the image profile in a depot that the vSphere Auto Deploy server can access. The image profile must include a base ESXi VIB.
- (Optional) If you have a reference host in your environment, export the host profile of the reference host and define a rule that applies the host profile to one or more hosts. See [Setting Up a vSphere Auto Deploy Reference Host](#).
- Specify rules for the deployment of the host and add the rules to the active rule set.

First Boot Overview

When a host that has not yet been provisioned with vSphere Auto Deploy boots (first boot), the host interacts with several vSphere Auto Deploy components.

- 1 When the administrator turns on a host, the host starts a PXE boot sequence.

The DHCP Server assigns an IP address to the host and instructs the host to contact the TFTP server.

- 2 The host contacts the TFTP server and downloads the iPXE file (executable boot loader) and an iPXE configuration file.

- 3 iPXE starts executing.

The configuration file instructs the host to make a HTTP boot request to the vSphere Auto Deploy server. The HTTP request includes hardware and network information.

- 4 In response, the vSphere Auto Deploy server performs these tasks:

- a Queries the rules engine for information about the host.
- b Streams the components specified in the image profile, the optional host profile, and optional vCenter Server location information.

- 5 The host boots using the image profile.

If the vSphere Auto Deploy server provided a host profile, the host profile is applied to the host.

- 6 vSphere Auto Deploy adds the host to the vCenter Server system that vSphere Auto Deploy is registered with.

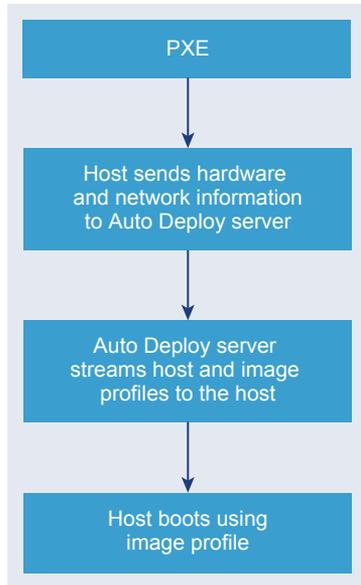
- a If a rule specifies a target folder or cluster on the vCenter Server system, the host is placed in that folder or cluster. The target folder must be under a data center.
- b If no rule exists that specifies a vCenter Server inventory location, vSphere Auto Deploy adds the host to the first datacenter displayed in the vSphere Web Client UI.

- 7 (Optional) If the host profile requires the user to specify certain information, such as a static IP address, the host is placed in maintenance mode when the host is added to the vCenter Server system.

You must reapply the host profile and update the host customization to have the host exit maintenance mode. When you update the host customization, answer any questions when prompted.

- 8 If the host is part of a DRS cluster, virtual machines from other hosts might be migrated to the host after the host has successfully been added to the vCenter Server system.

See [Provision a Host \(First Boot\)](#).

Figure 2-5. vSphere Auto Deploy Installation, First Boot

Subsequent Boots Without Updates

For hosts that are provisioned with vSphere Auto Deploy and managed by vCenter Server system, subsequent boots can become completely automatic.

- 1 The administrator reboots the host.
- 2 As the host boots up, vSphere Auto Deploy provisions the host with its image profile and host profile.
- 3 Virtual machines are brought up or migrated to the host based on the settings of the host.
 - Standalone host. Virtual machines are powered on according to autostart rules defined on the host.
 - DRS cluster host. Virtual machines that were successfully migrated to other hosts stay there. Virtual machines for which no host had enough resources are registered to the rebooted host.

If the vCenter Server system is unavailable, the host contacts the vSphere Auto Deploy server and is provisioned with an image profile. The host continues to contact the vSphere Auto Deploy server until vSphere Auto Deploy reconnects to the vCenter Server system.

vSphere Auto Deploy cannot set up vSphere distributed switches if vCenter Server is unavailable, and virtual machines are assigned to hosts only if they participate in an HA cluster. Until the host is reconnected to vCenter Server and the host profile is applied, the switch cannot be created. Because the host is in maintenance mode, virtual machines cannot start. See [Reprovision Hosts with Simple Reboot Operations](#).

Any hosts that are set up to require user input are placed in maintenance mode. See [Update the Host Customization in the vSphere Web Client](#).

Subsequent Boots With Updates

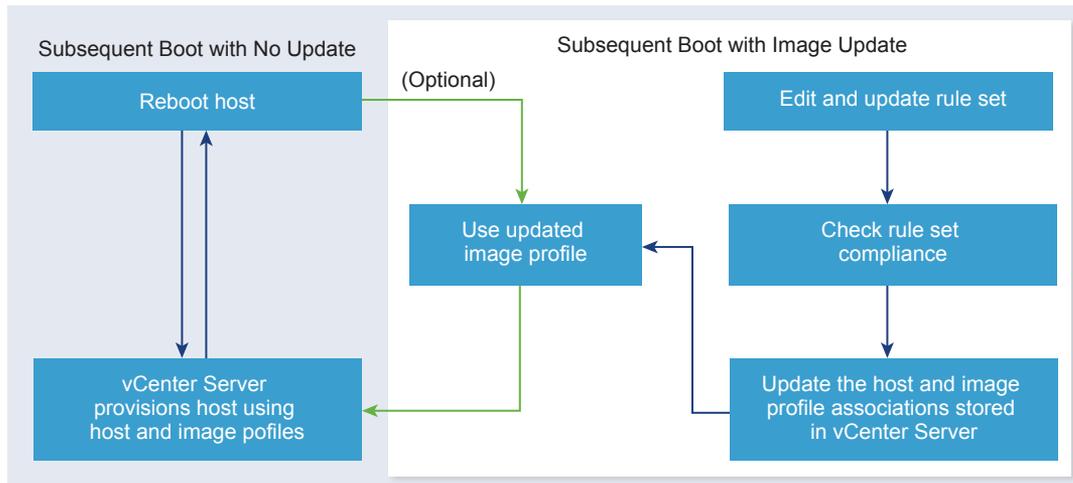
You can change the image profile, host profile, vCenter Server location, or script bundle for hosts. The process includes changing rules and testing and repairing the host's rule compliance.

- 1 The administrator uses the `Copy-DeployRule` PowerCLI cmdlet to copy and edit one or more rules and updates the rule set. See [Overview of the vSphere Auto Deploy Process by Using PowerCLI](#) for an example.
- 2 The administrator runs the `Test-DeployRulesetCompliance` cmdlet to check whether each host is using the information that the current rule set specifies.
- 3 The host returns a PowerCLI object that encapsulates compliance information.
- 4 The administrator runs the `Repair-DeployRulesetCompliance` cmdlet to update the image profile, host profile, or vCenter Server location the vCenter Server system stores for each host.
- 5 When the host reboots, it uses the updated image profile, host profile, vCenter Server location, or script bundle for the host.

If the host profile is set up to request user input, the host is placed in maintenance mode. Follow the steps in [Update the Host Customization in the vSphere Web Client](#).

See [Test and Repair Rule Compliance](#).

Figure 2-6. vSphere Auto Deploy Installation, Subsequent Boots



Provisioning of Systems that Have Distributed Switches

You can configure the host profile of a vSphere Auto Deploy reference host with a distributed switch.

When you configure the distributed switch, the boot configuration parameters policy is automatically set to match the network parameters required for host connectivity after a reboot.

When vSphere Auto Deploy provisions the ESXi host with the host profile, the host goes through a two-step process.

- 1 The host creates a standard virtual switch with the properties specified in the boot configuration parameters field.

- 2 The host creates the VMkernel NICs. The VMkernel NICs allow the host to connect to vSphere Auto Deploy and to the vCenter Server system.

When the host is added to vCenter Server, vCenter Server removes the standard switch and reapplies the distributed switch to the host.

Note Do not change the boot configuration parameters to avoid problems with your distributed switch.

Overview of the vSphere Auto Deploy Process by Using the vSphere Web Client

Getting started with vSphere Auto Deploy requires that you learn how vSphere Auto Deploy works, start the vSphere Auto Deploy and vSphere ESXi Image Builder vCenter Server services, create deploy rules that provision hosts, and power on your hosts to be booted with the image profile you specify.

The workflow for provisioning the hosts in your environment with vSphere Auto Deploy includes the following tasks:

- 1 Install vCenter Server and the vCenter Server components, or deploy the vCenter Server Appliance.
The vSphere Auto Deploy server is included with the management node.
- 2 Configure the vSphere Auto Deploy and vSphere ESXi Image Builder service startup types.
See [Prepare Your System for vSphere Auto Deploy](#).
- 3 Add or import a software depot to the vSphere Auto Deploy inventory.
See [Add a Software Depot](#) or [Import a Software Depot](#).
- 4 (Optional) If you want to create a custom image profile, clone or create an image profile by using the vSphere Web Client.
See [Clone an Image Profile](#) or [Create an Image Profile](#).
- 5 Create a deploy rule that assigns the image profile to one host, to multiple hosts specified by a pattern, or to all hosts..
See [Create a Deploy Rule](#).

Note vSphere Auto Deploy is optimized for provisioning hosts that have a fixed MAC address to IP address mapping in DHCP (sometimes called DHCP reservations). If you want to use static IP addresses, you must set up the host profile to prompt for host customization. For more information, see the *vSphere Host Profiles* documentation.

- 6 Power on the hosts that you want to provision.
- 7 Set up the host you provisioned as a reference host for your host profile.
You can specify the reference host syslog settings, firewall settings, storage, networking, and so on.
- 8 Extract a host profile from the reference host.
See the *Host Profiles* documentation.
- 9 To provision multiple hosts with the host profile, clone or edit the previously created rule by using the vSphere Web Client.

See [Clone a Deploy Rule](#) or [Editing a Deploy Rule](#).

10 Activate the new rule and deactivate the old one.

See [Activate, Deactivate, and Reorder Deploy Rules](#).

11 Remediate the host associations to apply the new rule to the host.

See [Remediate a Non-compliant Host](#).

12 Verify that the hosts you provisioned meet the following requirements.

- Each host is connected to the vCenter Server system.
- The hosts are not in maintenance mode.
- The hosts have no compliance failures.
- Each host with a host profile that requires user input has up-to-date host customization information.

Remediate host associations and compliance problems and reboot hosts until all hosts meet the requirements.

Read for an introduction to the boot process, differences between first and subsequent boots, and an overview of using host customization.

Overview of the vSphere Auto Deploy Process by Using PowerCLI

Getting started with vSphere Auto Deploy requires that you learn how vSphere Auto Deploy works, install the vSphere Auto Deploy server, install PowerCLI, write PowerCLI rules that provision hosts, and power on your hosts to be booted with the image profile you specify. You can customize of the image profile, host profile, and vCenter Server location.

See [Set Up vSphere Auto Deploy and Provision Hosts with vSphere PowerCLI](#) for a step-by-step exercise that helps you set up your first vSphere Auto Deploy environment on a Windows Server 2008 system.

To provision the hosts in your environment with vSphere Auto Deploy successfully, you can follow these steps.

- 1 Install vCenter Server and the vCenter Server components, or deploy the vCenter Server Appliance.
The vSphere Auto Deploy server is included with the management node.
- 2 Configure the vSphere Auto Deploy service startup type.
See [Prepare Your System for vSphere Auto Deploy](#).
- 3 Install PowerCLI, which includes vSphere Auto Deploy and vSphere ESXi Image Builder cmdlets.
See [Prepare Your System for vSphere Auto Deploy](#) and [Using vSphere Auto Deploy Cmdlets](#).
- 4 Find the image profile that includes the VIBs that you want to deploy to your hosts.
 - In most cases, you add the depots containing the required software to your PowerCLI session, and then select an image profile from one of those depots.

- To create a custom image profile, use vSphere ESXi Image Builder cmdlets to clone an existing image profile and add the custom VIBs to the clone. Add the custom image profile to the PowerCLI session.

You must use vSphere ESXi Image Builder for customization only if you have to add or remove VIBs. In most cases, you can add the depot where VMware hosts the image profiles to your PowerCLI session as a URL.

- 5 Start a PowerCLI session and connect to the vCenter Server system that vSphere Auto Deploy is registered with.
- 6 Use the `New-DeployRule` PowerCLI cmdlet to write a rule that assigns the image profile to one host, to multiple hosts specified by a pattern, or to all hosts.

```
New-DeployRule -Name "testrule" -Item image-profile -AllHosts
```

See [Assign an Image Profile to Hosts](#).

Note vSphere Auto Deploy is optimized for provisioning hosts that have a fixed MAC address to IP address mapping in DHCP (sometimes called DHCP reservations). If you want to use static IP addresses, you must set up the host profile to prompt for host customization. For more information, see the *vSphere Host Profiles* documentation.

- 7 Power on the hosts that you want to provision.
- 8 Set up the host you provisioned as a reference host for your host profile.
You can specify the reference host syslog settings, firewall settings, storage, networking, and so on.
- 9 Set up the host you provisioned as a reference host for your host profile.
You can specify the reference host syslog settings, firewall settings, storage, networking, and so on. See [Setting Up a vSphere Auto Deploy Reference Host](#).
- 10 Create and export a host profile for the reference host.
See the *Host Profiles* documentation.
- 11 To provision multiple hosts with the host profile, use the `Copy-DeployRule` cmdlet to edit the previously created rule.

You can revise the rule to assign not only an image profile but also a host profile, a vCenter Server location and a custom script bundle.

```
Copy-DeployRule -DeployRule "testrule" -ReplaceItem
  my_host_profile_from_reference_host,my_target_cluster
  -ReplacePattern "ipv4=192.XXX.1.10-192.XXX.1.20"
```

Where *my_host_profile_from_reference_host* is the name of the reference host profile, and *my_target_cluster* is the name of the target cluster.

- 12 Perform the test and repair compliance operations to remediate the hosts.

See [Test and Repair Rule Compliance](#).

13 Verify that the hosts you provisioned meet the following requirements.

- Each host is connected to the vCenter Server system.
- The hosts are not in maintenance mode.
- The hosts have no compliance failures.
- Each host with a host profile that requires user input has up-to-date host customization information.

Remediate host associations and compliance problems and reboot hosts until all hosts meet the requirements.

Read for an introduction to the boot process, differences between first and subsequent boots, and an overview of using host customization.

Preparing for vSphere Auto Deploy

Before you can start using vSphere Auto Deploy, you must prepare your environment. You start with server setup and hardware preparation. You must configure the vSphere Auto Deploy service startup type in the vCenter Server system that you plan to use for managing the hosts you provision, and install PowerCLI.

- [Prepare Your System for vSphere Auto Deploy](#)

Before you can PXE boot an ESXi host with vSphere Auto Deploy, you must install prerequisite software and set up the DHCP and TFTP servers that vSphere Auto Deploy interacts with.

- [Using vSphere Auto Deploy Cmdlets](#)

vSphere Auto Deploy cmdlets are implemented as Microsoft PowerShell cmdlets and included in PowerCLI. Users of vSphere Auto Deploy cmdlets can take advantage of all PowerCLI features.

- [Set Up Bulk Licensing](#)

You can use the vSphere Web Client or ESXi Shell to specify individual license keys, or you can set up bulk licensing by using PowerCLI cmdlets. Bulk licensing works for all ESXi hosts, but is especially useful for hosts provisioned with vSphere Auto Deploy.

Prepare Your System for vSphere Auto Deploy

Before you can PXE boot an ESXi host with vSphere Auto Deploy, you must install prerequisite software and set up the DHCP and TFTP servers that vSphere Auto Deploy interacts with.

Prerequisites

- Verify that the hosts that you plan to provision with vSphere Auto Deploy meet the hardware requirements for ESXi. See [ESXi Hardware Requirements](#).
- Verify that the ESXi hosts have network connectivity to vCenter Server and that all port requirements are met. See [Required Ports for vCenter Server and Platform Services Controller](#).

- If you want to use VLANs in your vSphere Auto Deploy environment, you must set up the end to end networking properly. When the host is PXE booting, the firmware driver must be set up to tag the frames with proper VLAN IDs. You must do this set up manually by making the correct changes in the UEFI/BIOS interface. You must also correctly configure the ESXi port groups with the correct VLAN IDs. Ask your network administrator how VLAN IDs are used in your environment.
- Verify that you have enough storage for the vSphere Auto Deploy repository. The vSphere Auto Deploy server uses the repository to store data it needs, including the rules and rule sets you create and the VIBs and image profiles that you specify in your rules.

Best practice is to allocate 2 GB to have enough room for four image profiles and some extra space. Each image profile requires approximately 350 MB. Determine how much space to reserve for the vSphere Auto Deploy repository by considering how many image profiles you expect to use.

- Obtain administrative privileges to the DHCP server that manages the network segment you want to boot from. You can use a DHCP server already in your environment, or install a DHCP server. For your vSphere Auto Deploy setup, replace the `gpxelinux.0` file name with `snponly64.efi.vmw-hardwired` for UEFI or `undionly.kpxe.vmw-hardwired` for BIOS. For more information on DHCP configurations, see [Sample DHCP Configurations](#).
- Secure your network as you would for any other PXE-based deployment method. vSphere Auto Deploy transfers data over SSL to prevent casual interference and snooping. However, the authenticity of the client or the vSphere Auto Deploy server is not checked during a PXE boot.
- If you want to manage vSphere Auto Deploy with PowerCLI cmdlets, verify that Microsoft .NET Framework 4.5 or 4.5.x and Windows PowerShell 3.0 or 4.0 are installed on a Windows machine. You can install PowerCLI on the Windows system on which vCenter Server is installed or on a different Windows system. See the *vSphere PowerCLI User's Guide*.
- Set up a remote Syslog server. See the *vCenter Server and Host Management* documentation for Syslog server configuration information. Configure the first host you boot to use the remote Syslog server and apply that host's host profile to all other target hosts. Optionally, install and use the vSphere Syslog Collector, a vCenter Server support tool that provides a unified architecture for system logging and enables network logging and combining of logs from multiple hosts.
- Install ESXi Dump Collector, set up your first host so that all core dumps are directed to ESXi Dump Collector, and apply the host profile from that host to all other hosts. See [Configure ESXi Dump Collector with ESXCLI](#).
- If the hosts that you plan to provision with vSphere Auto Deploy are with legacy BIOS, verify that the vSphere Auto Deploy server has an IPv4 address. PXE booting with legacy BIOS firmware is possible only over IPv4. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.

Procedure

- 1 Install vCenter Server or deploy the vCenter Server Appliance.

The vSphere Auto Deploy server is included with the management node.

- 2 Configure the vSphere Auto Deploy service startup type.
 - a Log in to your vCenter Server system by using the vSphere Web Client.
 - b On the vSphere Web Client Home page, click **Administration**.
 - c Under **System Configuration** click **Services**.
 - d Select **Auto Deploy**, click the **Actions** menu, and select **Edit Startup Type**.
 - On Windows, the vSphere Auto Deploy service is disabled. In the **Edit Startup Type** window, select **Manual** or **Automatic** to enable vSphere Auto Deploy.
 - On the vCenter Server Appliance, the vSphere Auto Deploy service by default is set to **Manual**. If you want the vSphere Auto Deploy service to start automatically upon OS startup, select **Automatic**.
- 3 (Optional) If you want to manage vSphere Auto Deploy with the vSphere Web Client, configure the vSphere ESXi Image Builder service startup type.
 - a Repeat [Step 2a](#) through [Step 2c](#).
 - b Select **ImageBuilder Service**, click the **Actions** menu, and select **Edit Startup Type**.
 - On Windows, the vSphere ESXi Image Builder service is disabled. In the **Edit Startup Type** window, select **Manual** or **Automatic** to enable the service.
 - On the vCenter Server Appliance, the vSphere Auto Deploy service by default is set to **Manual**. If you want the vSphere ESXi Image Builder service to start automatically upon OS startup, select **Automatic**.
 - c Log out of the vSphere Web Client and log in again.

The **Auto Deploy** icon is visible on the Home page of the vSphere Web Client.
- 4 (Optional) If you want to manage vSphere Auto Deploy with PowerCLI cmdlets, install PowerCLI.
 - a Download the latest version of PowerCLI from the VMware Web site.
 - b Navigate to the folder that contains the PowerCLI file you downloaded and double-click the executable file.

If the installation wizard detects an earlier version of PowerCLI on your system, it will attempt to upgrade your existing installation
 - c Follow the prompts in the wizard to complete the installation.
- 5 Configure the TFTP server.
 - a In a vSphere Web Client connected to the vCenter Server system, go to the inventory list and select the vCenter Server system.
 - b Click the **Manage** tab, select **Settings**, and click **Auto Deploy**.
 - c Click **Download TFTP Boot Zip** to download the TFTP configuration file and unzip the file to the directory in which your TFTP server stores files.

- 6 Set up your DHCP server to point to the TFTP server on which the TFTP ZIP file is located.
 - a Specify the TFTP Server's IP address in DHCP option 66, frequently called next-server.
 - b Specify the boot file name, which is `snponly64.efi.vmw-hardwired` for UEFI or `undionly.kpxe.vmw-hardwired` for BIOS in the DHCP option 67, frequently called boot-filename.
- 7 Set each host you want to provision with vSphere Auto Deploy to network boot or PXE boot, following the manufacturer's instructions.
- 8 (Optional) If you set up your environment to use Thumbprint mode, you can use your own Certificate Authority (CA) by replacing the OpenSSL certificate `rbd-ca.crt` and the OpenSSL private key `rbd-ca.key` with your own certificate and key file.
 - On Windows, the files are in the SSL subfolder of the vSphere Auto Deploy installation directory. For example, on Windows 7 the default is `C:\ProgramData\VMware\VMware vSphere Auto Deploy\ssl`.
 - On the vCenter Server Appliance, the files are in `/etc/vmware-rbd/ssl/`.
 By default, vCenter Server 6.0 and later uses VMware Certificate Authority (VMCA).

When you start a host that is set up for vSphere Auto Deploy, the host contacts the DHCP server and is directed to the vSphere Auto Deploy server, which provisions the host with the image profile specified in the active rule set.

What to do next

- Define a rule that assigns an image profile and optional host profile, host location or script bundle to the host. For Managing vSphere Auto Deploy with PowerCLI cmdlets, see the [Managing vSphere Auto Deploy with PowerCLI Cmdlets](#) section. For managing vSphere Auto Deploy with the vSphere Web Client, see the [Managing vSphere Auto Deploy with the vSphere Web Client](#) section.
- (Optional) Configure the first host that you provision as a reference host. Use the storage, networking, and other settings you want for your target hosts to share. Create a host profile for the reference host and write a rule that assigns both the already tested image profile and the host profile to target hosts.
- (Optional) If you want to have vSphere Auto Deploy overwrite existing partitions, set up a reference host to do auto partitioning and apply the host profile of the reference host to other hosts. See [Configure a Reference Host for Auto-Partitioning](#).
- (Optional) If you have to configure host-specific information, set up the host profile of the reference host to prompt for user input. For more information about host customizations, see the *vSphere Host Profiles* documentation.

Using vSphere Auto Deploy Cmdlets

vSphere Auto Deploy cmdlets are implemented as Microsoft PowerShell cmdlets and included in PowerCLI. Users of vSphere Auto Deploy cmdlets can take advantage of all PowerCLI features.

Experienced PowerShell users can use vSphere Auto Deploy cmdlets just like other PowerShell cmdlets. If you are new to PowerShell and PowerCLI, the following tips might be helpful.

You can type cmdlets, parameters, and parameter values in the PowerCLI shell.

- Get help for any cmdlet by running `Get-Helpcmdlet_name`.
- Remember that PowerShell is not case sensitive.
- Use tab completion for cmdlet names and parameter names.
- Format any variable and cmdlet output by using `Format-List` or `Format-Table`, or their short forms `fl` or `ft`. For more information, run the `Get-Help Format-List` cmdlet.

Passing Parameters by Name

You can pass in parameters by name in most cases and surround parameter values that contain spaces or special characters with double quotes.

```
Copy-DeployRule -DeployRule testrule -ReplaceItem MyNewProfile
```

Most examples in the *vSphere Installation and Setup* documentation pass in parameters by name.

Passing Parameters as Objects

You can pass parameters as objects if you want to perform scripting and automation. Passing in parameters as objects is useful with cmdlets that return multiple objects and with cmdlets that return a single object. Consider the following example.

- 1 Bind the object that encapsulates rule set compliance information for a host to a variable.

```
$tr = Test-DeployRuleSetCompliance MyEsxi42
```

- 2 View the `itemlist` property of the object to see the difference between what is in the rule set and what the host is currently using.

```
$tr.itemlist
```

- 3 Remediate the host to use the revised rule set by using the `Repair-DeployRuleSetCompliance` cmdlet with the variable.

```
Repair-DeployRuleSetCompliance $tr
```

The example remediates the host the next time you boot the host.

Set Up Bulk Licensing

You can use the vSphere Web Client or ESXi Shell to specify individual license keys, or you can set up bulk licensing by using PowerCLI cmdlets. Bulk licensing works for all ESXi hosts, but is especially useful for hosts provisioned with vSphere Auto Deploy.

Assigning license keys through the vSphere Web Client and assigning licensing by using PowerCLI cmdlets function differently.

Assign license keys with the vSphere Web Client

You can assign license keys to a host when you add the host to the vCenter Server system or when the host is managed by a vCenter Server system.

Assign license keys with LicenseDataManager PowerCLI

You can specify a set of license keys to be added to a set of hosts. The license keys are added to the vCenter Server database. Each time a host is added to the vCenter Server system or reconnects to it, the host is assigned a license key. A license key that is assigned through PowerCLI is treated as a default license key. When an unlicensed host is added or reconnected, it is assigned the default license key. If a host is already licensed, it keeps its license key.

The following example assigns licenses to all hosts in a data center. You can also associate licenses with hosts and clusters.

The following example is for advanced PowerCLI users who know how to use PowerShell variables.

Prerequisites

[Prepare Your System for vSphere Auto Deploy.](#)

Procedure

- 1 In a PowerCLI session, connect to the vCenter Server system you want to use and bind the associated license manager to a variable.

```
Connect-VIServer -Server 192.XXX.X.XX -User username -Password password  
$licenseDataManager = Get-LicenseDataManager
```

- 2 Run a cmdlet that retrieves the datacenter in which the hosts for which you want to use the bulk licensing feature are located.

```
$hostContainer = Get-Datacenter -Name Datacenter-X
```

You can also run a cmdlet that retrieves a cluster to use bulk licensing for all hosts in a cluster, or retrieves a folder to use bulk licensing for all hosts in a folder.

- 3 Create a new `LicenseData` object and a `LicenseKeyEntry` object with associated type ID and license key.

```
$licenseData = New-Object VMware.VimAutomation.License.Types.LicenseData
$licenseKeyEntry = New-Object VMware.VimAutomation.License.Types.LicenseKeyEntry
$licenseKeyEntry.TypeId = "vmware-vmware"
$licenseKeyEntry.LicenseKey = "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX"
```

- 4 Associate the `LicenseKeys` attribute of the `LicenseData` object you created in step 3 with the `LicenseKeyEntry` object.

```
$licenseData.LicenseKeys += $licenseKeyEntry
```

- 5 Update the license data for the data center with the `LicenseData` object and verify that the license is associated with the host container.

```
$licenseDataManager.UpdateAssociatedLicenseData($hostContainer.Uid, $licenseData)
$licenseDataManager.QueryAssociatedLicenseData($hostContainer.Uid)
```

- 6 Provision one or more hosts with vSphere Auto Deploy and assign them to the data center or to the cluster that you assigned the license data to.
- 7 You can use the vSphere Web Client to verify that the host is successfully assigned to the default license XXXXX-XXXXX-XXXXX-XXXXX-XXXXX.

All hosts that you assigned to the data center are now licensed automatically.

Managing vSphere Auto Deploy with PowerCLI Cmdlets

You can manage vSphere Auto Deploy with PowerCLI cmdlets to create rules that associate hosts with image profiles, host profiles, custom scripts and locations on the vCenter Server target. You can also update hosts by testing rule compliance and repairing compliance issues.

vSphere Auto Deploy PowerCLI Cmdlet Overview

You specify the rules that assign image profiles and host profiles to hosts using a set of PowerCLI cmdlets that are included in PowerCLI.

If you are new to PowerCLI, read the PowerCLI documentation and review [Using vSphere Auto Deploy Cmdlets](#). You can get help for any command at the PowerShell prompt.

- Basic help: `Get-Help cmdlet_name`
- Detailed help: `Get-Help cmdlet_name -Detailed`

Note When you run vSphere Auto Deploy cmdlets, provide all parameters on the command line when you invoke the cmdlet. Supplying parameters in interactive mode is not recommended.

Table 2-16. Rule Engine PowerCLI Cmdlets

Command	Description
Get-DeployCommand	Returns a list of vSphere Auto Deploy cmdlets.
New-DeployRule	Creates a new rule with the specified items and patterns.
Set-DeployRule	Updates an existing rule with the specified items and patterns. You cannot update a rule that is part of a rule set.
Get-DeployRule	Retrieves the rules with the specified names.
Copy-DeployRule	Clones and updates an existing rule.
Add-DeployRule	Adds one or more rules to the working rule set and, by default, also to the active rule set. Use the NoActivate parameter to add a rule only to the working rule set.
Remove-DeployRule	Removes one or more rules from the working rule set and from the active rule set. Run this command with the -Delete parameter to completely delete the rule.
Set-DeployRuleset	Explicitly sets the list of rules in the working rule set.
Get-DeployRuleset	Retrieves the current working rule set or the current active rule set.
Switch-ActiveDeployRuleset	Activates a rule set so that any new requests are evaluated through the rule set.
Get-VMHostMatchingRules	Retrieves rules matching a pattern. For example, you can retrieve all rules that apply to a host or hosts. Use this cmdlet primarily for debugging.
Test-DeployRulesetCompliance	Checks whether the items associated with a specified host are in compliance with the active rule set.
Repair-DeployRulesetCompliance	Given the output of Test-DeployRulesetCompliance, this cmdlet updates the image profile, host profile, and location for each host in the vCenter Server inventory. The cmdlet might apply image profiles, apply host profiles, or move hosts to prespecified folders or clusters on the vCenter Server system.
Apply-EsxImageProfile	Associates the specified image profile with the specified host.
Get-VMHostImageProfile	Retrieves the image profile in use by a specified host. This cmdlet differs from the Get-EsxImageProfile cmdlet in vSphere ESXi Image Builder.
Repair-DeployImageCache	Use this cmdlet only if the vSphere Auto Deploy image cache is accidentally deleted.
Get-VMHostAttributes	Retrieves the attributes for a host that are used when the vSphere Auto Deploy server evaluates the rules.
Get-DeployMachineIdentity	Returns a string value that vSphere Auto Deploy uses to logically link an ESXi host in vCenter Server to a physical machine.
Set-DeployMachineIdentity	Logically links a host object in the vCenter Server database to a physical machine. Use this cmdlet to add hosts without specifying rules.

Table 2-16. Rule Engine PowerCLI Cmdlets (Continued)

Command	Description
Get-DeployOption	Retrieves the vSphere Auto Deploy global configuration options. This cmdlet currently supports the <code>vlan-id</code> option, which specifies the default VLAN ID for the ESXi Management Network of a host provisioned with vSphere Auto Deploy. vSphere Auto Deploy uses the value only if the host boots without a host profile.
Set-DeployOption	Sets the value of a global configuration option. Currently supports the <code>vlan-id</code> option for setting the default VLAN ID for the ESXi Management Network.
Add-ProxyServer	Adds a proxy server to the vSphere Auto Deploy database. Run the command with the <code>-Address</code> parameter to specify the IPv4 or IPv6 address. The address can include a port number.
List-ProxyServer	Lists the proxy servers that are currently registered with vSphere Auto Deploy.
Delete-ProxyServer	Deletes one or more proxy servers from the list of proxy servers that are registered with vSphere Auto Deploy. You can run the command with the <code>-id</code> parameter from the list of proxy servers or with the <code>-Address</code> parameter by specifying the IPv4 or IPv6 address of the proxy server you want to delete.
Add-ScriptBundle	Adds one or more script bundles to the vSphere Auto Deploy server.
Get-ScriptBundle	Retrieves the list of script bundles available on the vSphere Auto Deploy server and the scripts they contain.

Assign an Image Profile to Hosts

Before you can provision a host, you must create rules that assign an image profile to each host that you want to provision by using vSphere Auto Deploy.

vSphere Auto Deploy extensibility rules enforce that VIBs at the CommunitySupported level can only contain files from certain predefined locations, such as the ESXCLI plug-in path, jumpstart plug-in path, and so on. If you add a VIB that is in a different location to an image profile, a warning results. You can override the warning by using the `force` option.

If you call the `New-DeployRule` cmdlet on an image profile that includes VIBs at the CommunitySupported level which violate the rule, set `$DeployNoSignatureCheck = $true` before adding the image profile. With that setting, the system ignores signature validation and does not perform the extensibility rules check.

Note Image profiles that include VIBs at the CommunitySupported level are not supported on production systems.

Prerequisites

- Prepare your system for vSphere Auto Deploy. For more information, see *vSphere Installation and Setup*.

Procedure

- 1 In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

- 2 Determine the location of a public software depot, or define a custom image profile by using vSphere ESXi Image Builder.
- 3 Run `Add-EsxSoftwareDepot` to add the software depot that contains the image profile to the PowerCLI session.

Depot Type	Cmdlet
Remote depot	Run <code>Add-EsxSoftwareDepot <i>depot_url</i></code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file path. b Run <code>Add-EsxSoftwareDepot C:\<i>file_path</i>\my_offline_depot.zip</code>.

- 4 In the depot, find the image profile that you want to use by running the `Get-EsxImageProfile` cmdlet.

By default, the ESXi depot includes one base image profile that includes VMware tools and has the string `standard` in its name, and one base image profile that does not include VMware tools.

- 5 Define a rule in which hosts with certain attributes, for example a range of IP addresses, are assigned to the image profile.

```
New-DeployRule -Name "testrule" -Item "My Profile25" -Pattern "vendor=Acme,Zven", "ipv4=192.XXX.1.10-192.XXX.1.20"
```

Double quotes are required if a name contains spaces, optional otherwise. Specify `-AllHosts` instead of a pattern to apply the item to all hosts.

The cmdlet creates a rule named `testrule`. The rule assigns the image profile named `My Profile25` to all hosts with a vendor of `Acme` or `Zven` that also have an IP address in the specified range.

- 6 Add the rule to the rule set.

```
Add-DeployRule testrule
```

By default, the rule is added to both the working rule set and the active rule set. If you use the `NoActivate` parameter, the working rule set does not become the active rule set.

When the host boots from iPXE, it reports attributes of the machine to the console. Use the same format of the attributes when writing deploy rules.

```
*****
* Booting through VMware AutoDeploy...
*
* Machine attributes:
* . asset=No Asset Tag
* . domain=vmware.com
* . hostname=myhost.mycompany.com
* . ipv4=XX.XX.XXX.XXX
* . mac=XX:XX:XX:XX:XX:XX
* . model=MyVendorModel
* . oemstring=Product ID: XXXXXX-XXX
* . serial=XX XX XX XX XX XX...
* . uuid=XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXX
* . vendor=MyVendor
*****
```

What to do next

- For hosts already provisioned with vSphere Auto Deploy, perform the compliance testing and repair operations to provision them with the new image profile. See [Test and Repair Rule Compliance](#).
- Turn on unprovisioned hosts to provision them with the new image profile.

Write a Rule and Assign a Host Profile to Hosts

vSphere Auto Deploy can assign a host profile to one or more hosts. The host profile might include information about storage configuration, network configuration, or other characteristics of the host. If you add a host to a cluster, that cluster's host profile is used.

In many cases, you assign a host to a cluster instead of specifying a host profile explicitly. The host uses the host profile of the cluster.

Prerequisites

- Prepare your system for vSphere Auto Deploy. For more information, see *vSphere Installation and Setup*.
- Export the host profile that you want to use.

Procedure

- 1 In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

- 2 Using the vSphere Web Client, set up a host with the settings you want to use and create a host profile from that host.
- 3 Find the name of the host profile by running `Get-VMhostProfile` PowerCLI cmdlet, passing in the ESXi host from which you create a host profile.
- 4 At the PowerCLI prompt, define a rule in which host profiles are assigned to hosts with certain attributes, for example a range of IP addresses.

```
New-DeployRule -Name "testrule2" -Item my_host_profile -Pattern "vendor=Acme,Zven", "ipv4=192.XXX.1.10-192.XXX.1.20"
```

The specified item is assigned to all hosts with the specified attributes. This example specifies a rule named `testrule2`. The rule assigns the specified host profile `my_host_profile` to all hosts with an IP address inside the specified range and with a manufacturer of Acme or Zven.

- 5 Add the rule to the rule set.

```
Add-DeployRule testrule2
```

By default, the working rule set becomes the active rule set, and any changes to the rule set become active when you add a rule. If you use the `NoActivate` parameter, the working rule set does not become the active rule set.

What to do next

- Assign a host already provisioned with vSphere Auto Deploy to the new host profile by performing compliance test and repair operations on those hosts. For more information, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the host profile.

Write a Rule and Assign a Host to a Folder or Cluster

vSphere Auto Deploy can assign a host to a folder or cluster. When the host boots, vSphere Auto Deploy adds it to the specified location on the vCenter Server. Hosts assigned to a cluster inherit the cluster's host profile.

Prerequisites

- [Prepare Your System for vSphere Auto Deploy](#)
- Verify that the folder you select is in a data center or in a cluster. You cannot assign the host to a standalone top-level folder.

Procedure

- 1 In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

- 2 Define a rule in which hosts with certain attributes, for example a range of IP addresses, are assigned to a folder or a cluster.

```
New-DeployRule -Name testrule3 -Item "my folder" -Pattern "vendor=Acme,Zven", "ipv4=192.XXX.1.10-192.XXX.1.20"
```

This example passes in the folder by name. You can instead pass in a folder, cluster, or data center object that you retrieve with the `Get-Folder`, `Get-Cluster`, or `Get-Datacenter` cmdlet.

- 3 Add the rule to the rule set.

```
Add-DeployRule testrule3
```

By default, the working rule set becomes the active rule set, and any changes to the rule set become active when you add a rule. If you use the `NoActivate` parameter, the working rule set does not become the active rule set.

What to do next

- Assign a host already provisioned with vSphere Auto Deploy to the new folder or cluster location by performing test and repair compliance operation. See [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to add them to the specified vCenter Server location.

Configure a Stateless System by Running a Custom Script

You can use vSphere Auto Deploy to configure one or more hosts by associating custom scripts with a vSphere Auto Deploy rule.

The scripts run in alphabetical order after the initial ESXi boot workflow of the host.

Prerequisites

- Prepare your system for vSphere Auto Deploy. For more information, see *vSphere Installation and Setup*.
- Verify that the script bundle you want to associate with a vSphere Auto Deploy rule is in `.tgz` format, with a maximum size of 10 MB, and written in Python or BusyBox ash scripting language.

Procedure

- 1 In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

- 2 Run the `Add-ScriptBundle` cmdlet to add the script bundle that contains the necessary scripts to the vSphere Auto Deploy inventory.

```
Add-ScriptBundle c:/temp/MyScriptBundle.tgz
```

The name of the script bundle without the `.tgz` extension is the name identifier or object of the script bundle item. You can update an existing script bundle by using the `-Update` parameter with the `Add-ScriptBundle` cmdlet.

- 3 (Optional) Run the `Get-ScriptBundle` cmdlet to verify that the script bundle is added to the vSphere Auto Deploy inventory.
- 4 Define a rule in which hosts with certain attributes, for example a range of IP addresses, are assigned to the script bundle.

```
New-DeployRule -Name "testrule4" -Item "MyScriptBundle" -Pattern "vendor=Acme,Zven", "ipv4=192.XXX.1.10-192.XXX.1.20"
```

Double quotes are required if a name contains spaces, optional otherwise. Specify `-AllHosts` instead of a pattern to apply the item to all hosts.

You create a rule named `testrule4`. The rule assigns the script bundle named My Script Bundle to all hosts with a vendor of Acme or Zven that also have an IP address in the specified range. You can use the name identifier of the script bundle or the object returned by the `Get-ScriptBundle` cmdlet to identify the script bundle you want to associate with the rule.

- 5 Add the rule to the rule set.

```
Add-DeployRule testrule4
```

By default, the rule is added to both the working rule set and the active rule set. If you use the `NoActivate` parameter, the working rule set does not become the active rule set.

What to do next

- For hosts already provisioned with vSphere Auto Deploy, perform the compliance testing and repair operations to provision them with the new scripts. See [Test and Repair Rule Compliance](#).
- Turn on unprovisioned hosts to provision them with the new scripts.

Test and Repair Rule Compliance

When you add a rule to the vSphere Auto Deploy rule set or make changes to one or more rules, hosts are not updated automatically. vSphere Auto Deploy applies the new rules only when you test their rule compliance and perform remediation.

Prerequisites

- Prepare your system for vSphere Auto Deploy. For more information, see *vSphere Installation and Setup*.
- Verify that your infrastructure includes one or more ESXi hosts provisioned with vSphere Auto Deploy, and that the host on which you installed PowerCLI can access those ESXi hosts.

Procedure

- 1 In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

- 2 Use PowerCLI to check which vSphere Auto Deploy rules are currently available.

```
Get-DeployRule
```

The system returns the rules and the associated items and patterns.

- 3 Make a change to one of the available rules.

For example, you can change the image profile and the name of the rule.

```
Copy-DeployRule -DeployRule testrule -ReplaceItem MyNewProfile
```

You cannot edit a rule already added to the active rule set. Instead, you can copy the rule and replace the item or pattern you want to change.

- 4 Verify that you can access the host for which you want to test rule set compliance.

```
Get-VMHost -Name MyEsxi42
```

- 5 Run the cmdlet that tests rule set compliance for the host, and bind the return value to a variable for later use.

```
$tr = Test-DeployRuleSetCompliance MyEsxi42
```

- 6 Examine the differences between the contents of the rule set and configuration of the host.

```
$str.itemlist
```

The system returns a table of current and expected items if the host for which you want to test the new rule set compliance is compliant with the active rule set.

CurrentItem	ExpectedItem
-----	-----
<i>My Profile 25</i>	<i>MyNewProfile</i>

- 7 Remediate the host to use the revised rule set the next time you boot the host.

```
Repair-DeployRuleSetCompliance $str
```

What to do next

If the rule you changed specified the inventory location, the change takes effect when you repair compliance. For all other changes, reboot your host to have vSphere Auto Deploy apply the new rule and to achieve compliance between the rule set and the host.

Register a Caching Proxy Server Address with vSphere Auto Deploy

Simultaneously booting large number of stateless hosts places a significant load on the vSphere Auto Deploy server. You can load balance the requests between the vSphere Auto Deploy server and one or more proxy servers that you register with vSphere Auto Deploy.

Prerequisites

- Prepare your system for vSphere Auto Deploy. For more information, see *vSphere Installation and Setup*.

Procedure

- 1 In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

- 2 Register a caching proxy server addresses with vSphere Auto Deploy by running the `Add-ProxyServer` cmdlet.

```
Add-ProxyServer -Address 'https://proxy_server_ip_address:port_number'
```

You can run the cmdlet multiple times to register multiple proxy servers. The address can contain a port number.

- 3 (Optional) Run the `List-ProxyServer` cmdlet to verify that the caching proxy server is registered with vSphere Auto Deploy.

Managing vSphere Auto Deploy with the vSphere Web Client

You can add ESXi hosts to the vSphere Auto Deploy inventory and create, monitor, and manage the vSphere Auto Deploy rules and ESXi host associations by using the vSphere Web Client.

Create a Deploy Rule

Before you provision ESXi hosts with vSphere Auto Deploy, you must create rules that assign image profiles, host profiles, and host locations to the hosts. An ESXi host can match more than one vSphere Auto Deploy rule criteria, when this is the case, the rule order is considered.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- If you want to include an image profile to the rule, verify that the software depot you need is added to the inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).

Procedure

1 Start the New Deploy Rule Wizard

You can create a new vSphere Auto Deploy rule by using the **New Deploy Rule** wizard.

2 Name the Rule and Define Matching Criteria in the New Deploy Rule Wizard

When you start the **New Deploy Rule** wizard, you must first enter a rule name and select a pattern to apply the rule to some or all hosts in the inventory.

3 Select an Image Profile in the New Deploy Rule Wizard

In the **New Deploy Rule** wizard, you can optionally assign an image profile to the hosts that match the rule criteria.

4 Select a Host Profile in the New Deploy Rule Wizard

In the **New Deploy Rule** wizard, you can optionally assign a host profile to the hosts that match the rule criteria.

5 Select Host Location in the New Deploy Rule Wizard

In the **New Deploy Rule** wizard, you can optionally add the hosts that match the criteria of the rule to a specific location.

6 View the Summary of the New Deploy Rule Wizard

In the **New Deploy Rule** wizard, you can review the settings of the new vSphere Auto Deploy rule before completing the wizard.

What to do next

- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).

- Edit a vSphere Auto Deploy rule. See [Editing a Deploy Rule](#).
- View the image profile, host profile, and location associations of a host. See [View Host Associations](#).
- Remediate non-compliant hosts. See [Remediate a Non-compliant Host](#).
- Change the image profile association of a host. See [Edit the Image Profile Association of a Host](#).

Start the New Deploy Rule Wizard

You can create a new vSphere Auto Deploy rule by using the **New Deploy Rule** wizard.

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.
By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.
- 2 On the **Deploy Rules** tab, click **New Deploy Rule**.

The **New Deploy Rule** wizard appears.

Name the Rule and Define Matching Criteria in the New Deploy Rule Wizard

When you start the **New Deploy Rule** wizard, you must first enter a rule name and select a pattern to apply the rule to some or all hosts in the inventory.

Procedure

- 1 On the Name and hosts page of the wizard, enter a name for the new rule.
- 2 Select a pattern to apply the rule to the hosts in the inventory.
You can select to apply the rule to all the hosts in the inventory or to apply the rule only to hosts that match a specific pattern. You can select one or more patterns.
For example, the rule can apply only to hosts in a vCenter Single Sign-On domain, with a specific host name, or that match a specific IPv4 range.
- 3 Click **Next**.

Select an Image Profile in the New Deploy Rule Wizard

In the **New Deploy Rule** wizard, you can optionally assign an image profile to the hosts that match the rule criteria.

Prerequisites

If you want to include an image profile to the rule, verify that the software depot you need is added to the inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).

Procedure

- 1 On the Select image profile page of the wizard, select an image profile.

Option	Action
If you do not want to assign an image profile to the selected hosts	Select the No image profile check box.
If you want to assign an image profile to the selected hosts	<ol style="list-style-type: none"> 1 Select a software depot from the drop-down menu. 2 Select an image profile from the list. 3 (Optional) If you want to bypass the acceptance level verification for the image profile, select the Skip image profile signature check check box.

- 2 Click **Next**.

Select a Host Profile in the New Deploy Rule Wizard

In the **New Deploy Rule** wizard, you can optionally assign a host profile to the hosts that match the rule criteria.

Procedure

- 1 On the Select host profile page of the wizard, select a host profile.

Option	Action
If you do not want to assign a host profile to the selected hosts	Select the Do not include a host profile check box.
If you want to assign a host profile to the selected hosts	Select a host profile from the list.

- 2 Click **Next**.

Select Host Location in the New Deploy Rule Wizard

In the **New Deploy Rule** wizard, you can optionally add the hosts that match the criteria of the rule to a specific location.

Procedure

- 1 On the Select host location page of the wizard, select a location for the hosts that match the rule.

Option	Action
If you do not want to select a host location	Select the Do not include a location check box.
If you want to select a specific location for the selected hosts	Select a data center, folder, or cluster as host location.

- 2 Click **Next**.

View the Summary of the New Deploy Rule Wizard

In the **New Deploy Rule** wizard,, you can review the settings of the new vSphere Auto Deploy rule before completing the wizard.

Procedure

- 1 On the Ready to complete page, review the summary information for the new rule.
- 2 Click **Finish**.

You can view the newly created rule listed on the **Deploy Rules** tab.

Clone a Deploy Rule

You can use a vSphere Auto Deploy rule as a template and modify only parts of the rule instead of creating a new one.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).
- If you want to include an image profile to the rule, verify that the software depot you need is added to the inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).

Procedure

1 [Start the Clone Deploy Rule Wizard](#)

You can clone an existing vSphere Auto Deploy rule by using the **Clone Deploy Rule** wizard.

2 [Name the Rule and Define Matching Criteria in the Clone Deploy Rule Wizard](#)

When you start the **Clone Deploy Rule** wizard to clone a vSphere Auto Deploy rule, you must first choose whether to keep the default name of the cloned rule and whether to change the matching criteria of the rule.

3 [Select an Image Profile in the Clone Deploy Rule Wizard](#)

In the **Clone Deploy Rule** wizard, you can optionally assign an image profile to the hosts that match the rule criteria or keep the same image profile that the cloned rule uses.

4 [Select a Host Profile in the Clone Deploy Rule Wizard](#)

In the **Clone Deploy Rule** wizard, you can optionally assign a host profile to the hosts that match the rule criteria or keep the same host profile used in the cloned rule.

5 [Select Host Location in the Clone Deploy Rule](#)

In the **Clone Deploy Rule** wizard, you can optionally add the hosts that match the criteria of the rule to a specific location or keep the location that the cloned rule uses.

6 [View the Summary of the Clone Deploy Rule Wizard](#)

You can review the settings of the cloned vSphere Auto Deploy rule before completing the wizard.

What to do next

- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).
- Edit a vSphere Auto Deploy rule. See [Editing a Deploy Rule](#).

Start the Clone Deploy Rule Wizard

You can clone an existing vSphere Auto Deploy rule by using the **Clone Deploy Rule** wizard.

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.
By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.
- 2 On the **Deploy Rules** tab, select a rule from the list.
- 3 Click the **Clone** icon.

The **Clone Deploy Rule** wizard appears.

Name the Rule and Define Matching Criteria in the Clone Deploy Rule Wizard

When you start the **Clone Deploy Rule** wizard to clone a vSphere Auto Deploy rule, you must first choose whether to keep the default name of the cloned rule and whether to change the matching criteria of the rule.

Procedure

- 1 On the Name and hosts page of the wizard, enter a name for the new rule.
- 2 Select a pattern to apply the rule to the hosts in the inventory.
You can select to apply the rule to all the hosts in the inventory or to apply the rule only to hosts that match a specific pattern. You can select one or more patterns.
For example, the rule can apply only to hosts in a vCenter Single Sign-On domain, with a specific host name, or that match a specific IPv4 range.
- 3 Click **Next**.

Select an Image Profile in the Clone Deploy Rule Wizard

In the **Clone Deploy Rule** wizard, you can optionally assign an image profile to the hosts that match the rule criteria or keep the same image profile that the cloned rule uses.

Prerequisites

If you want to include an image profile to the rule, verify that the software depot you need is added to the inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).

Procedure

- 1 On the Select image profile page of the wizard, select an image profile.

Option	Action
If you do not want to change the image profile	Select the Same image profile option.
If you do not want to assign an image profile to the selected hosts	Select the No image profile option.
If you want to assign a new image profile to the selected hosts	<ol style="list-style-type: none"> 1 Select the Browse for Image Profile option. 2 Select a software depot from the drop-down menu. 3 Select an image profile from the list. 4 (Optional) If you want to bypass the acceptance level verification for the image profile, select the Skip image profile signature check check box.

- 2 Click **Next**.

Select a Host Profile in the Clone Deploy Rule Wizard

In the **Clone Deploy Rule** wizard, you can optionally assign a host profile to the hosts that match the rule criteria or keep the same host profile used in the cloned rule.

Procedure

- ◆ On the Select host profile page of the wizard, select a host profile.

Option	Action
If you want to keep the host profile used in the cloned rule	Click Next .
If you do not want to assign a host profile to the selected hosts	Select the Do not include a host profile check box and click Next .
If you want to assign a new host profile to the selected hosts	Select a host profile from the list and click Next .

Select Host Location in the Clone Deploy Rule

In the **Clone Deploy Rule** wizard, you can optionally add the hosts that match the criteria of the rule to a specific location or keep the location that the cloned rule uses.

Procedure

- 1 On the Select host location page of the wizard, select a location for the hosts that match the rule.

Option	Action
If you want to keep the host location used in the cloned rule	Leave the default location.
If you do not want the rule to include a host location	Select the Do not include a host profile check box.
If you want to select a new location for the selected hosts	Select a data center, folder, or cluster as host location.

- 2 Click **Next**.

View the Summary of the Clone Deploy Rule Wizard

You can review the settings of the cloned vSphere Auto Deploy rule before completing the wizard.

Procedure

- 1 On the Ready to complete page, review the summary information for the new rule.
- 2 Click **Finish**.

You can view the newly created rule listed on the **Deploy Rules** tab.

Editing a Deploy Rule

You can edit a vSphere Auto Deploy rule only if it is in inactive state in the inventory. You can edit the name of the rule, the matching hosts, the assigned image profile, host profile, and host location.

- [Edit the Name and Matching Hosts of a Rule](#)

If a rule in the inventory is in inactive state, you can edit its name and change the selection of hosts that match the rule criteria.

- [Edit a Rule to Assign a Different Image Profile to Hosts](#)

If a rule in the inventory is in inactive state, you can edit the rule and assign a different image profile to the hosts matching it.

- [Edit a Rule to Assign a Different Host Profile to Hosts](#)

If a rule in the inventory is in inactive state, you can edit the rule and assign a different host profile to the hosts that match the criteria for the rule.

- [Edit the Host Location of a Rule](#)

If a rule in the inventory is in inactive state, you can edit the rule and assign a different host location to the hosts that match the rule criteria.

Edit the Name and Matching Hosts of a Rule

If a rule in the inventory is in inactive state, you can edit its name and change the selection of hosts that match the rule criteria.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.
By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.
- 2 On the **Deploy Rules** tab, from the list of rules in the inventory select the rule that you want to edit and click **Edit**.

The Edit Deploy Rule dialog box appears.

3 Select the **Name and hosts** page and enter a new name of the rule.

4 Select a pattern to apply the rule to the hosts in the inventory.

You can select to apply the rule to all the hosts in the inventory or to apply the rule only to hosts that match a specific pattern. You can select one or more patterns.

For example, the rule can apply only to hosts in a vCenter Single Sign-On domain, with a specific host name, or that match a specific IPv4 range.

5 Click **OK**.

Edit a Rule to Assign a Different Image Profile to Hosts

If a rule in the inventory is in inactive state, you can edit the rule and assign a different image profile to the hosts matching it.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).

Procedure

1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.

2 On the **Deploy Rules** tab, from the list of rules in the inventory select the rule that you want to edit and click **Edit**.

The Edit Deploy Rule dialog box appears.

3 Select the **Select image profile** page to assign an image profile to the hosts that match the rule criteria.

Option	Action
If you do not want to change the image profile	Select the Same image profile option.
If you do not want to assign an image profile to the selected hosts	Select the No image profile option.
If you want to assign a new image profile to the selected hosts	<ol style="list-style-type: none"> 1 Select the Browse for Image Profile option. 2 Select a software depot from the drop-down menu. 3 Select an image profile from the list. 4 (Optional) If you want to bypass the acceptance level verification for the image profile, select the Skip image profile signature check check box.

4 Click **OK**.

Edit a Rule to Assign a Different Host Profile to Hosts

If a rule in the inventory is in inactive state, you can edit the rule and assign a different host profile to the hosts that match the criteria for the rule.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.

- 2 On the **Deploy Rules** tab, from the list of rules in the inventory select the rule that you want to edit and click **Edit**.

The Edit Deploy Rule dialog box appears.

- 3 Select the **Select host profile** page and assign a new host profile to the hosts matching the rule.

Option	Action
If you do not want to assign a host profile to the selected hosts	Select the Do not include a host profile check box.
If you want to assign a host profile to the selected hosts	Select a host profile from the list.

- 4 Click **OK**.

Edit the Host Location of a Rule

If a rule in the inventory is in inactive state, you can edit the rule and assign a different host location to the hosts that match the rule criteria.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.

- 2 On the **Deploy Rules** tab, from the list of rules in the inventory select the rule that you want to edit and click **Edit**.

The Edit Deploy Rule dialog box appears.

- 3 Select the **Select host location** page and select a host location for the hosts matching the rule.

Option	Action
If you do not want to select a host location	Select the Do not include a location check box.
If you want to select a specific location for the selected hosts	Select a data center, folder, or cluster as host location.

- 4 Click **OK**.

Activate, Deactivate, and Reorder Deploy Rules

After you create a vSphere Auto Deploy rule, the rule is in inactive state. You must activate the rule for it to take effect. You can use the Activate and Reorder wizard to activate, deactivate, and change the order of the rules.

The upper list on the **Activate and Reorder** page of the wizard displays the rules in the active rule set. The lower list displays the inactive rules.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.
By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.
- 2 On the **Deploy Rules** tab, click **Activate/Deactivate rules**.
The Activate and Reorder wizard appears.
- 3 (Optional) If you want to deactivate an active rule, select the rule from the active rules list and click the **Deactivate** button.
- 4 From the list of inactive rules, select the rule that you want to activate and click the **Activate** button.
- 5 (Optional) If you want to reorder the rules in the active rule list, select a rule that you want to move up or down in the list and click the **Move up** or **Move down** icon above the list of active rules.

The rules are listed by priority. For example, if two or more rules apply to the same host but are set to provision the host with different image profiles, host profiles, and locations, the rule that is highest in the list takes effect on the host.

- 6 (Optional) If you want to test an inactive rule before activation, select the **Test rules before activation** check box and click **Next**.
 - a On the Select test targets page of the wizard, from the **Filter** tab select the hosts on which to test the inactive rule and click **Next**.

The **Selected** tab displays only the selected hosts.
 - b On the Preview test results page of the wizard, select a host from the list to view the current status of the host and the changes that are expected after the activation of the rule.

If the host is compliant with the rule, you do not need to remediate the host after you activate the rule.
 - c (Optional) If you want to remediate the selected hosts after the rule activation, select the **Remediate listed host associations after rule activation** check box.
- 7 Click **Next**.
- 8 Review the list of active rules and click **Finish**.

On the **Deploy Rules** tab, the rule is listed as active in the Status column.

What to do next

- View the image profile, host profile, and location associations of a host. See [View Host Associations](#).
- Remediate non-compliant hosts. See [Remediate a Non-compliant Host](#).

View Host Associations

Some of the hosts in the vSphere Auto Deploy inventory might not be compliant with the active deploy rules. To verify that one or more ESXi hosts are compliant with the active rule set, you must check the host associations compliance.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).
- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.

2 Check the host associations compliance

The **Check Host Associations Compliance** window displays the current status of the host and whether the host is compliant with the active rule set. You can view the currently assigned image profile, host profile, host location, script bundle, and the associations that will take effect after a remediation of the host. You can assign a script bundle to a host only by using PowerCLI cmdlets.

Option	Steps
If you want to check the host associations compliance of a single host	<ol style="list-style-type: none"> 1 On the Deployed Hosts tab, select an ESXi host. 2 Click Check Host Associations Compliance. 3 Check if the host associations are compliant with the current active rule set. 4 Close the Check Host Associations Compliance window. <ul style="list-style-type: none"> ■ If you want to remediate the host, click Remediate. ■ If you do not want to remediate the host, click Close.
If you want to check the host associations compliance of multiple hosts	<ol style="list-style-type: none"> 1 On the Deployed Hosts tab, use Shift+left-click or Ctrl+left-click to select multiple ESXi hosts. 2 Click Check Host Associations Compliance. 3 Confirm that you want to check the compliance of all selected hosts. 4 Review the compliance status of the hosts in the left pane. 5 (Optional) In the left pane, select a host to view the compliance status details in the right pane. 6 (Optional) Select one or multiple hosts and click Remediate Selected Hosts to remediate them. <ul style="list-style-type: none"> ■ Click the check box of each host you want to select. ■ Click the Hosts check box to select all hosts. 7 Click Close to close the Check Host Associations Compliance window.

What to do next

- Remediate non-compliant hosts. See [Remediate a Non-compliant Host](#).
- Edit the image profile association of a host. See [Edit the Image Profile Association of a Host](#).
- Edit a vSphere Auto Deploy rule. See [Editing a Deploy Rule](#).

Edit the Image Profile Association of a Host

You can edit the image profile association of a single host if the host is not associated with a vSphere Auto Deploy rule or if you do not want to change the image profile association of multiple hosts by editing a rule.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).
- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.

- 2 On the **Deployed Hosts** tab, select an ESXi host.

- 3 Click **Edit Image Profile Association**.

The Edit Image Profile Association dialog box appears.

- 4 Edit the image profile association of the host.

Option	Action
If you do not want to change the image profile	Select the Same image profile option.
If you want to assign a new image profile to the selected hosts	<ol style="list-style-type: none"> 1 Select the Browse for Image Profile option. 2 Select a software depot from the drop-down menu. 3 Select an image profile from the list. 4 (Optional) If you want to bypass the acceptance level verification for the image profile, select the Skip image profile signature check check box.

- 5 Click **OK**.

The new image profile is listed in the Image Profile column after a refresh of the page.

What to do next

- View the image profile, host profile, and location associations of a host. See [View Host Associations](#).
- If the host is associated with a rule and you want to revert to the image profile defined in the rule, remediate the host. See [Remediate a Non-compliant Host](#).

Remediate a Non-compliant Host

When you add a rule to the vSphere Auto Deploy active rule set or make changes to one or more rules, hosts are not updated automatically. You must remediate the host associations to apply the new rules to the host.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).
- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).
- If the remediation of a host, results in a change in its location, the host must be placed in maintenance mode.

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.

- 2 On the **Deployed Hosts** tab, select an ESXi host.

You can use Shift+left-click or Ctrl+left-click to select multiple hosts

- 3 Click **Remediate Host Associations**.

If you remediate a host that has an edited image profile association, the host reverts to the settings defined in the rule that it matches.

You can monitor the progress of the remediation process in the Recent Tasks pane.

What to do next

- View the image profile, host profile, and location associations of a host. See [View Host Associations](#).
- Change the image profile association of a host. See [Edit the Image Profile Association of a Host](#).

Add a Host to the vSphere Auto Deploy Inventory

You can view the hosts that do not match any vSphere Auto Deploy rule and manually add a host to the vSphere Auto Deploy inventory.

To add a host to the current vSphere Auto Deploy inventory of deployed hosts, you can create a new rule or edit an existing rule to include a host that is not deployed with vSphere Auto Deploy and associate it with a specific image profile, host profile and location. Alternatively, you can add a host manually to the inventory by assigning to it an image profile, host profile, and location.

Prerequisites

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- To assign an image profile to the host, add the software depot that you need to the inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).

Procedure

- 1 [Start the Add to Inventory Wizard](#)

You can add a host that does not correspond to any vSphere Auto Deploy rule to the list of deployed hosts by using the Add to Inventory wizard.

- 2 [Select an Image Profile in the Add to Inventory Wizard](#)

You can assign an image profile to a host that you want to add to the vSphere Auto Deploy inventory.

- 3 [Select a Host Profile in the Add to Inventory Wizard](#)

You can optionally assign a host profile to a host that you want to add to the vSphere Auto Deploy inventory.

4 [Select a Host Location in the Add to Inventory Wizard](#)

You can assign a location to a host that you want to add to the vSphere Auto Deploy inventory.

5 [View the Summary of the Add to Inventory Wizard](#)

You can review the host associations before you complete the Add to Inventory wizard.

What to do next

- Edit a vSphere Auto Deploy rule. See [Editing a Deploy Rule](#).
- View the image profile, host profile, and location associations of a host. See [View Host Associations](#).
- Remediate non-compliant hosts. See [Remediate a Non-compliant Host](#).

Start the Add to Inventory Wizard

You can add a host that does not correspond to any vSphere Auto Deploy rule to the list of deployed hosts by using the Add to Inventory wizard.

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.

- 2 On the **Discovered Hosts** tab, select one or more hosts that you want to provision with an image profile, host profile, and location.
- 3 Select **Add to Inventory**.

The Add to Inventory wizard appears.

Select an Image Profile in the Add to Inventory Wizard

You can assign an image profile to a host that you want to add to the vSphere Auto Deploy inventory.

Prerequisites

To assign an image profile to the host, add the software depot that you need to the inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).

Procedure

- 1 Select a software depot from the drop-down menu.
- 2 Select an image profile from the list of image profiles in the selected depot.
- 3 If you want to bypass the acceptance level verification of the image profile, select the **Skip image profile signature check** check box.
- 4 Select an image profile.
- 5 Click **Next**.

Select a Host Profile in the Add to Inventory Wizard

You can optionally assign a host profile to a host that you want to add to the vSphere Auto Deploy inventory.

Procedure

- 1 On the Select host profile page of the wizard, select a host profile.

Option	Action
If you do not want to assign a host profile to the selected hosts	Select the Do not include a host profile check box.
If you want to assign a host profile to the selected hosts	Select a host profile from the list.

- 2 Click **Next**.

Select a Host Location in the Add to Inventory Wizard

You can assign a location to a host that you want to add to the vSphere Auto Deploy inventory.

Procedure

- 1 Select a data center, folder, or cluster as the location for the host.
- 2 Click **Next**.

View the Summary of the Add to Inventory Wizard

You can review the host associations before you complete the Add to Inventory wizard.

Procedure

- 1 On the Ready to complete page, review the selected host associations.
- 2 Click **Finish**.

Provisioning ESXi Systems with vSphere Auto Deploy

vSphere Auto Deploy can provision hundreds of physical hosts with ESXi software. You can provision hosts that did not previously run ESXi software (first boot), reboot hosts, or reprovision hosts with a different image profile, host profile, custom script, or folder or cluster location.

The vSphere Auto Deploy process differs depending on the state of the host and on the changes that you want to make.

Provision a Host (First Boot)

Provisioning a host that has never been provisioned with vSphere Auto Deploy (first boot) differs from subsequent boot processes. You must prepare the host and fulfill all other prerequisites before you can provision the host. You can optionally define a custom image profile with vSphere ESXi Image Builder by using the vSphere Web Client or PowerCLI cmdlets.

Prerequisites

- Make sure your host meets the hardware requirements for ESXi hosts.
See [ESXi Hardware Requirements](#).
- Prepare the system for vSphere Auto Deploy. See [Preparing for vSphere Auto Deploy](#).
- Write rules that assign an image profile to the host and optionally assign a host profile and a vCenter Server location to the host. See [Managing vSphere Auto Deploy with PowerCLI Cmdlets](#) or [Managing vSphere Auto Deploy with the vSphere Web Client](#).

When the setup is complete, the vSphere Auto Deploy service is enabled, DHCP setup is complete, and rules for the host that you want to provision are in the active rule set.

Procedure

- 1 Turn on the host.

The host contacts the DHCP server and downloads iPXE from the location the server points it to. Next, the vSphere Auto Deploy server provisions the host with the image specified by the rule engine. The vSphere Auto Deploy server might also apply a host profile to the host if one is specified in the rule set. Finally, vSphere Auto Deploy adds the host to the vCenter Server system that is specified in the rule set.

- 2 (Optional) If vSphere Auto Deploy applies a host profile that requires user input such as an IP address, the host is placed in maintenance mode. Reapply the host profile with the vSphere Web Client and provide the user input when prompted.

After the first boot process, the host is running and managed by a vCenter Server system. The vCenter Server stores the host's image profile, host profile, and location information.

You can now reboot the host as needed. Each time you reboot, the host is reprovisioned by the vCenter Server system.

What to do next

Reprovision hosts as needed. See [Reprovisioning Hosts](#).

If you want to change the image profile, host profile, custom script, or location of the host, update the rules and activate them by using the vSphere Web Client or perform a test and repair compliance operation in a PowerCLI session. See [Rules and Rule Sets](#) or [Test and Repair Rule Compliance](#).

Reprovisioning Hosts

vSphere Auto Deploy supports multiple reprovisioning options. You can perform a simple reboot or reprovision with a different image profile or a different host profile.

A first boot using vSphere Auto Deploy requires that you set up your environment and add rules to the rule set. See [Preparing for vSphere Auto Deploy](#).

The following reprovisioning operations are available.

- Simple reboot.
- Reboot of hosts for which the user answered questions during the boot operation.
- Reprovision with a different image profile.
- Reprovision with a different host profile.

Reprovision Hosts with Simple Reboot Operations

A simple reboot of a host that is provisioned with vSphere Auto Deploy requires only that all prerequisites are still met. The process uses the previously assigned image profile, host profile, custom script, and vCenter Server location.

Prerequisites

- Verify that the setup you performed during the first boot operation is in place. See [Provision a Host \(First Boot\)](#).
- Verify that all associated items like are available. An item can be an image profile, host profile, custom script or vCenter Server inventory location.
- Verify that the host has the identifying information (asset tag, IP address) it had during previous boot operations.

Procedure

- 1 Place the host in maintenance mode.

Host Type	Action
Host is part of a DRS cluster	VMware DRS migrates virtual machines to appropriate hosts when you place the host in maintenance mode.
Host is not part of a DRS cluster	You must migrate all virtual machines to different hosts and place each host in maintenance mode.

- 2 Reboot the host.

The host shuts down. When the host reboots, it uses the image profile that the vSphere Auto Deploy server provides. The vSphere Auto Deploy server also applies the host profile stored on the vCenter Server system.

Reprovision a Host with a New Image Profile by Using PowerCLI

You can use vSphere Auto Deploy to reprovision a host with a new image profile in a PowerCLI session by changing the rule for the host and performing a test and repair compliance operation.

Several options for reprovisioning hosts exist.

- If the VIBs that you want to use support live update, you can use an `esxcli software vib` command. In that case, you must also update the rule set to use an image profile that includes the new VIBs.
- During testing, you can apply an image profile to an individual host with the `Apply-ESXImageProfile` cmdlet and reboot the host so the change takes effect. The `Apply-ESXImageProfile` cmdlet updates the association between the host and the image profile but does not install VIBs on the host.
- In all other cases, use this procedure.

Prerequisites

- Verify that the image profile you want to use to reprovision the host is available. Use vSphere ESXi Image Builder in a PowerCLI session. See [Customizing Installations with vSphere ESXi Image Builder](#).
- Verify that the setup you performed during the first boot operation is in place.

Procedure

- 1 At the PowerShell prompt, run the `Connect-VIServer` PowerCLI cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

- 2 Determine the location of a public software depot that contains the image profile that you want to use, or define a custom image profile with vSphere ESXi Image Builder.
- 3 Run `Add-EsxSoftwareDepot` to add the software depot that contains the image profile to the PowerCLI session.

Depot Type	Cmdlet
Remote depot	Run <code>Add-EsxSoftwareDepot <i>depot_url</i></code> .
ZIP file	<ol style="list-style-type: none"> a Download the ZIP file to a local file path or create a mount point local to the PowerCLI machine. b Run <code>Add-EsxSoftwareDepot C:\file_path\my_offline_depot.zip</code>.

- 4 Run `Get-ESXImageProfile` to see a list of image profiles, and decide which profile you want to use.

- 5 Run `Copy-DeployRule` and specify the `ReplaceItem` parameter to change the rule that assigns an image profile to hosts.

The following cmdlet replaces the current image profile that the rule assigns to the host with the `my_new_imageprofile` profile. After the cmdlet completes, `myrule` assigns the new image profile to hosts. The old version of `myrule` is renamed and hidden.

```
Copy-DeployRule myrule -ReplaceItem my_new_imageprofile
```

- 6 Test the rule compliance for each host that you want to deploy the image to.

- a Verify that you can access the host for which you want to test rule set compliance.

```
Get-VMHost -Name ESXi_hostname
```

- b Run the cmdlet that tests rule set compliance for the host, and bind the return value to a variable for later use.

```
$tr = Test-DeployRuleSetCompliance ESXi_hostname
```

- c Examine the differences between the contents of the rule set and configuration of the host.

```
$tr.itemlist
```

The system returns a table of current and expected items if the host for which you want to test the new rule set compliance is compliant with the active rule set.

CurrentItem	ExpectedItem
-----	-----
<code>my_old_imageprofile</code>	<code>my_new_imageprofile</code>

- d Remediate the host to use the revised rule set the next time you boot the host.

```
Repair-DeployRuleSetCompliance $tr
```

- 7 Reboot the host to provision it with the new image profile.

Reprovision a Host with a New Image Profile by Using the vSphere Web Client

You can use vSphere Auto Deploy to reprovision a host with a new image profile with the vSphere Web Client by changing the rule that the host corresponds to and activating the rule.

Prerequisites

- Verify that the image profile you want to use to reprovision the host is available. See [Create an Image Profile](#).
- Verify that the setup you performed during the first boot operation is in place.

Procedure

- 1 On the vSphere Web Client Home page, click **Auto Deploy**.

By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.

- 2 On the **Deploy Rules** tab, from the list of rules in the inventory select the rule that you want to edit and click **Edit**.

The Edit Deploy Rule dialog box appears.

- 3 Select the **Select image profile** page to assign an image profile to the hosts that match the rule criteria.

Option	Action
If you do not want to change the image profile	Select the Same image profile option.
If you do not want to assign an image profile to the selected hosts	Select the No image profile option.
If you want to assign a new image profile to the selected hosts	<ol style="list-style-type: none"> 1 Select the Browse for Image Profile option. 2 Select a software depot from the drop-down menu. 3 Select an image profile from the list. 4 (Optional) If you want to bypass the acceptance level verification for the image profile, select the Skip image profile signature check check box.

- 4 Click **Activate/Deactivate rules**.

- 5 From the list of inactive rules, select the rule that you want to activate and click the **Activate** button.

- 6 (Optional) If you want to reorder the rules in the active rule list, select a rule that you want to move up or down in the list and click the **Move up** or **Move down** icon above the list of active rules.

The rules are listed by priority. For example, if two or more rules apply to the same host but are set to provision the host with different image profiles, host profiles, and locations, the rule that is highest in the list takes effect on the host.

- 7 (Optional) If you want to test an inactive rule before activation, select the **Test rules before activation** check box and click **Next**.

- a On the Select test targets page of the wizard, from the **Filter** tab select the hosts on which to test the inactive rule and click **Next**.

The **Selected** tab displays only the selected hosts.

- b On the Preview test results page of the wizard, select a host from the list to view the current status of the host and the changes that are expected after the activation of the rule.

If the host is compliant with the rule, you do not need to remediate the host after you activate the rule.

- c (Optional) If you want to remediate the selected hosts after the rule activation, select the **Remediate listed host associations after rule activation** check box.

- 8 Click **Next**.

- 9 Review the list of active rules and click **Finish**.
- 10 Reboot the host to provision it with the new image profile.

Update the Host Customization in the vSphere Web Client

If a host required user input during a previous boot, the answers are saved with the vCenter Server. If you want to prompt the user for new information, you must remediate the host.

Prerequisites

Attach a host profile that prompts for user input to the host.

Procedure

- 1 Migrate all virtual machines to different hosts, and place the host into maintenance mode.

Host Type	Action
Host is part of a DRS cluster	VMware DRS migrates virtual machines to appropriate hosts when you place the host in maintenance mode.
Host is not part of a DRS cluster	You must migrate all virtual machines to different hosts and place each host in maintenance mode.

- 2 On the vSphere Web Client Home page, click **Auto Deploy**.
By default, only the Administrator role has privileges to use the vSphere Auto Deploy service.
- 3 On the **Deployed Hosts** tab, select an ESXi host.
- 4 Click **Remediate Host Associations**.
You can monitor the progress of the remediation process in the Recent Tasks pane.
- 5 When prompted, provide the user input.
- 6 Direct the host to exit maintenance mode.

The host customization is saved and takes effect the next time you boot the host.

Using vSphere Auto Deploy for Stateless Caching and Stateful Installs

The vSphere Auto Deploy stateless caching feature lets you cache the host's image. The vSphere Auto Deploy stateful installs feature lets you install hosts over the network. After the initial network boot, these hosts boot like other ESXi hosts.

The stateless caching solution is primarily intended for situations when several hosts boot simultaneously. The locally cached image helps prevent a bottleneck that results if several hundreds of hosts connect to the vSphere Auto Deploy server simultaneously. After the boot operation is complete, hosts connect to vSphere Auto Deploy to complete the setup.

The stateful installs feature lets you provision hosts with the image profile over the network without having to set up the PXE boot infrastructure.

- [Introduction to Stateless Caching and Stateful Installs](#)

You can use the System Cache Configuration host profile to provision hosts with vSphere Auto Deploy stateless caching and stateful installs.

- [Understanding Stateless Caching and Stateful Installs](#)

When you want to use vSphere Auto Deploy with stateless caching or stateful installs, you must set up a host profile, apply the host profile, and set the boot order.

- [Configure a Host Profile to Use Stateless Caching](#)

When a host is set up to use stateless caching, the host uses a cached image if the vSphere Auto Deploy Server is not available. To use stateless caching, you must configure a host profile. You can apply that host profile to other hosts that you want to set up for stateless caching.

- [Configure a Host Profile to Enable Stateful Installs](#)

To set up a host provisioned with vSphere Auto Deploy to boot from disk, you must configure a host profile. You can apply that host profile to other hosts that you want to set up for stateful installs.

Introduction to Stateless Caching and Stateful Installs

You can use the System Cache Configuration host profile to provision hosts with vSphere Auto Deploy stateless caching and stateful installs.

Examples of Stateless Caching and Stateful Installs

Hosts provisioned with vSphere Auto Deploy cache the image (stateless caching)

Set up and apply a host profile for stateless caching. You can cache the image on a local disk, a remote disk, or a USB drive. Continue provisioning this host with vSphere Auto Deploy. If the vSphere Auto Deploy server becomes unavailable, for example because hundreds of hosts attempt to access it simultaneously, the host boots from the cache. The host attempts to reach the vSphere Auto Deploy server after the boot operation to complete configuration.

Hosts provisioned with vSphere Auto Deploy become stateful hosts

Set up and apply a host profile for stateful installs. When you provision a host with vSphere Auto Deploy, the image is installed on the local disk, a remote disk, or a USB drive. For subsequent boots, you boot from the disk. The host no longer uses vSphere Auto Deploy.

Preparation

To successfully use stateless caching or stateful installs, decide how to configure the system and set the boot order.

Table 2-17. Preparation for Stateless Caching or Stateful Installs

Requirement or Decision	Description
Decide on VMFS partition overwrite	<p>When you install ESXi by using the interactive installer, you are prompted whether you want to overwrite an existing VMFS datastore. The System Cache Configuration host profile provides an option to overwrite existing VMFS partitions.</p> <p>The option is not available if you set up the host profile to use a USB drive.</p>
Decide whether you need a highly available environment	<p>If you use vSphere Auto Deploy with stateless caching, you can set up a highly available vSphere Auto Deploy environment to guarantee that virtual machines are migrated on newly provisioned hosts and that the environment supports vNetwork Distributed Switch even if the vCenter Server system becomes temporarily unavailable.</p>
Set the boot order	<p>The boot order you specify for your hosts depends on the feature you want to use.</p> <ul style="list-style-type: none"> ■ To set up vSphere Auto Deploy with stateless caching, configure your host to first attempt to boot from the network, and to then attempt to boot from disk. If the vSphere Auto Deploy server is not available, the host boots using the cache. ■ To set up vSphere Auto Deploy for stateful installs on hosts that do not currently have a bootable disk, configure your hosts to first attempt to boot from disk, and to then attempt to boot from the network. <p>Note If you currently have a bootable image on the disk, configure the hosts for one-time PXE boot, and provision the host with vSphere Auto Deploy to use a host profile that specifies stateful installs.</p>

Stateless Caching and Loss of Connectivity

If the ESXi hosts that run your virtual machines lose connectivity to the vSphere Auto Deploy server, the vCenter Server system, or both, some limitations apply the next time you reboot the host.

- If vCenter Server is available but the vSphere Auto Deploy server is unavailable, hosts do not connect to the vCenter Server system automatically. You can manually connect the hosts to the vCenter Server, or wait until the vSphere Auto Deploy server is available again.
- If both vCenter Server and vSphere Auto Deploy are unavailable, you can connect to each ESXi host by using the VMware Host Client, and add virtual machines to each host.
- If vCenter Server is not available, vSphere DRS does not work. The vSphere Auto Deploy server cannot add hosts to the vCenter Server. You can connect to each ESXi host by using the VMware Host Client, and add virtual machines to each host.
- If you make changes to your setup while connectivity is lost, the changes are lost when the connection to the vSphere Auto Deploy server is restored.

Understanding Stateless Caching and Stateful Installs

When you want to use vSphere Auto Deploy with stateless caching or stateful installs, you must set up a host profile, apply the host profile, and set the boot order.

When you apply a host profile that enables caching to a host, vSphere Auto Deploy partitions the specified disk. What happens next depends on how you set up the host profile and how you set the boot order on the host.

- vSphere Auto Deploy caches the image when you apply the host profile if **Enable stateless caching on the host** is selected in the System Cache Configuration host profile. No reboot is required. When you later reboot, the host continues to use the vSphere Auto Deploy infrastructure to retrieve its image. If the vSphere Auto Deploy server is not available, the host uses the cached image.
- vSphere Auto Deploy installs the image if **Enable stateful installs on the host** is selected in the System Cache Configuration host profile. When you reboot, the host initially boots using vSphere Auto Deploy to complete the installation. A reboot is then issued automatically, after which the host boots from disk, similar to a host that was provisioned with the installer. vSphere Auto Deploy no longer provisions the host.

You can apply the host profile from the vSphere Web Client, or write a vSphere Auto Deploy rule in a PowerCLI session that applies the host profile.

Using the vSphere Web Client to Set Up vSphere Auto Deploy for Stateless Caching or Stateful Installs

You can create a host profile on a reference host and apply that host profile to additional hosts or to a vCenter Server folder or cluster. The following workflow results.

- 1 You provision a host with vSphere Auto Deploy and edit that host's System Image Cache Configuration host profile.
- 2 You place one or more target hosts in maintenance mode, apply the host profile to each host, and instruct the host to exit maintenance mode.
- 3 What happens next depends on the host profile you selected.
 - If the host profile enabled stateless caching, the image is cached to disk. No reboot is required.
 - If the host profile enabled stateful installs, the image is installed. When you reboot, the host uses the installed image.

Using PowerCLI to Set Up vSphere Auto Deploy for Stateless Caching or Stateful Installs

You can create a host profile for a reference host and write a vSphere Auto Deploy rule that applies that host profile to other target hosts in a PowerCLI session. The following workflow results.

- 1 You provision a reference host with vSphere Auto Deploy and create a host profile to enable a form of caching.
- 2 You write a rule that provisions additional hosts with vSphere Auto Deploy and that applies the host profile of the reference host to those hosts.

- 3 vSphere Auto Deploy provisions each host with the image profile or by using the script bundle associated with the rule. The exact effect of applying the host profile depends on the host profile you selected.
 - For stateful installs, vSphere Auto Deploy proceeds as follows:
 - During first boot, vSphere Auto Deploy installs the image on the host.
 - During subsequent boots, the host boots from disk. The hosts do not need a connection to the vSphere Auto Deploy server.
 - For stateless caching, vSphere Auto Deploy proceeds as follows:
 - During first boot, vSphere Auto Deploy provisions the host and caches the image.
 - During subsequent boots, vSphere Auto Deploy provisions the host. If vSphere Auto Deploy is unavailable, the host boots from the cached image, however, setup can only be completed when the host can reach the vSphere Auto Deploy server.

Configure a Host Profile to Use Stateless Caching

When a host is set up to use stateless caching, the host uses a cached image if the vSphere Auto Deploy Server is not available. To use stateless caching, you must configure a host profile. You can apply that host profile to other hosts that you want to set up for stateless caching.

Prerequisites

- Decide which disk to use for caching and determine whether the caching process will overwrite an existing VMFS partition.
- In production environments, protect the vCenter Server system and the vSphere Auto Deploy server by including them in a highly available environment. Having the vCenter Server in a management cluster guarantees that VDS and virtual machine migration are available. If possible, also protect other elements of your infrastructure. See [Set Up Highly Available vSphere Auto Deploy Infrastructure](#).
- Set up your environment for vSphere Auto Deploy. See [Preparing for vSphere Auto Deploy](#).
- Verify that a disk with at least 1GB of free space is available. If the disk is not yet partitioned, partitioning happens when you apply the host profile.
- Set up the host to attempt a network boot first and to boot from disk if network boot fails. See your hardware vendor's documentation.
- Create a host profile. See the *Host Profiles* documentation.

Procedure

- 1 On the vSphere Web Client Home page, click **Host Profiles**.
- 2 Select the host profile you want to configure and select the **Manage** tab.
- 3 Click **Edit Host Profile**.
- 4 Leave the name and description and click **Next**.

- 5 On the Edit host profile page of the wizard, select **Advanced Configuration Settings > System Image Cache Configuration > System Image Cache Configuration**.
- 6 In the **System Image Cache Profile Settings** drop-down menu, choose a policy option.

Option	Description
Enable stateless caching on the host	Caches the image to disk.
Enable stateless caching to a USB disk on the host	Caches the image to a USB disk attached to the host.

- 7 (Optional) If you select **Enable stateless caching on the host**, specify the information about the disk to use.

Option	Description
Arguments for first disk	<p>By default, the system attempts to replace an existing ESXi installation, and then attempts to write to the local disk.</p> <p>You can use the Arguments for first disk field to specify a comma-separated list of disks to use, in order of preference. You can specify more than one disk. Use esx for the first disk with ESX installed on it, use model and vendor information, or specify the name of the vmkernel device driver. For example, to have the system first look for a disk with the model name ST3120814A, second for any disk that uses the mptsas driver, and third for the local disk, specify ST3120814A,mptsas,local as the value of this field.</p> <p>The first disk setting in the host profile specifies the search order for determining which disk to use for the cache. The search order is specified as a comma delimited list of values. The default setting local,esx,local specifies that vSphere Auto Deploy should first look for an existing local cache disk. The cache disk is identified as a disk with an existing ESXi software image. If vSphere Auto Deploy cannot find an existing cache disk, it searches for an available local disk device. When searching for an available disk vSphere Auto Deploy uses the first empty disk that does not have an existing VMFS partition.</p> <p>You can use the first disk argument only to specify the search order. You cannot explicitly specify a disk. For example, you cannot specify a specific LUN on a SAN.</p>
Check to overwrite any VMFS volumes on the selected disk	If you select this check box, the system overwrites existing VMFS volumes if not enough space is available to store the image, image profile, and host profile.
Check to ignore any SSD devices connected to the host	If you select this check box, the system ignores any existing SSD devices and does not store image profiles and host profiles on them.

- 8 Click **Finish** to complete the host profile configuration.

What to do next

Apply the host profile to individual hosts by using the host profiles feature in the vSphere Web Client. See the *Host Profiles* documentation. Alternatively, you can create a rule to assign the host profile to hosts with the vSphere Web Client or by using PowerCLI. See [Write a Rule and Assign a Host Profile to Hosts](#).

- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).

- For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the new host profile.

Configure a Host Profile to Enable Stateful Installs

To set up a host provisioned with vSphere Auto Deploy to boot from disk, you must configure a host profile. You can apply that host profile to other hosts that you want to set up for stateful installs.

You can configure the host profile on a single host. You can also create a host profile on a reference host and apply that host profile to other hosts.

Prerequisites

- Decide which disk to use for storing the image, and determine whether the new image will overwrite an existing VMFS partition.
- Set up your environment for vSphere Auto Deploy. See [Preparing for vSphere Auto Deploy](#).
- Verify that a disk with at least 1GB of free space is available. If the disk is not yet partitioned, partitioning happens when you apply the host profile.
- Set up the host to boot from disk. See your hardware vendor's documentation.
- Create a host profile. See the *Host Profiles* documentation.

Procedure

- 1 On the vSphere Web Client Home page, click **Host Profiles**.
- 2 Select the host profile you want to configure and select the **Manage** tab.
- 3 Click **Edit Host Profile**.
- 4 Leave the name and description and click **Next**.
- 5 On the Edit host profile page of the wizard, select **Advanced Configuration Settings > System Image Cache Configuration > System Image Cache Configuration**.
- 6 In the **System Image Cache Profile Settings** drop-down menu, choose a policy option.

Option	Description
Enable stateful installs on the host	Caches the image to a disk.
Enable stateful installs to a USB disk on the host	Caches the image to a USB disk attached to the host.

- 7 (Optional) If you select **Enable stateful installs on the host**, specify information about the disk to use.

Option	Description
Arguments for first disk	<p>By default, the system attempts to replace an existing ESXi installation, and then attempts to write to the local disk.</p> <p>You can use the Arguments for first disk field to specify a comma-separated list of disks to use, in order of preference. You can specify more than one disk. Use esx for the first disk with ESX installed on it, use model and vendor information, or specify the name of the vmkernel device driver. For example, to have the system first look for a disk with the model name ST3120814A, second for any disk that uses the mptsas driver, and third for the local disk, specify ST3120814A,mptsas,local as the value of this field.</p> <p>The first disk setting in the host profile specifies the search order for determining which disk to use for the cache. The search order is specified as a comma delimited list of values. The default setting localesex,local specifies that vSphere Auto Deploy should first look for an existing local cache disk. The cache disk is identified as a disk with an existing ESXi software image. If vSphere Auto Deploy cannot find an existing cache disk, it searches for an available local disk device. When searching for an available disk vSphere Auto Deploy uses the first empty disk that does not have an existing VMFS partition.</p> <p>You can use the first disk argument only to specify the search order. You cannot explicitly specify a disk. For example, you cannot specify a specific LUN on a SAN.</p>
Check to overwrite any VMFS volumes on the selected disk	If you select this check box, the system overwrites existing VMFS volumes if not enough space is available to store the image, image profile, and host profile.
Check to ignore any SSD devices connected to the host	If you select this check box, the system ignores any existing SSD devices and does not store image profiles and host profiles on them.

- 8 Click **Finish** to complete the host profile configuration.

What to do next

Apply the host profile to individual hosts by using the host profiles feature in the vSphere Web Client. See the *Host Profiles* documentation. Alternatively, you can create a rule to assign the host profile to hosts with the vSphere Web Client or by using PowerCLI. See [Write a Rule and Assign a Host Profile to Hosts](#).

- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
- For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the new host profile.

Setting Up a vSphere Auto Deploy Reference Host

In an environment where no state is stored on the host, a reference host helps you set up multiple hosts with the same configuration. You configure the reference host with the logging, coredump, and other settings that you want, save the host profile, and write a rule that applies the host profile to other hosts as needed.

You can configure the storage, networking, and security settings on the reference host and set up services such as syslog and NTP.

Understanding Reference Host Setup

A well-designed reference host connects to all services such as syslog, NTP, and so on. The reference host setup might also include security, storage, networking, and ESXi Dump Collector. You can apply such a host's setup to other hosts by using host profiles.

The exact setup of your reference host depends on your environment, but you might consider the following customization.

NTP Server Setup

When you collect logging information in large environments, you must make sure that log times are coordinated. Set up the reference host to use the NTP server in your environment that all hosts can share. You can specify an NTP server by running the `vicfg-ntp` command. You can start and stop the NTP service for a host with the `vicfg-ntp` command, or the vSphere Web Client.

Syslog Server Setup

All ESXi hosts run a syslog service (`vmtoolsd`), which logs messages from the VMkernel and other system components to a file. You can specify the log host and manage the log location, rotation, size, and other attributes by running the `esxcli system syslog vCLI` command or by using the vSphere Web Client. Setting up logging on a remote host is especially important for hosts provisioned with vSphere Auto Deploy that have no local storage. You can optionally install the vSphere Syslog Collector to collect logs from all hosts.

Core Dump Setup

You can set up your reference host to send core dumps to a shared SAN LUN, or you can install ESXi Dump Collector in your environment and configure the reference host to use ESXi Dump Collector. See [Configure ESXi Dump Collector with ESXCLI](#). You can either install ESXi Dump Collector by using the vCenter Server installation media or use the ESXi Dump Collector that is included in the vCenter Server Appliance. After setup is complete, VMkernel memory is sent to the specified network server when the system encounters a critical failure.

Security Setup

In most deployments, all hosts that you provision with vSphere Auto Deploy must have the same security settings. You can, for example, set up the firewall to allow certain services to access the ESXi system, set up the security configuration, user configuration, and user group configuration for

the reference host with the vSphere Web Client or with vCLI commands. Security setup includes shared user access settings for all hosts. You can achieve unified user access by setting up your reference host to use Active Directory. See the *vSphere Security* documentation.

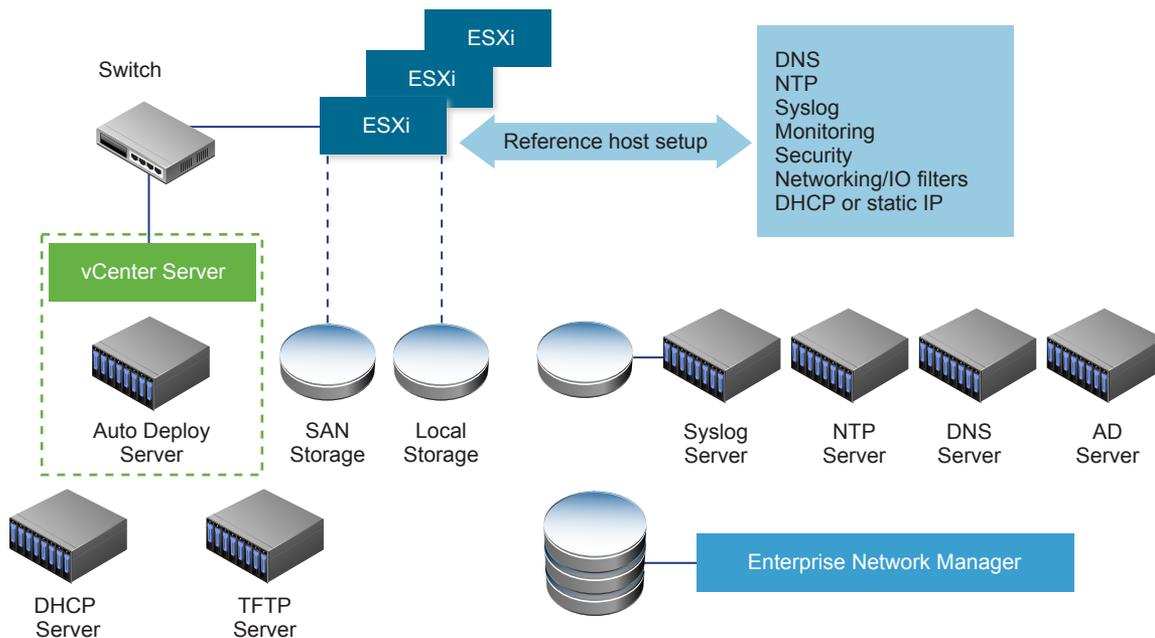
Note If you set up Active Directory by using host profiles, the passwords are not protected. Use the vSphere Authentication Service to set up Active Directory to avoid exposing the Active Directory password.

Networking and Storage Setup

If you reserve a set of networking and storage resources for use by hosts provisioned with vSphere Auto Deploy, you can set up your reference host to use those resources.

In very large deployments, the reference host setup supports an Enterprise Network Manager, which collects all information coming from the different monitoring services that are running in the environment.

Figure 2-7. vSphere Auto Deploy Reference Host Setup



[Options for Configuration of a vSphere Auto Deploy Reference Host](#) explains how to perform this setup.

Watch the video "Auto Deploy Reference Hosts" for information about the reference host setup:



vSphere Auto Deploy Reference Hosts

(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_auto_deploy_reference_hosts)

Options for Configuration of a vSphere Auto Deploy Reference Host

You can configure a reference host by using the vSphere Web Client, vCLI, or host profiles.

To set up a reference host, you can use the approach that suits you best.

vSphere Web Client	The vSphere Web Client supports setup of networking, storage, security, and most other aspects of an ESXi host. Set up your environment and create a host profile from the reference host for use by vSphere Auto Deploy.
vSphere Command-Line Interface	You can use vCLI commands for setup of many aspects of your host. vCLI is suitable for configuring many of the services in the vSphere environment. Commands include <code>vicfg-ntp</code> for setting up an NTP server, <code>esxcli system syslog</code> for setting up a syslog server, <code>esxcli network route</code> for adding routes and set up the default route, and <code>esxcli system coredump</code> for configuring Esxi Dump Collector.
Host Profiles Feature	Best practice is to set up a host with vSphere Web Client or vCLI and create a host profile from that host. You can instead use the Host Profiles feature in the vSphere Web Client and save that host profile.

vSphere Auto Deploy applies all common settings from the host profile to all target hosts. If you set up the host profile to prompt for user input, all hosts provisioned with that host profile come up in maintenance mode. You must reapply the host profile or reset host customizations to be prompted for the host-specific information.

Configure ESXi Dump Collector with ESXCLI

Hosts provisioned with vSphere Auto Deploy do not have a local disk to store core dumps on. You can configure ESXi Dump Collector by using ESXCLI commands and keep core dumps on a network server for use during debugging.

A core dump is the state of working memory in the event of host failure. By default, a core dump is saved to the local disk. ESXi Dump Collector is especially useful for vSphere Auto Deploy, but is supported for any ESXi host. ESXi Dump Collector supports other customization, including sending core dumps to the local disk and is included with the vCenter Server management node.

If you intend to use IPv6, and if both the ESXi host and ESXi Dump Collector are on the same local link, both can use either local link scope IPv6 addresses or global scope IPv6 addresses.

If you intend to use IPv6, and if ESXi and ESXi Dump Collector are on different hosts, both require global scope IPv6 addresses. The traffic routes through the default IPv6 gateway.

Prerequisites

Install vCLI if you want to configure the host to use ESXi Dump Collector. In troubleshooting situations, you can use ESXCLI in the ESXi Shell instead.

Procedure

- 1 Set up an ESXi system to use ESXi Dump Collector by running `esxcli system coredump` in the local ESXi Shell or by using vCLI.

```
esxcli system coredump network set --interface-name vmk0 --server-ip 10xx.xx.xx.xx --server-port 6500
```

You must specify a VMkernel NIC and the IP address and optional port of the server to send the core dumps to. You can use an IPv4 address or an IPv6 address. If you configure an ESXi system that is running on a virtual machine that is using a vSphere standard switch, you must select a VMkernel port that is in promiscuous mode.

- 2 Enable ESXi Dump Collector.

```
esxcli system coredump network set --enable true
```

- 3 (Optional) Verify that ESXi Dump Collector is configured correctly.

```
esxcli system coredump network check
```

The host on which you have set up ESXi Dump Collector is configured to send core dumps to the specified server by using the specified VMkernel NIC and optional port.

What to do next

- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
- For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the new host profile.

Configure ESXi Dump Collector from the Host Profiles Feature in the vSphere Web Client

Hosts provisioned with vSphere Auto Deploy do not have a local disk to store core dumps on. You can configure a reference host to use ESXi Dump Collector by using the Host Profiles feature in the vSphere Web Client.

Best practice is to set up hosts to use ESXi Dump Collector with the `esxcli system coredump` command and save the host profile. For more information, see [Configure ESXi Dump Collector with ESXCLI](#).

Prerequisites

- Verify that you have created the host profile on which you want to configure a coredump policy. For more information on how to create a host profile, see the *vSphere Host Profiles* documentation.

- Verify that at least one partition has sufficient storage capability for core dumps from multiple hosts provisioned with vSphere Auto Deploy.

Procedure

- 1 In the vSphere Web Client, click **Policies and Profiles**, and select **Host Profiles**.
- 2 Right-click the host profile you want to modify and select **Edit Settings**.
- 3 Leave the name and description unchanged and click **Next**.
- 4 On the Edit host profile page of the wizard, select **Networking Configuration > Network Coredump Settings**.
- 5 Select the **Enabled** check box.
- 6 Specify the host NIC to use, the Network Coredump Server IP, and the Network Coredump Server Port.
- 7 Click **Finish** to save the host profile settings.

What to do next

- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
- For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the new host profile.

Configure Syslog from the Host Profiles Feature in the vSphere Web Client

Hosts provisioned with vSphere Auto Deploy usually do not have sufficient local storage to save system logs. You can specify a remote syslog server for those hosts by setting up a reference host, saving the host profile, and applying that host profile to other hosts as needed.

Best practice is to set up the syslog server on the reference host with the vSphere Web Client or the `esxcli system syslog` command and to save the host profile. You can also set up syslog from the Host Profiles feature in the vSphere Web Client.

Prerequisites

- If you intend to use a remote syslog host, set up that host before you customize host profiles.
- Verify that you have access to a vSphere Web Client that can connect to the vCenter Server system.

Procedure

- 1 In the vSphere Web Client, click **Policies and Profiles**, and select **Host Profiles**.
- 2 (Optional) If no reference host exists in your environment, click the **Extract Profile from Host** icon to create a host profile.
- 3 Right-click the host profile you want to modify and select **Edit Settings**.

- 4 Leave the name and description unchanged and click **Next**.
- 5 On the Edit host profile page of the wizard, select **Advanced Configuration Settings > Advanced Options > Advanced configuration options**.

You can select specific sub-profiles and edit the syslog settings.

- 6 If you are setting up an ESXi 5.0 host that did not have a previously configured syslog server, you have to create an advanced configuration option.
 - a Click the **Add sub-profile** icon.
 - b Select the new sub-profile **Advanced configuration option** at the top of the list.
 - c From the **Advanced option** drop-down list select **Configure a fixed option**.
 - d Specify Syslog.global.loghost as the option, and your host as the value.

If you are configuring an ESXi host version 5.1 or later or an ESXi 5.0 host that has syslog configured, Syslog.global.loghost is already in the list of advanced options.

- 7 Click **Finish** to save the host profile settings.

What to do next

- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
- For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the new host profile.

Enable NTP Client on a Reference Host in the vSphere Web Client

When you collect logging information in large environments, you must ensure that log times are coordinated. You can set up the reference host to use the NTP server in your environment, extract the host profile and create a vSphere Auto Deploy rule to apply it to other hosts.

Procedure

- 1 In the vSphere Web Client navigator, browse to the host that you want to use as a reference host.
- 2 Select the **Manage** tab and select **Settings**.
- 3 Under **System**, select **Time Configuration** and click **Edit**.
- 4 Select the **Use Network Time Protocol (Enable NTP client)** radio button.

This option synchronizes the time and date of the host with an NTP server. The NTP service on the host periodically takes the time and date from the NTP server.

- 5 From the **NTP Service Startup Policy** drop-down list, select **Start and stop with host**.
- 6 In the **NTP Servers** text box, type the IP addresses or host names of the NTP servers that you want to use.

7 Click **OK**.

What to do next

- Extract a host profile from the reference host. See the *Host Profiles* documentation.
- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).

Configure Networking for Your vSphere Auto Deploy Host in the vSphere Web Client

You can set up networking for your vSphere Auto Deploy reference host and apply the host profile to all other hosts to guarantee a fully functional networking environment.

Prerequisites

Provision the host you want to use as your reference host with an ESXi image by using vSphere Auto Deploy.

Procedure

- 1 In the vSphere Web Client navigator, browse to the host that you want to use as a reference host.
- 2 Select the **Manage** tab and select **Networking**.
- 3 Perform the networking setup.

If you are using virtual switches and not vSphere Distributed Switch, do not add other VMkernel NICs to vSwitch0.

- 4 After the reference host is configured, reboot the system to verify that vmk0 is connected to the Management Network.
- 5 If no host profile exists for your reference host, create a host profile.

What to do next

- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
- For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the new host profile.

Configure a Reference Host for Auto-Partitioning

By default, vSphere Auto Deploy provisions hosts only if a partition is available on the host. The auto-partitioning option creates a VMFS datastore on your host's local storage. You can set up a reference host to auto-partition all hosts that you provision with vSphere Auto Deploy.

Caution If you change the default auto-partitioning behavior, vSphere Auto Deploy overwrites existing partitions regardless of their content. If you turn on this option, ensure that no unintended data loss results.

To ensure that local SSDs remain unpartitioned during auto-partitioning, you must set the parameter `skipPartitioningSsds=TRUE` on the reference host.

For more information about preventing SSD formatting during auto-partitioning, see the *vSphere Storage* documentation.

Prerequisites

- Provision the host that you want to use as your reference host with an ESXi image by using vSphere Auto Deploy.
- Verify that you have access to vSphere Web Client that can connect to the vCenter Server system.

Procedure

- 1 In the vSphere Web Client navigator, browse to the host that you want to use as a reference host.
- 2 Select the **Manage** tab and select **Settings**.
- 3 Under **System**, select **Advanced System Settings** and click **Edit**.
- 4 Scroll to `VMkernel.Boot.autoPartition` and select the **Enabled** check box.
- 5 (Optional) If you want the local SSDs to remain unpartitioned, scroll to `VMkernel.Boot.skipPartitioningSsds` and select the **Enabled** check box.
- 6 Click **OK**.
- 7 If no host profile exists for your reference host, create a host profile.

Auto-partitioning is performed when the hosts boot.

What to do next

- Use vSphere Auto Deploy to create a rule that applies the host profile of your reference host to all hosts immediately when they boot. To create a rule with the vSphere Web Client, see [Create a Deploy Rule](#). For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).

vSphere Auto Deploy Best Practices and Security Consideration

Follow best practices when installing vSphere Auto Deploy and when using vSphere Auto Deploy with other vSphere components. Set up a highly available vSphere Auto Deploy infrastructure in large production environments or when using stateless caching. Follow all security guidelines that you would follow in a PXE boot environment, and consider the recommendations in this chapter.

vSphere Auto Deploy Best Practices

You can follow several vSphere Auto Deploy best practices, set up networking, configure vSphere HA, and otherwise optimize your environment for vSphere Auto Deploy.

See the VMware Knowledge Base for additional best practice information.

vSphere Auto Deploy and vSphere HA Best Practices

You can improve the availability of the virtual machines running on hosts provisioned with vSphere Auto Deploy by following best practices.

Some environments configure the hosts provisioned with vSphere Auto Deploy with a distributed switch or configure virtual machines running on the hosts with Auto Start Manager. In such environments, deploy the vCenter Server system so that its availability matches the availability of the vSphere Auto Deploy server. Several approaches are possible.

- Install vCenter Server on a Windows virtual machine or physical server or deploy the vCenter Server Appliance. Auto Deploy is deployed together with the vCenter Server system.
- Deploy the vCenter Server system on a virtual machine. Run the vCenter Server virtual machine in a vSphere HA enabled cluster and configure the virtual machine with a vSphere HA restart priority of high. Include two or more hosts in the cluster that are not managed by vSphere Auto Deploy and pin the vCenter Server virtual machine to these hosts by using a rule (vSphere HA DRS required VM to host rule). You can set up the rule and then disable DRS if you do not want to use DRS in the cluster. The greater the number of hosts that are not managed by vSphere Auto Deploy, the greater your resilience to host failures.

Note This approach is not suitable if you use Auto Start Manager. Auto Start Manager is not supported in a cluster enabled for vSphere HA.

vSphere Auto Deploy Networking Best Practices

Prevent networking problems by following vSphere Auto Deploy networking best practices.

vSphere Auto Deploy and IPv6

Because vSphere Auto Deploy takes advantage of the iPXE infrastructure, if the hosts that you plan to provision with vSphere Auto Deploy are with legacy BIOS, the vSphere Auto Deploy server must have an IPv4 address. PXE booting with legacy BIOS firmware is possible only over IPv4. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.

IP Address Allocation

Use DHCP reservations for address allocation. Fixed IP addresses are supported by the host customization mechanism, but providing input for each host is not recommended.

VLAN Considerations

Use vSphere Auto Deploy in environments that do not use VLANs.

If you intend to use vSphere Auto Deploy in an environment that uses VLANs, make sure that the hosts that you want to provision can reach the DHCP server. How hosts are assigned to a VLAN depends on the setup at your site. The VLAN ID might be assigned by the switch or the router, or might be set in the host's BIOS or through the host profile. Contact your network administrator to determine the steps for allowing hosts to reach the DHCP server.

vSphere Auto Deploy and VMware Tools Best Practices

When you provision hosts with vSphere Auto Deploy, you can select an image profile that includes VMware Tools, or select the smaller image associated with the image profile that does not contain VMware Tools.

You can download two image profiles from the VMware download site.

- `xxxxx-standard`: An image profile that includes the VMware Tools binaries, required by the guest operating system running inside a virtual machine. The image is usually named `esxi-version-xxxxx-standard`.
- `xxxxx-no-tools`: An image profile that does not include the VMware Tools binaries. This image profile is usually smaller has a lower memory overhead, and boots faster in a PXE-boot environment. This image is usually named `esxi-version-xxxxx-no-tools`.

With vSphere 5.0 Update 1 and later, you can deploy ESXi using either image profile.

- If the network boot time is of no concern, and your environment has sufficient extra memory and storage overhead, use the image that includes VMware Tools.
- If you find the network boot time too slow when using the standard image, or if you want to save some space on the hosts, you can use the image profile that does not include VMware Tools, and place the VMware Tools binaries on shared storage. See, [Provision ESXi Host by Using an Image Profile Without VMware Tools](#).

vSphere Auto Deploy Load Management Best Practices

Simultaneously booting large numbers of hosts places a significant load on the vSphere Auto Deploy server. Because vSphere Auto Deploy is a Web server at its core, you can use existing Web server scaling technologies to help distribute the load. For example, one or more caching reverse proxy servers can be used with vSphere Auto Deploy. The reverse proxies serve up the static files that make up the majority of an ESXi boot image. Configure the reverse proxy to cache static content and pass all requests through to the vSphere Auto Deploy server. For more information, watch the video "Using Reverse Web Proxy Servers for vSphere Auto Deploy Scalability":



Using Reverse Web Proxy Servers for vSphere Auto Deploy Scalability

(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_reverse_web_proxy_for_auto_deploy_scalability)

Use multiple TFTP servers to point to different proxy servers. Use one TFTP server for each reverse proxy server. After that, set up the DHCP server to send different hosts to different TFTP servers.

When you boot the hosts, the DHCP server redirects them to different TFTP servers. Each TFTP server redirects hosts to a different server, either the vSphere Auto Deploy server or a reverse proxy server, significantly reducing the load on the vSphere Auto Deploy server.

After a massive power outage, bring up the hosts on a per-cluster basis. If you bring multiple clusters online simultaneously, the vSphere Auto Deploy server might experience CPU bottlenecks. All hosts might come up after a delay. The bottleneck is less severe if you set up the reverse proxy.

vSphere Auto Deploy Logging and Troubleshooting Best Practices

To resolve problems that you encounter with vSphere Auto Deploy, use the vSphere Auto Deploy logging information from the vSphere Web Client and set up your environment to send logging information and core dumps to remote hosts.

vSphere Auto Deploy Logs

Download the vSphere Auto Deploy logs by going to the vSphere Auto Deploy page in the vSphere Web Client. See, [Download vSphere Auto Deploy Logs](#).

Setting Up Syslog

Set up a remote syslog server. See the *vCenter Server and Host Management* documentation for syslog server configuration information. Configure the first host you boot to use the remote syslog server and apply that host's host profile to all other target hosts. Optionally, install and use the vSphere Syslog Collector, a vCenter Server support tool that provides a unified architecture for system logging, enables network logging, and lets you combine logs from multiple hosts.

Setting Up ESXi Dump Collector

Hosts provisioned with vSphere Auto Deploy do not have a local disk to store core dumps on. Install ESXi Dump Collector and set up your first host so that all core dumps are directed to ESXi Dump Collector, and apply the host profile from that host to all other hosts. See [Configure ESXi Dump Collector with ESXCLI](#).

Using vSphere Auto Deploy in a Production Environment

When you move from a proof of concept setup to a production environment, take care to make the environment resilient.

- Protect the vSphere Auto Deploy server. See [vSphere Auto Deploy and vSphere HA Best Practices](#).
- Protect all other servers in your environment, including the DHCP server and the TFTP server.
- Follow VMware security guidelines, including those outlined in [vSphere Auto Deploy Security Considerations](#).

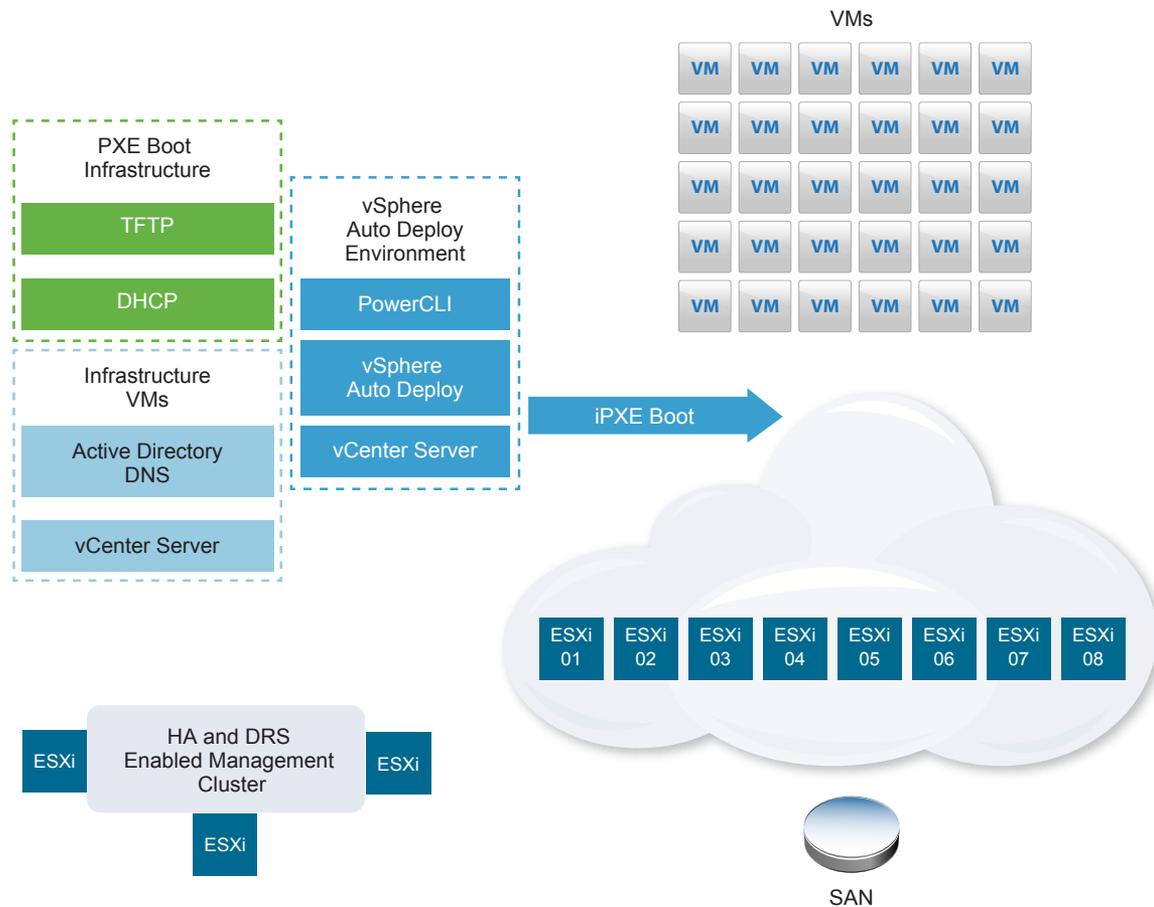
Set Up Highly Available vSphere Auto Deploy Infrastructure

In many production situations, a highly available vSphere Auto Deploy infrastructure is required to prevent data loss. Such infrastructure is also a prerequisite for using vSphere Auto Deploy with stateless caching.



Highly Available vSphere Auto Deploy Infrastructure
http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_ha_auto_deploy_infrastructure

Figure 2-8. Highly Available vSphere Auto Deploy Infrastructure



Prerequisites

For the management cluster, install ESXi on three hosts. Do not provision the management cluster hosts with vSphere Auto Deploy.

Watch the video "Highly Available vSphere Auto Deploy Infrastructure" for information about the implementation of a highly available vSphere Auto Deploy infrastructure:

Procedure

- 1 Enable vSphere HA and vSphere DRS on the management cluster.
- 2 Set up the following virtual machines on the management cluster.

Infrastructure Component	Description
PXE boot infrastructure	TFTP and DHCP servers.
Infrastructure VM	Active Directory, DNS, vCenter Server.
vSphere Auto Deploy environment	PowerCLI, vSphere Auto Deploy server, vCenter Server. Set up this environment on a single virtual machine or on three separate virtual machines in production systems.

The vCenter Server on the infrastructure virtual machine differs from the vCenter Server in the vSphere Auto Deploy environment.

- 3 Set up vSphere Auto Deploy to provision other hosts as needed.

Because the components on the management cluster are protected with vSphere HA, high availability is supported.

vSphere Auto Deploy Security Considerations

When you use vSphere Auto Deploy, pay careful attention to networking security, boot image security, and potential password exposure through host profiles to protect your environment.

Networking Security

Secure your network just as you secure the network for any other PXE-based deployment method. vSphere Auto Deploy transfers data over SSL to prevent casual interference and snooping. However, the authenticity of the client or of the Auto Deploy server is not checked during a PXE boot.

You can greatly reduce the security risk of Auto Deploy by completely isolating the network where Auto Deploy is used.

Boot Image and Host Profile Security

The boot image that the vSphere Auto Deploy server downloads to a machine can have the following components.

- The VIB packages that the image profile consists of are always included in the boot image.
- The host profile and host customization are included in the boot image if Auto Deploy rules are set up to provision the host with a host profile or host customization.
 - The administrator (root) password and user passwords that are included with host profile and host customization are MD5 encrypted.

- Any other passwords associated with profiles are in the clear. If you set up Active Directory by using host profiles, the passwords are not protected.

Use the vSphere Authentication Proxy to avoid exposing the Active Directory passwords. If you set up Active Directory using host profiles, the passwords are not protected.

- The host's public and private SSL key and certificate are included in the boot image.

Using the Device Alias Configuration Host Profile

In vSphere 5.5 and later, you can persistently map a device (bus address) to a device name (alias). You can modify the mapping by using the Device Alias Configuration host profile. Using persistent mapping can help avoid compliance warnings for stateless hosts, and is also useful for stateful hosts.

The Device Alias Configuration host profile is selected by default, which means that aliases are assigned to each device. For example, if a host does not recognize one of the NICs during the boot process, the NIC aliases no longer change. That can help for management with scripts, and if you apply a host profile from a reference host.

Note To avoid errors, do not disable or edit the Device Alias Configuration host profile.

To ensure uniform, persistent, and stable device naming across all hosts, use the device alias profile with homogeneous hosts only. These are hosts that are identically configured with the same network and storage cards in the PCI bus.

Note Always bring the BIOS up to the latest level. For systems with earlier versions of the BIOS, the BIOS might not provide accurate location information for on-board devices. ESXi applies heuristics for this case to keep the alias stable, even for these devices, this might not work under all conditions, for example if changes are made in the BIOS setting or if the devices fail.

Device Alias Configuration Compliance Failures

For hosts are not fully homogenous, for example, the hosts contain different PCI cards or have different BIOS levels, if you apply the host profile from a reference host, a compliance check might result in a compliance failure. The compliance check ignores extra devices on the host that were not on the reference host. Select the host with the fewest devices as the reference host.

If the compliance check shows that he hosts are not fully homogeneous, the compliance failure cannot be remediated without modifying the hardware itself.

If the compliance check shows that the device aliases, for example, names such as vmhba3, are different from those on the reference host, remediation might be possible.

- To remediate a host that is not provisioned with vSphere Auto Deploy, perform host profile remediation and reboot the host.
- To remediate a host that is provisioned with vSphere Auto Deploy, reprovision a host.

Upgrading Systems for Device Alias Profiles

In ESXi versions earlier than 5.5, the Device Alias Configuration profile does not exist. Consider the following problems when you upgrade from previous versions of ESXi to ESXi 5.5 and later:

- For installed hosts, that is, hosts not provisioned with vSphere Auto Deploy, upgrading the ESXi host preserves aliases. After they are upgraded, aliases remain stable as long as the BIOS provides the information.
- When you upgrade a cluster of ESXi host provisioned with vSphere Auto Deploy image, the aliases do not change because ESXi 5.5 uses the same algorithm to generate aliases as earlier versions. Generate a new host profile for the reference host. This host profile includes the Device Alias Configuration profile. Set up vSphere Auto Deploy to apply the reference host's host profile to all other hosts for consistent device naming across your cluster.
- When upgrading a system, do not flash the BIOS, because this action can change aliases. Flashing the BIOS to the latest level is more appropriate for a new install.

Provision ESXi Host by Using an Image Profile Without VMware Tools

When you provision ESXi hosts with vSphere Auto Deploy, you can select to provision the host by using the image profile that does not contain VMware Tools binaries. This image profile is usually smaller, has a lower memory overhead, and boots faster in a PXE-boot environment.

If you find the network boot time too slow when using the standard image, or if you want to save some space on the hosts, you can use the image profile that does not include VMware Tools, and place the VMware Tools binaries on a shared storage.

Prerequisites

Download the `xxxxx-no-tools` image profile from the VMware download site.

Procedure

- 1 Boot an ESXi host that was not provisioned with vSphere Auto Deploy.
- 2 Copy the `/productLocker` directory from the ESXi host to a shared storage.
- 3 Change the `UserVars.ProductLockerLocation` variable to point to the `/productLocker` directory.
 - a In the vSphere Web Client, select the reference host and click the **Manage** tab.
 - b Select **Settings** and click **Advanced System Settings**.
 - c Filter the settings for **uservars**, and select **UserVars.ProductLockerLocation**.
 - d Click the **pen** icon and edit the location so it points to the shared storage.
- 4 Create a host profile from the reference host.
- 5 Create a vSphere Auto Deploy rule that assigns the `xxxxx-no-tools` image profile and host profile from the reference host to all other hosts.
- 6 Boot your target hosts with the rule so they pick up the product locker location from the reference host.

Download vSphere Auto Deploy Logs

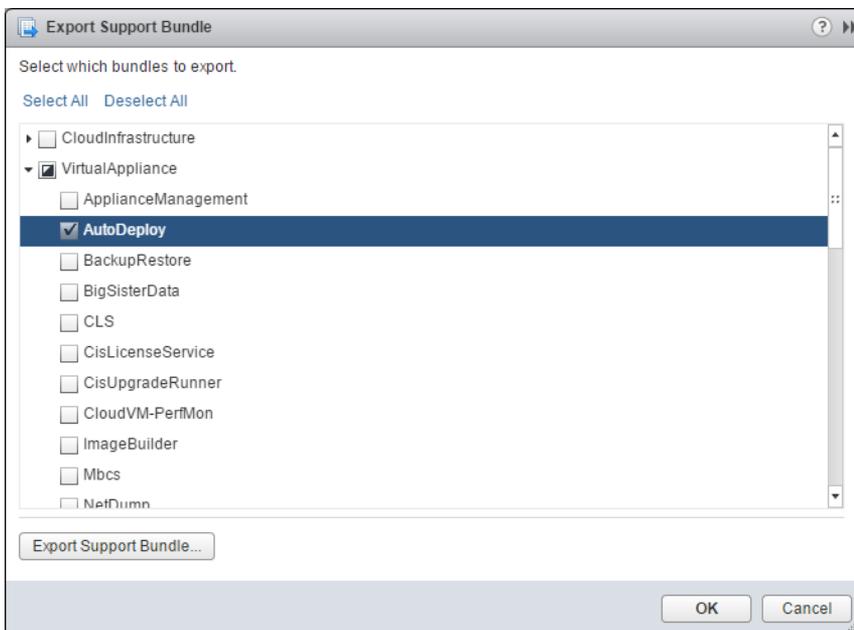
You can use the vSphere Auto Deploy logging information from the vSphere Web Client to resolve problems that you encounter with vSphere Auto Deploy.

Prerequisites

Use the vSphere Web Client to log in to the vCenter Server instance that vSphere Auto Deploy is registered with.

Procedure

- 1 From **Administration** select **Deployment > System Configuration**.
- 2 Click one of the Nodes for which you want to retrieve a support bundle. The support bundle holds the services logs.
- 3 From the **Actions** menu, select the **Export Support Bundles...** option.
- 4 Select only **VirtualAppliance > Auto Deploy**.
- 5 Click the **Export Support Bundle...** button to download the log files.



Set Up vSphere Auto Deploy and Provision Hosts with vSphere PowerCLI

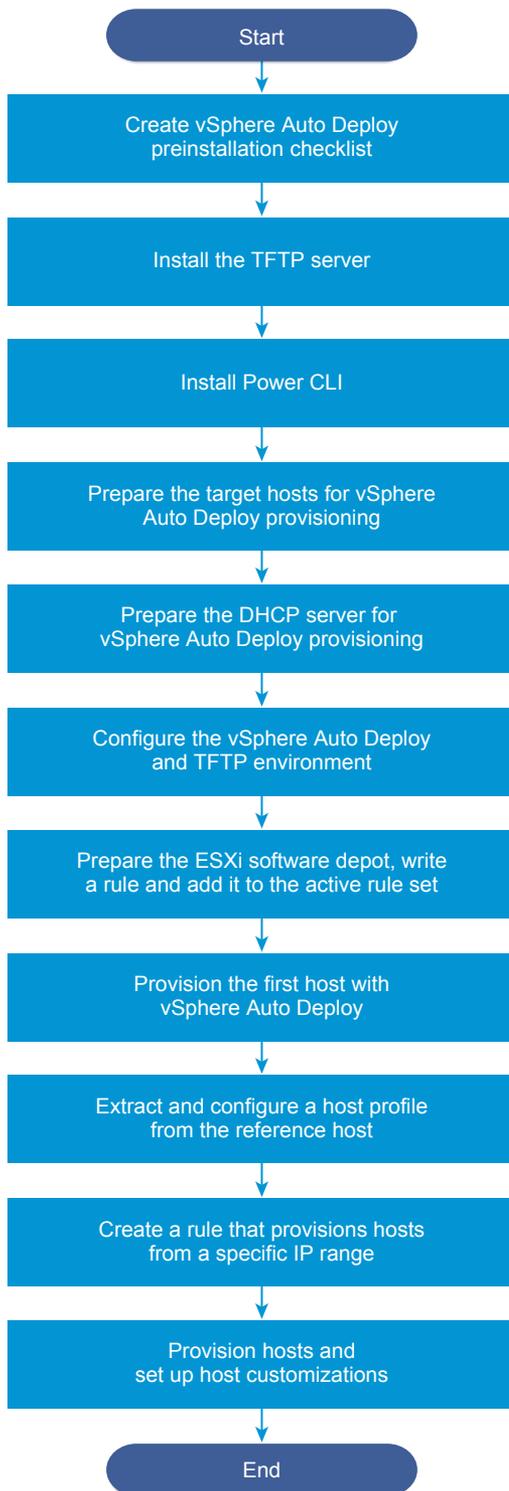
In this scenario you are going to set up and configure a working vSphere Auto Deploy environment that includes four hosts. You will create rules and provision two of the hosts with an image profile and the other two with the same image profile and a host profile that is set up to prompt for user input.

This scenario can provide you with the basis for a production environment. The task descriptions assume that you are using a flat network with no VLAN tagging between the physical hosts and the rest of your environment.

To perform the tasks in this scenario, you should have the following background knowledge and privileges.

- Experience with vSphere (vCenter Server and ESXi).
- Basic knowledge of Microsoft PowerShell and PowerCLI.
- Administrator rights to the target Windows and vCenter Server systems.

Follow the tasks in the order presented in this scenario. Some steps can be performed in a different order, but the order used here limits repeated manipulation of some components.

Figure 2-9. vSphere Auto Deploy Setup and Hosts Provisioning Workflow

vSphere Auto Deploy takes advantage of the iPXE infrastructure and PXE booting with legacy BIOS firmware is possible only over IPv4. If the hosts that you want to provision with vSphere Auto Deploy are with legacy BIOS, the vSphere Auto Deploy server must have an IPv4 address. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.

Procedure

1 vSphere Auto Deploy Preinstallation Checklist

Before you can start the tasks in this vSphere Auto Deploy scenario, make sure that your environment meets the hardware and software requirements, and that you have the necessary permissions for the components included in the setup.

2 Install the TFTP Server

To set up a vSphere Auto Deploy infrastructure, you must install a TFTP server in your environment. vSphere Auto Deploy relies on a TFTP server for sending the boot image to the hosts that it provisions.

3 Install PowerCLI

Before you can manage vSphere Auto Deploy with rules that you create with PowerCLI cmdlets, you must install PowerCLI.

4 Prepare the vSphere Auto Deploy Target Hosts

You must configure the BIOS settings of the four hosts and reconfirm the MAC address of the primary network device to prepare the target hosts for provisioning with vSphere Auto Deploy.

5 Prepare the DHCP Server for vSphere Auto Deploy Provisioning

When you prepare the vSphere Auto Deploy target hosts, you must set up the DHCP server in this scenario to serve each target host with an iPXE binary.

6 Configure the vSphere Auto Deploy and TFTP Environment in the vSphere Web Client

After you prepare the DHCP server, you must start the vSphere Auto Deploy vCenter Server service and configure the TFTP server. You must download a TFTP Boot ZIP file from your vSphere Auto Deploy server. The customized FTP server serves the boot images that vSphere Auto Deploy provides.

7 Prepare the ESXi Software Depot and Write a Rule

After you configure the vSphere Auto Deploy infrastructure, you must add an ESXi software depot, specify an image profile, write a rule, and add it to the active rule set.

8 Provision the First Host with vSphere Auto Deploy

After creating a rule and adding it to the active rule set, you can provision the first host and check its vCenter Server location to complete verification of the image provisioning of your setup.

9 Extract and Configure a Host Profile from the Reference Host

After provisioning the first host, you can extract and configure a host profile that can be used to apply the same configuration to other target hosts. Configuration that differs for different hosts, such as a static IP address, can be managed through the host customization mechanism.

10 Create a Rule that Provisions Hosts from a Specific IP Range

After creating a host profile from a reference host, you can create a rule that applies the previously verified image profile and the host profile that you extracted to target hosts from a specific IP range.

11 Provision Hosts and Set Up Host Customizations

With the rule in place that provisions hosts using an image profile and a host profile, you can provision specific target hosts. If any host profile items are set to prompt the user for input, the host comes up in maintenance mode. You apply the host profile or check host compliance to be prompted for the information. The system associates the host customization with the host.

vSphere Auto Deploy Preinstallation Checklist

Before you can start the tasks in this vSphere Auto Deploy scenario, make sure that your environment meets the hardware and software requirements, and that you have the necessary permissions for the components included in the setup.

This scenario is customized for vCenter Server 6.0 and later. For earlier versions of vCenter Server, go to the corresponding VMware Documentation Center.

For your setup, your system must meet specific software and hardware requirements.

Table 2-18. Preinstallation Checklist

Required Software and Hardware	Details
Operating System	A Windows Server 2008 R2 system or later supported Windows system with Microsoft PowerShell preinstalled. For a full list of supported operating systems, see Supported host Operating Systems for VMware vCenter Server installation .
vCenter Server	vCenter Server version 6.0 or later to be installed on a Windows system. You can also install PowerCLI on a different Windows system. The vSphere Auto Deploy server is part of vCenter Server. You must enable and start the vSphere Auto Deploy service on the vCenter Server system. You can perform many of the setup tasks by logging in to the Windows system, either directly into the console or by using Remote Desktop (RDP). See Prepare Your System for vSphere Auto Deploy .
Storage	At least 4 GB of free space on the Windows system where vCenter Server is running. Preferably a second volume or hard drive. Storage for ESXi datastores NFS, iSCSI, or FibreChannel, with servers and storage arrays that are configured so the servers can detect the LUNs. <ul style="list-style-type: none"> ■ A list of target IP addresses for NFS or iSCSI. ■ A list of target volume information for NFS or iSCSI.
Host information (for four ESXi hosts)	A list of target IP addresses for NFS or iSCSI. A list of target volume information for NFS or iSCSI. <ul style="list-style-type: none"> ■ Default route, net mask, and primary and secondary DNS server IP addresses. ■ IP address and net mask for the VMkernel primary management network. ■ IP address and net mask for other VMkernel networks such as storage, vSphere FT, or VMware vMotion. vSphere Auto Deploy does not overwrite existing partitions by default.

Table 2-18. Preinstallation Checklist (Continued)

Required Software and Hardware	Details
PowerCLI	PowerCLI installer binaries downloaded from the Downloads page on the VMware Web site. See the <i>vSphere PowerCLI User's Guide</i> for detailed instructions for PowerCLI installation.
ESXi software depot	The location of the ESXi software depot on the Downloads page of the VMware Web site. You use a URL to point to the image profile stored at that location, or you download a ZIP file to work with a local depot. Do not download the ESXi image.
TFTP server	TFTP installer software such as WinAgents TFTP server. The TFTP server included in Windows Server 2008 is closely tied to Windows network deployment and is not suitable.
DHCP server	The DHCP server included with Windows Server 2008.
DNS server	A working DNS server. You must add entries in both Forward (A Record) and Reverse (PTR Record) Zones for each target host.

You also need information about and administrator privileges to the core servers of the environment, including the ActiveDirectory server, DNS server, DHCP server, NTP server, and so on.

You must have complete control of the broadcast domain of the subnet in which you will deploy the setup. Ensure that no other DHCP, DNS, or TFTP server are on this subnet.

Install the TFTP Server

To set up a vSphere Auto Deploy infrastructure, you must install a TFTP server in your environment. . vSphere Auto Deploy relies on a TFTP server for sending the boot image to the hosts that it provisions.

This task only installs the TFTP server. You later download a configuration file to the server. See [Configure the vSphere Auto Deploy and TFTP Environment in the vSphere Web Client](#).

Procedure

- 1 Log in to the console of the Windows system on which vCenter Server is installed with administrator privileges, either directly or by using RDP.
- 2 Download and install the TFTP server software.

This sample setup uses the TFTP server from WinAgents. The TFTP server that is included with Windows 2008 is closely tied to Windows network deployment and not suitable for vSphere Auto Deploy.
- 3 Configure the TFTP root directory as D:\Drive or a similar location (for example, D:\TFTP_Root\).

What to do next

Install PowerCLI, to manage vSphere Auto Deploy with PowerCLI cmdlets.

Install PowerCLI

Before you can manage vSphere Auto Deploy with rules that you create with PowerCLI cmdlets, you must install PowerCLI.

In this scenario, you install PowerCLI on the same system as the vCenter Server system. You can also install PowerCLI on a different Windows system.

Prerequisites

- Verify that Microsoft .NET Framework 4.5 or 4.5.x is installed, or install it from the Microsoft Web site.
- Verify that Windows PowerShell 3.0 or 4.0 is installed, or install it from the Microsoft Web site.

Procedure

- 1 Log in with administrator privileges to the console of the Windows system on which vCenter Server is installed, either directly or by using RDP.
- 2 Download PowerCLI from the Download page of the VMware Web site and install the PowerCLI software.
- 3 Confirm that PowerCLI is working.
 - a Start a PowerCLI session.
 - b (Optional) If an SSL error appears, check the thumbprint and ignore the error.
 - c Run the `Get-DeployCommand` cmdlet.

PowerCLI displays a list of cmdlets and their definitions in the PowerCLI window.

- 4 (Optional) If `Get-DeployCommand` does not return the list of cmdlets, check your PowerCLI version and uninstall and reinstall it if necessary.

What to do next

Configure the settings of your target hosts to prepare them for provisioning with vSphere Auto Deploy.

Prepare the vSphere Auto Deploy Target Hosts

You must configure the BIOS settings of the four hosts and reconfirm the MAC address of the primary network device to prepare the target hosts for provisioning with vSphere Auto Deploy.

Prerequisites

Hosts that you want to provision with vSphere Auto Deploy must meet the requirements for ESXi.

See [ESXi Hardware Requirements](#).

Procedure

- 1 Change the BIOS settings of each of the four physical hosts to force the hosts to boot from the primary network device.
- 2 Reconfirm the MAC address of the primary network device.

What to do next

Set up the DHCP server to serve each target host with an iPXE binary.

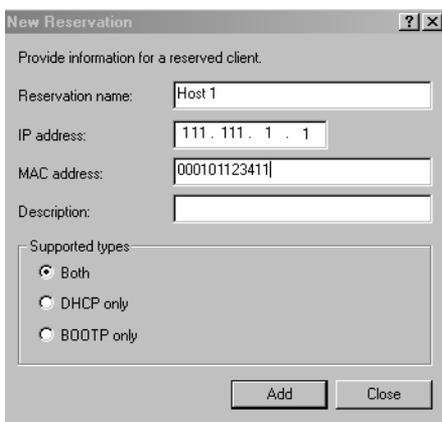
Prepare the DHCP Server for vSphere Auto Deploy Provisioning

When you prepare the vSphere Auto Deploy target hosts, you must set up the DHCP server in this scenario to serve each target host with an iPXE binary.

The environment in this scenario uses Active Directory with DNS and DHCP. The DHCP server is included with Windows 2008.

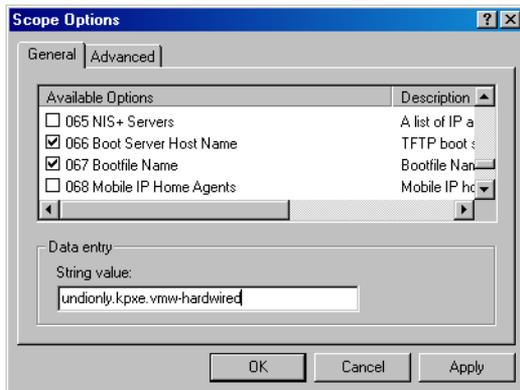
Procedure

- 1 Log in to your DHCP Server with administrator privileges.
- 2 Create a DHCP scope for your IP address range.
 - a Click **Start > Settings > Control Panel > Administrative Tools** and click **DHCP**.
 - b Navigate to **DHCP > hostname > IPv4**.
 - c Right-click **IPv4** and select **New Scope**.
 - d On the Welcome screen, click **Next**, and specify a name and description for the scope.
 - e Specify an IP address range and click **Next**.
 - f Click **Next** until you reach the Configure DHCP Options screen and select **No, I will configure this option later**.
- 3 Create a DHCP reservation for each target ESXi host.
 - a In the DHCP window, navigate to **DHCP > hostname > IPv4 > Autodeploy Scope > Reservations**.
 - b Right-click **Reservations** and select **New Reservation**.
 - c In the New Reservation window, specify a name, IP address, and the MAC address for one of the hosts. Do not include the colon (:) in the MAC address.



- d Repeat the process for each of the other hosts.

- 4 Set up the DHCP Server to point the hosts to the TFTP Server.
 - a In the DHCP window, navigate to **DHCP > hostname > IPv4 > Autodeploy Scope > Scope Options**.
 - b Right click **Scope Options** and choose **Configure Options**.
 - c In the Scope Options window, click the **General** tab.
 - d Click **066 Boot Server Host Name** and enter the address of the TFTP server that you installed in the String value field below the Available Options.



- e Click **067 Bootfile Name** and enter **undionly.kpxe.vmw-hardwired**.
The undionly.kpxe.vmw-hardwired iPXE binary will be used to boot the ESXi hosts.
 - f Click **Apply** and click **OK** to close the window.
- 5 In the DHCP window, right-click **DHCP > hostname > IPv4 > Scope > Activate** and click **Activate**.
- 6 Do not log out from the DHCP Server if you are using Active Directory for DHCP and DNS, or log out otherwise.

What to do next

start the vCenter Server service of vSphere Auto Deploy and configure the TFTP server.

Configure the vSphere Auto Deploy and TFTP Environment in the vSphere Web Client

After you prepare the DHCP server, you must start the vSphere Auto Deploy vCenter Server service and configure the TFTP server. You must download a TFTP Boot ZIP file from your vSphere Auto Deploy server. The customized FTP server serves the boot images that vSphere Auto Deploy provides.

Procedure

- 1 Use the vSphere Web Client to connect to the vCenter Server system that manages the vSphere Auto Deploy server.
- 2 When the Certificate warning appears, continue to the vCenter Server system.

- 3 Start the vSphere Auto Deploy service.
 - a On the vSphere Web Client Home page, click **Administration**.
 - b Under **System Configuration** click **Services**.
 - c Select **Auto Deploy**, click the **Actions** menu, and select **Start**.
 On Windows, the vSphere Auto Deploy service can be disabled. You can enable the service by changing the vSphere Auto Deploy service startup type.
- 4 In the inventory, navigate to the vCenter Server system.
- 5 On the **Manage** tab, select **Settings**, and click **Auto Deploy**.
- 6 Click the **Download TFTP Boot Zip** link to download the TFTP configuration file.
- 7 Save the file `Deploy-tftp.zip` to the `TFTP_Root` directory that you created when you installed the TFTP Server, and unzip the file.

What to do next

Add a software depot to your inventory and use an image profile from the depot to create a rule for host provisioning.

Prepare the ESXi Software Depot and Write a Rule

After you configure the vSphere Auto Deploy infrastructure, you must add an ESXi software depot, specify an image profile, write a rule, and add it to the active rule set.

vSphere Auto Deploy provisions hosts with image profiles that define the set of VIBs that an ESXi installation process uses. Image profiles are stored in software depots. You must make sure the correct image profile is available before you start provisioning hosts. When you add a software depot to a PowerCLI session, it is available only during the current session. It does not persist across sessions.

The steps in this task instruct you to run PowerCLI cmdlets. For additional information about the vSphere Auto Deploy cmdlets that you can run in a PowerCLI session, see [vSphere Auto Deploy PowerCLI Cmdlet Overview](#).

Prerequisites

Verify that you can access the ESXi hosts that you want to provision from the system on which you run PowerCLI.

Procedure

- 1 Log in as an administrator to the console of the Windows system on which vCenter Server is installed, either directly or by using RDP.

This task assumes that you installed PowerCLI on the system on which the vCenter Server system is running.

- 2 In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate issues occur. In a development environment, you can ignore the warning.

- 3 Enter the vCenter Server credentials.
- 4 Run `Add-EsxSoftwareDepot` to add the online depot to the PowerCLI session.

```
Add-EsxSoftwareDepot https://hostupdate.vmware.com/software/VUM/PRODUCTION/main/vmw-depot-index.xml
```

Adding the software depot is required each time you start a new PowerCLI session.

- 5 Validate that you successfully added the software depot by checking the contents of the depot with the `Get-EsxImageProfile` cmdlet.

The cmdlet returns information about all image profiles in the depot.

- 6 Create a new rule by running the `New-DeployRule` cmdlet.

```
New-DeployRule -Name "InitialBootRule" -Item ESXi-6.0.0-2494585-standard -AllHosts
```

The cmdlet creates a rule that assigns the specified image profile to all hosts in the inventory.

- 7 Add the new rule to the active rule set to make the rule available to the vSphere Auto Deploy server.

```
Add-DeployRule -DeployRule "InitialBootRule"
```

What to do next

Provision your first host with vSphere Auto Deploy and verify its image provisioning.

Provision the First Host with vSphere Auto Deploy

After creating a rule and adding it to the active rule set, you can provision the first host and check its vCenter Server location to complete verification of the image provisioning of your setup.

Procedure

- 1 Open a console session to the physical host that you want to use as the first ESXi target host, boot the host, and look for messages that indicate a successful iPXE boot.

During the boot process, DHCP assigns an IP address to the host. The IP address matches the name you specified earlier in the DNS server. The host contacts the vSphere Auto Deploy server and downloads the ESXi binaries from the HTTP URL indicated in the iPXE tramp file that you downloaded earlier to the TFTP_Root directory. Each instance of vSphere Auto Deploy produces a custom set of files for the TFTP Server.

- 2 Use the vSphere Web Client to connect to the vCenter Server system that manages the vSphere Auto Deploy server.

- 3 On the vSphere Web Client Home page, click **Hosts and Clusters**.
- 4 Verify that the newly provisioned host is now in the vCenter Server inventory at the datacenter level.
By default, vSphere Auto Deploy adds hosts at the datacenter level when the boot process completes.

What to do next

Extract a host profile from the host and configure it to require user input.

Extract and Configure a Host Profile from the Reference Host

After provisioning the first host, you can extract and configure a host profile that can be used to apply the same configuration to other target hosts. Configuration that differs for different hosts, such as a static IP address, can be managed through the host customization mechanism.

vSphere Auto Deploy can provision each host with the same host profile. vSphere Auto Deploy can also use host customization that allows you to specify different information for different hosts. For example, if you set up a VMkernel port for vMotion or for storage, you can specify a static IP address for the port by using the host customization mechanism.

Procedure

- 1 Use the vSphere Web Client to connect to the vCenter Server system that manages the vSphere Auto Deploy server.
- 2 Click **Policies and Profiles** and select **Host Profiles**.
- 3 Click **Extract profile from a host**.
- 4 On the Select Host page of the wizard, select the reference host that you configured earlier and click **Next**.
- 5 On the Name and Description page of the wizard, name the profile ESXiGold, add a description, and click **Next**.
- 6 Review the host profile settings and click **Finish**.
- 7 Right-click the new ESXiGold host profile and click **Edit Settings**.
- 8 Leave the name and description unchanged and click **Next**.
- 9 On the Edit host profile page of the wizard, select **Security and Services > Security Settings > Security > User Configuration > root**.
- 10 In **Password policy** on the right panel, select **User Input Password configuration**.
- 11 Click **Finish** to save the host profile settings.

What to do next

Create a vSphere Auto Deploy rule to apply the host profile to other ESXi hosts.

Create a Rule that Provisions Hosts from a Specific IP Range

After creating a host profile from a reference host, you can create a rule that applies the previously verified image profile and the host profile that you extracted to target hosts from a specific IP range.

Procedure

- 1 Log in with administrator privileges to the console of the Windows system on which vCenter Server is installed, either directly or by using RDP.
- 2 In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate issues occur. In a development environment, you can ignore the warning.

- 3 Run `Add-EsxSoftwareDepot` to add the online depot to the PowerCLI session.

```
Add-EsxSoftwareDepot https://hostupdate.vmware.com/software/VUM/PRODUCTION/main/vmw-depot-index.xml
```

Adding the software depot is required each time you start a new PowerCLI session.

- 4 Display the rules in the active rule set by running the `Get-DeployRuleset` cmdlet.
- 5 Create a rule that instructs vSphere Auto Deploy to provision the set of hosts from a specified IP range with the image profile that you previously selected and the host profile that you created from the reference host.

```
New-DeployRule -name "Production01Rule" -item "image_profile", ESXiGold -Pattern "ipv4=IP_range"
```

- 6 Add the new rule to the active rule set.

```
Add-DeployRule -DeployRule "Production01Rule"
```

- 7 Check the active rule set by running the `Get-DeployRuleset` command.

PowerCLI displays information similar to the following example.

```
Name:           Production01Rule
PatternList:    {ipv4=address_range}
ItemList:       {ESXi-version-XXXXXX-standard, Compute01, ESXiGold}
```

What to do next

Provision the hosts and set up the host customizations.

Provision Hosts and Set Up Host Customizations

With the rule in place that provisions hosts using an image profile and a host profile, you can provision specific target hosts. If any host profile items are set to prompt the user for input, the host comes up in maintenance mode. You apply the host profile or check host compliance to be prompted for the information. The system associates the host customization with the host.

Procedure

- 1 Boot the remaining hosts.

vSphere Auto Deploy boots the hosts, applies the host profile, and adds the hosts to the vCenter Server inventory. The hosts remain in maintenance mode because the host profile from the reference host is set up to require user input for each host.

- 2 Use the vSphere Web Client to connect to the vCenter Server system that manages the vSphere Auto Deploy server.
- 3 Click **Policies and Profiles** and select **Host Profiles**.
- 4 Right-click the previously created ESXiGold profile and click **Edit Host Customizations**.
- 5 Enter the required host customizations and save them.
- 6 Apply the host profile to each of the hosts and get the hosts out of maintenance mode. Alternatively, you can reboot each host.

When the reboot progress completes, all hosts are running with the image you specify and use the configuration in the reference host profile. The cluster shows that all hosts are fully compliant.

All hosts are now configured with the shared information through the reference host profile and with the host-specific information through the host customization mechanism. The next time you boot the hosts, they receive the complete Host Profile information, including the host-specific information, and boot up completely configured and out of Maintenance Mode.

Setting Up ESXi

These topics provide information about using the direct console user interface and configuring defaults for ESXi.

ESXi Autoconfiguration

When you turn on the ESXi host for the first time or after resetting the configuration defaults, the host enters an autoconfiguration phase. This phase configures system network and storage devices with default settings.

By default, Dynamic Host Configuration Protocol (DHCP) configures IP, and all visible blank internal disks are formatted with the virtual machine file system (VMFS) so that virtual machines can be stored on the disks.

About the Direct Console ESXi Interface

Use the direct console interface for initial ESXi configuration and troubleshooting.

Connect a keyboard and monitor to the host to use the direct console. After the host completes the autoconfiguration phase, the direct console appears on the monitor. You can examine the default network configuration and change any settings that are not compatible with your network environment.

Key operations available to you in the direct console include:

- Configuring hosts
- Setting up administrative access
- Troubleshooting

You can also use vSphere Web Client to manage the host by using vCenter Server.

Table 2-19. Navigating in the Direct Console

Action	Key
View and change the configuration	F2
Change the user interface to high-contrast mode	F4
Shut down or restart the host	F12
View the VMkernel log	Alt+F12
Switch to the shell console	Alt+F1
Switch to the direct console user interface	Alt+F2
Move the selection between fields	Arrow keys
Select a menu item	Enter
Toggle a value	Spacebar
Confirm sensitive commands, such as resetting configuration defaults	F11
Save and exit	Enter
Exit without saving	Esc
Exit system logs	q

Configure the Keyboard Layout for the Direct Console

You can configure the layout for the keyboard that you use with the direct console.

Procedure

- 1 From the direct console, select **Configure Keyboard** and press Enter.
- 2 Select the layout to use.
- 3 Press the spacebar to toggle selections on and off.
- 4 Press Enter.

Create a Security Banner for the Direct Console

A security banner is a message that is displayed on the direct console **Welcome** screen.

Procedure

- 1 From the vSphere Web Client, connect to the vCenter Server.
- 2 Select the host in the inventory.
- 3 Click the **Manage** tab.
- 4 Click **Settings**.
- 5 Under System, select **Advanced System Settings**.
- 6 Select **Annotations**.
- 7 Click the Edit icon.
- 8 Enter a security message.

The message is displayed on the direct console **Welcome** screen.

Redirecting the Direct Console to a Serial Port

To manage your ESXi host remotely from a serial console, you can redirect the direct console to a serial port.

vSphere supports the VT100 terminal type and the PuTTY terminal emulator to view the direct console over the serial port.

You can redirect the direct console to a serial port in several ways.

- [Redirect the Direct Console to a Serial Port by Setting the Boot Options Manually](#)

When you redirect the direct console to a serial port by setting the boot options, the change does not persist for subsequent boots.

- [Redirect the Direct Console to a Serial Port from the vSphere Web Client](#)

You can manage the ESXi host remotely from a console that is connected to the serial port by redirecting the direct console to either of the serial ports com1 or com2. When you use the vSphere Web Client to redirect the direct console to a serial port, the boot option that you set persists after subsequent reboots.

- [Redirect the Direct Console to a Serial Port in a Host Deployed with Auto Deploy](#)

After you redirect the direct console to a serial port, you can make that setting part of the host profile that persists when you reprovise the host with Auto Deploy.

Redirect the Direct Console to a Serial Port by Setting the Boot Options Manually

When you redirect the direct console to a serial port by setting the boot options, the change does not persist for subsequent boots.

Prerequisites

Verify that the serial port is not in use for serial logging and debugging.

Procedure

- 1 Start the host.
- 2 When the Loading VMware Hypervisor window appears, press Shift+O to edit boot options.
- 3 Disable the logPort and gdbPort on com1 and set tty2Port to com1 by entering the following boot options:

```
"gdbPort=none logPort=none tty2Port=com1";
```

To use com2 instead, replace com1 with com2.

The direct console is redirected to the serial port until you reboot the host. To redirect the direct console for subsequent boots, see [Redirect the Direct Console to a Serial Port from the vSphere Web Client](#)

Redirect the Direct Console to a Serial Port from the vSphere Web Client

You can manage the ESXi host remotely from a console that is connected to the serial port by redirecting the direct console to either of the serial ports com1 or com2. When you use the vSphere Web Client to redirect the direct console to a serial port, the boot option that you set persists after subsequent reboots.

Prerequisites

- Verify that you can access the host from the vSphere Web Client.
- Verify that the serial port is not in use for serial logging and debugging, or for ESX Shell (tty1Port).

Procedure

- 1 From the vSphere Web Client, connect to the vCenter Server.
- 2 Select the host in the inventory.
- 3 Click the **Manage** tab.
- 4 Click **Settings**.
- 5 Under System, select **Advanced System Settings**.
- 6 Make sure that the **VMkernel.Boot.logPort** and **VMkernel.Boot.gdbPort** fields are not set to use the com port that you want to redirect the direct console to.
- 7 Set **VMkernel.Boot.tty2Port** to the serial port to redirect the direct console to: **com1** or **com2**.
- 8 Reboot the host.

You can now manage the ESXi host remotely from a console that is connected to the serial port.

Redirect the Direct Console to a Serial Port in a Host Deployed with Auto Deploy

After you redirect the direct console to a serial port, you can make that setting part of the host profile that persists when you reprovision the host with Auto Deploy.

Prerequisites

The serial port must not already be in use for serial logging and debugging.

Procedure

- 1 From the vSphere Web Client, connect to the vCenter Server.
- 2 Select the host in the inventory.
- 3 Click the **Manage** tab.
- 4 Select **Settings**.
- 5 Select **Advanced System Settings**.
- 6 Make sure that the **VMkernel.Boot.logPort** and **VMkernel.Boot.gdbPort** fields are not set to use the com port that you want to redirect the direct console to.
- 7 Set **VMkernel.Boot.tty2Port** to the serial port to redirect the direct console to: **com1** or **com2**.
- 8 Click **OK**.
- 9 Save the host profile and attach the host to the profile. See the *vSphere Host Profiles* documentation.

The setting to redirect the direct console to a serial port is stored by vCenter Server and persists when you reprovision the host with Auto Deploy.

Enable ESXi Shell and SSH Access with the Direct Console User Interface

Use the direct console user interface to enable the ESXi Shell.

Procedure

- 1 From the Direct Console User Interface, press F2 to access the System Customization menu.
- 2 Select **Troubleshooting Options** and press Enter.
- 3 From the Troubleshooting Mode Options menu, select a service to enable.
 - Enable ESXi Shell
 - Enable SSH
- 4 Press Enter to enable the service.
- 5 (Optional) Set the timeout for the ESXi Shell.

By default, timeouts for the ESXi Shell is 0 (disabled).

The availability timeout setting is the number of minutes that can elapse before you must log in after the ESXi Shell is enabled. After the timeout period, if you have not logged in, the shell is disabled.

Note If you are logged in when the timeout period elapses, your session will persist. However, the ESXi Shell will be disabled, preventing other users from logging in.

a From the Troubleshooting Mode Options menu, select **Modify ESXi Shell and SSH timeouts** and press Enter.

b Enter the availability timeout in minutes.

The availability timeout is the number of minutes that can elapse before you must log in after the ESXi Shell is enabled.

c Press Enter.

d Enter the idle timeout.

The idle timeout is the number of minutes that can elapse before the user is logged out of an idle interactive sessions. Changes to the idle timeout apply the next time a user logs in to the ESXi Shell and do not affect existing sessions.

6 Press Esc until you return to the main menu of the Direct Console User Interface.

Managing ESXi Remotely

You can use the VMware Host Client, the vSphere Web Client and vCenter Server to manage your ESXi hosts.

For instructions about downloading and installing vCenter Server and the vCenter Server components or for downloading and deploying the vCenter Server Appliance, see [Chapter 4 Installing vCenter Server and Platform Services Controller on Windows](#) and [Chapter 3 Deploying the vCenter Server Appliance and Platform Services Controller Appliance](#). For information about installing the VMware Host Client, see *vSphere Single Host Management*.

Set the Password for the Administrator Account

You can use the direct console to set the password for the administrator account (root).

The administrative user name for the ESXi host is root. By default, the administrative password is not set.

Procedure

- 1 From the direct console, select **Configure Password**.
- 2 (Optional) If a password is already set up, type the password in the **Old Password** line and press Enter.
- 3 In the **New Password** line, type a new password and press Enter.
- 4 Retype the new password and press Enter.

Configuring the BIOS Boot Settings

If your server has multiple drives, you might need to configure the BIOS settings.

The BIOS boot configuration determines how your server boots. Generally, the CD-ROM device is listed first.

Note If you are using ESXi Embedded, the BIOS boot configuration determines whether your server boots into the ESXi boot device or another boot device. Generally, the USB flash device is listed first in the BIOS boot settings on the machine that hosts ESXi.

You can change the boot setting by configuring the boot order in the BIOS during startup or by selecting a boot device from the boot device selection menu. When you change the boot order in the BIOS, the new setting affects all subsequent reboots. When you select a boot device from the boot device selection menu, the selection affects the current boot only.

Some servers do not have a boot device selection menu, in which case you must change the boot order in the BIOS even for one-time boots, and then change it back again during a subsequent reboot.

Change the BIOS Boot Setting for ESXi

Configure the BIOS boot setting for ESXi if you want the server to boot into ESXi by default.

ESXi Installable and ESXi Embedded cannot exist on the same host.

Procedure

- 1 While the ESXi host is powering on, press the key required to enter your host's BIOS setup.

Depending on your server hardware, the key might be a function key or Delete. The option to enter the BIOS setup might be different for your server.

- 2 Select the BIOS boot setting.

Option	Description
If you are using the installable version of ESXi	Select the disk on which you installed the ESXi software and move it to the first position in the list. The host boots into ESXi.
If you are using ESXi Embedded	Select the USB flash device and move it to the first position in the list. The host starts in ESXi mode.

Configure the Boot Setting for Virtual Media

If you are using remote management software to set up ESXi, you might need to configure the boot setting for virtual media.

Virtual media is a method of connecting a remote storage media such as CD-ROM, USB mass storage, ISO image, and floppy disk to a target server that can be anywhere on the network. The target server has access to the remote media, and can read from and write to it as if it were physically connected to the server's USB port.

Prerequisites

ESXi Installable and ESXi Embedded cannot exist on the same host.

Procedure

- 1 Connect the media to the virtual device.

For example, if you are using a Dell server, log in to the Dell Remote Access Controller (DRAC) or a similar remote management interface and select a physical floppy or CD-ROM drive, or provide a path to a floppy image or CD-ROM image.

- 2 Reboot the server.
- 3 While the server is powering on, enter the device selection menu.

Depending on your server hardware, the key might be a function key or Delete.

- 4 Follow the instructions to select the virtual device.

The server boots from the configured device once and goes back to the default boot order for subsequent boots.

Configuring Network Settings

ESXi requires one IP address for the management network. To configure basic network settings, use the vSphere Web Client or the direct console.

Use the vSphere Web Client if you are satisfied with the IP address assigned by the DHCP server.

Use the direct console for network configuration in the following cases:

- You are not satisfied with the IP address assigned by the DHCP server.
- You are not allowed to use the IP address assigned by the DHCP server.
- ESXi does not have an IP address. This situation might occur if the autoconfiguration phase did not succeed in configuring DHCP.
- The wrong network adapter was selected during the autoconfiguration phase.

Network Access to Your ESXi Host

The default behavior is to configure the ESXi management network using DHCP. You can override the default behavior and use static IP settings for the management network after the installation is completed.

Table 2-20. Network Configuration Scenarios Supported by ESXi

Scenario	Approach
You want to accept the DHCP-configured IP settings.	In the ESXi direct console, you can find the IP address assigned through DHCP to the ESXi management interface. You can use that IP address to connect to the host from the vSphere Web Client and customize settings, including changing the management IP address.
One of the following is true: <ul style="list-style-type: none"> ■ You do not have a DHCP server. ■ The ESXi host is not connected to a DHCP server. ■ Your connected DHCP server is not functioning properly. 	<p>During the autoconfiguration phase, the software assigns the link local IP address, which is in the subnet 169.254.x.x/16. The assigned IP address appears on the direct console.</p> <p>You can override the link local IP address by configuring a static IP address using the direct console.</p>
The ESXi host is connected to a functioning DHCP server, but you do not want to use the DHCP-configured IP address.	<p>During the autoconfiguration phase, the software assigns a DHCP-configured IP address.</p> <p>You can make the initial connection by using the DHCP-configured IP address. Then you can configure a static IP address.</p> <p>If you have physical access to the ESXi host, you can override the DHCP-configured IP address by configuring a static IP address using the direct console.</p>
Your security deployment policies do not permit unconfigured hosts to be powered on the network.	Follow the setup procedure in Configure the Network Settings on a Host That Is Not Attached to the Network .

ESXi Networking Security Recommendations

Isolation of network traffic is essential to a secure ESXi environment. Different networks require different access and level of isolation.

Your ESXi host uses several networks. Use appropriate security measures for each network, and isolate traffic for specific applications and functions. For example, ensure that VMware vSphere vMotion[®] traffic does not travel over networks where virtual machines are located. Isolation prevents snooping. Having separate networks is also recommended for performance reasons.

- vSphere infrastructure networks are used for features such as vSphere vMotion, VMware vSphere Fault Tolerance, and storage. Isolate these networks for their specific functions. It is often not necessary to route these networks outside a single physical server rack.
- A management network isolates client traffic, command-line interface (CLI) or API traffic, and third-party software traffic from other traffic. This network should be accessible only by system, network, and security administrators. Use jump box or virtual private network (VPN) to secure access to the management network. Strictly control access within this network.

- Virtual machine traffic can flow over one or many networks. You can enhance the isolation of virtual machines by using virtual firewall solutions that set firewall rules at the virtual network controller. These settings travel with a virtual machine as it migrates from host to host within your vSphere environment.

Choose Network Adapters for the Management Network

Traffic between an ESXi host and any external management software is transmitted through an Ethernet network adapter on the host. You can use the direct console to choose the network adapters that are used by the management network.

Examples of external management software include the vCenter Server and SNMP client. Network adapters on the host are named `vmnicN`, where N is a unique number identifying the network adapter, for example, `vmnic0`, `vmnic1`, and so forth.

During the autoconfiguration phase, the ESXi host chooses `vmnic0` for management traffic. You can override the default choice by manually choosing the network adapter that carries management traffic for the host. In some cases, you might want to use a Gigabit Ethernet network adapter for your management traffic. Another way to help ensure availability is to select multiple network adapters. Using multiple network adapters enables load balancing and failover capabilities.

Procedure

- 1 From the direct console, select **Configure Management Network** and press Enter.
- 2 Select **Network Adapters** and press Enter.
- 3 Select a network adapter and press Enter.

After the network is functional, you can use the vSphere Web Client to connect to the ESXi host through vCenter Server.

Set the VLAN ID

You can set the virtual LAN (VLAN) ID number of the ESXi host.

Procedure

- 1 From the direct console, select **Configure Management Network** and press Enter.
- 2 Select **VLAN** and press Enter.
- 3 Enter a VLAN ID number from 1 through 4094.

Configuring IP Settings for ESXi

By default, DHCP sets the IP address, subnet mask, and default gateway.

For future reference, write down the IP address.

For DHCP to work, your network environment must have a DHCP server. If DHCP is not available, the host assigns the link local IP address, which is in the subnet 169.254.x.x/16. The assigned IP address appears on the direct console. If you do not have physical monitor access to the host, you can access the direct console using a remote management application. See [Using Remote Management Applications](#)

When you have access to the direct console, you can optionally configure a static network address. The default subnet mask is 255.255.0.0.

Configure IP Settings from the Direct Console

If you have physical access to the host or remote access to the direct console, you can use the direct console to configure the IP address, subnet mask, and default gateway.

Procedure

- 1 Select **Configure Management Network** and press Enter.
- 2 Select **IP Configuration** and press Enter.
- 3 Select **Set static IP address and network configuration**.
- 4 Enter the IP address, subnet mask, and default gateway and press Enter.

Configure IP Settings from the vSphere Web Client

If you do not have physical access to the host, you can use the vSphere Web Client to configure static IP settings.

Procedure

- 1 Log in to the vCenter Server from the vSphere Web Client.
- 2 Select the host in the inventory.
- 3 On the **Configure** tab, expand **Networking**.
- 4 Select **Virtual adapters**.
- 5 Select **vmk0 Management Network** and click the edit icon.
- 6 Select **IPv4 settings**.
- 7 Select **Use static IPv4 settings**.
- 8 Enter or change the static IPv4 address settings.
- 9 (Optional) Set static IPv6 addresses.
 - a Select **IPv6 settings**.
 - b Select **Static IPv6 addresses**.
 - c Click the add icon.
 - d Type the IPv6 address and click **OK**.
- 10 Click **OK**.

Configuring DNS for ESXi

You can select either manual or automatic DNS configuration of the ESXi host.

The default is automatic. For automatic DNS to work, your network environment must have a DHCP server and a DNS server.

In network environments where automatic DNS is not available or not desirable, you can configure static DNS information, including a host name, a primary name server, a secondary name server, and DNS suffixes.

Configure DNS Settings from the Direct Console

If you have physical access to the host or remote access to the direct console, you can use the direct console to configure DNS information.

Procedure

- 1 Select **Configure Management Network** and press Enter.
- 2 Select **DNS Configuration** and press Enter.
- 3 Select **Use the following DNS server addresses and hostname**.
- 4 Enter the primary server, an alternative server (optional), and the host name.

Configure DNS Suffixes

If you have physical access to the host, you can use the direct console to configure DNS information. By default, DHCP acquires the DNS suffixes.

Procedure

- 1 From the direct console, select **Configure Management Network**.
- 2 Select **Custom DNS Suffixes** and press Enter.
- 3 Enter new DNS suffixes.

Configure the Network Settings on a Host That Is Not Attached to the Network

Some highly secure environments do not permit unconfigured hosts on the network to be powered on. You can configure the host before you attach the host to the network.

Prerequisites

Verify that no network cables are connected to the host.

Procedure

- 1 Power on the host.
- 2 Use the direct console user interface to configure the password for the administrator account (root).
- 3 Use the direct console user interface to configure a static IP address.

- 4 Connect a network cable to the host.
- 5 (Optional) Use the vSphere Web Client to connect to a vCenter Server system.
- 6 (Optional) Add the host to the vCenter Server inventory.

Test the Management Network

You can use the direct console to do simple network connectivity tests.

The direct console performs the following tests.

- Pings the default gateway
- Pings the primary DNS name server
- Pings the secondary DNS nameserver
- Resolves the configured host name

Procedure

- 1 From the direct console, select **Test Management Network** and press Enter.
- 2 Press Enter to start the test.

Restart the Management Agents

The management agents synchronize VMware components and let you access the ESXi host by using the vSphere Web Client and vCenter Server. They are installed with the vSphere software. You might need to restart the management agents if remote access is interrupted.

Restarting the management agents restarts all management agents and services that are installed and running in `/etc/init.d` on the ESXi host. Typically, these agents include `hostd`, `ntpd`, `sfcdbd`, `sldap`, `wsman`, and `vobd`. The software also restarts the Fault Domain Manager (FDM) if installed.

Users accessing this host by using the vSphere Web Client and vCenter Server lose connectivity when you restart management agents.

Procedure

- 1 From the direct console, select **Troubleshooting Options** and press Enter.
- 2 Select **Restart Management Agents** and press Enter.
- 3 Press F11 to confirm the restart.

The ESXi host restarts the management agents and services.

Restart the Management Network

Restarting the management network interface might be required to restore networking or to renew a DHCP lease.

Restarting the management network will result in a brief network outage that might temporarily affect running virtual machines.

If a renewed DHCP lease results in a new network identity (IP address or host name), remote management software will be disconnected.

Procedure

- 1 From the direct console, select **Restart Management Network** and press Enter.
- 2 Press F11 to confirm the restart.

Test Connectivity to Devices and Networks

You can use the direct console to perform some simple network connectivity tests. In addition to the management network, you can specify other devices and networks.

Procedure

- 1 From the direct console, select **Test Management Network** and press Enter.
- 2 Type addresses to ping or another DNS host name to resolve.
- 3 Press Enter to start the test.

Restoring the Standard Switch

A vSphere Distributed Switch functions as a single virtual switch across all associated hosts. Virtual machines can maintain a consistent network configuration as they migrate across multiple hosts. If you migrate an existing standard switch, or virtual adapter, to a Distributed Switch and the Distributed Switch becomes unnecessary or stops functioning, you can restore the standard switch to ensure that the host remains accessible.

When you restore the standard switch, a new virtual adapter is created and the management network uplink that is currently connected to Distributed Switch is migrated to the new virtual switch.

You might need to restore the standard switch for the following reasons:

- The Distributed Switch is not needed or is not functioning.
- The Distributed Switch needs to be repaired to restore connectivity to vCenter Server and the hosts need to remain accessible.
- You do not want vCenter Server to manage the host. When the host is not connected to vCenter Server, most Distributed Switch features are unavailable to the host.

Prerequisites

Verify that your management network is connected to a distributed switch.

Procedure

- 1 From the direct console, select **Restore Standard Switch** and press Enter.
If the host is on a standard switch, this selection is dimmed, and you cannot select it.
- 2 Press F11 to confirm.

Storage Behavior

When you start ESXi, the host enters an autoconfiguration phase during which system storage devices are configured with defaults.

When you reboot the ESXi host after installing the ESXi image, the host configures the system storage devices with default settings. By default, all visible blank internal disks are formatted with VMFS, so you can store virtual machines on the disks. In ESXi Embedded, all visible blank internal disks with VMFS are also formatted by default.

Caution ESXi overwrites any disks that appear to be blank. Disks are considered to be blank if they do not have a valid partition table or partitions. If you are using software that uses such disks, in particular if you are using logical volume manager (LVM) instead of, or in addition to, conventional partitioning schemes, ESXi might cause local LVM to be reformatted. Back up your system data before you power on ESXi for the first time.

On the hard drive or USB device that the ESXi host is booting from, the disk-formatting software retains existing diagnostic partitions that the hardware vendor creates. In the remaining space, the software creates the partitions described in [Table 2-21](#).

Table 2-21. Partitions Created by ESXi on the Host Drive

ESXi Version	Partitions Created
ESXi Installable	<p>For fresh installations, several new partitions are created for the boot banks, the scratch partition, and the locker. Fresh ESXi installations use GUID Partition Tables (GPT) instead of MSDOS-based partitioning. The partition table itself is fixed as part of the binary image, and is written to the disk at the time the system is installed. The ESXi installer leaves the scratch and VMFS partitions blank and ESXi creates them when the host is rebooted for the first time after installation or upgrade. One 4GB VFAT scratch partition is created for system swap. See About the Scratch Partition. The VFAT scratch partition is created only on the disk from which the ESXi host is booting.</p> <hr/> <p>Note To create the VMFS volume and a scratch partition with the installation, the ESXi installer requires a minimum of 5.2GB of free space on the installation disk.</p> <hr/> <p>The installer affects only the installation disk. The installer does not affect other disks of the server. When you install on a disk, the installer overwrites the entire disk. When the installer autoconfigures storage, the installer does not overwrite hardware vendor partitions. During ESXi installation, the installer creates a 110MB diagnostic partition for core dumps.</p>
ESXi Embedded	<p>One 110MB diagnostic partition for core dumps, if this partition is not present on another disk. The VFAT scratch and diagnostic partitions are created only on the disk from which the ESXi host is booting. On other disks, the software creates one VMFS5 partition per blank disk, using the whole disk. Only blank disks are formatted.</p>
Both ESXi Installable and ESXi Embedded	<p>One VMFS5 partition on the remaining free space.</p>

You might want to override this default behavior if, for example, you use shared storage devices instead of local storage. To prevent automatic disk formatting, detach the local storage devices from the host under the following circumstances:

- Before you start the host for the first time.
- Before you start the host after you reset the host to the configuration defaults.

To override the VMFS formatting if automatic disk formatting already occurred, you can remove the datastore. See the *vCenter Server and Host Management* documentation.

About the Scratch Partition

For new installations of ESXi, during the autoconfiguration phase, a 4-GB VFAT scratch partition is created if the partition is not present on another disk.

Note Partitioning for hosts that are upgraded to ESXi 5.x from ESXi versions earlier than version 5.0 differs significantly from partitioning for new installations of ESXi 5.x. See the *vSphere Upgrade* documentation.

When ESXi boots, the system tries to find a suitable partition on a local disk to create a scratch partition.

The scratch partition is not required. It is used to store vm-support output, which you need when you create a support bundle. If the scratch partition is not present, vm-support output is stored in a ramdisk. In low-memory situations, you might want to create a scratch partition if one is not present.

For the installable version of ESXi, the partition is created during installation and is selected. VMware recommends that you do not modify the partition.

Note To create the VMFS volume and scratch partition, the ESXi installer requires a minimum of 5.2 GB of free space on the installation disk.

For ESXi Embedded, if a partition is not found, but an empty local disk exists, the system formats it and creates a scratch partition. If no scratch partition is created, you can configure one, but a scratch partition is not required. You can also override the default configuration. You might want to create the scratch partition on a remote NFS-mounted directory.

Note The installer can create multiple VFAT partitions. The VFAT designation does not always indicate that the partition is a scratch partition. Sometimes, a VFAT partition can simply lie idle.

Set the Scratch Partition from the vSphere Web Client

If a scratch partition is not set up, you might want to configure one, especially if low memory is a concern. When a scratch partition is not present, vm-support output is stored in a ramdisk.

Prerequisites

The directory to use for the scratch partition must exist on the host.

Procedure

- 1 From the vSphere Web Client, connect to the vCenter Server.
- 2 Select the host in the inventory.
- 3 Click the **Manage** tab.
- 4 Select **Settings**.
- 5 Select **Advanced System Settings**.

The setting **ScratchConfig.CurrentScratchLocation** shows the current location of the scratch partition.

- 6 In the **ScratchConfig.ConfiguredScratchLocation** text box, enter a directory path that is unique for this host.

For example, `/vmfs/volumes/DatastoreUUID/DatastoreFolder`.

- 7 Reboot the host for the changes to take effect.

Configuring System Logging

The ESXi hosts run the syslog service (`vm syslogd`), which writes messages from the VMkernel and other system components to log files.

You can configure the amount and location of the log. You can also create and apply log filters to modify the logging policy of an ESXi host.

Configure Syslog on ESXi Hosts

You can use the vSphere Web Client or the `esxcli system syslog vCLI` command to configure the syslog service.

For information about using the `esxcli system syslog` command and other vCLI commands, see *Getting Started with vSphere Command-Line Interfaces*.

Procedure

- 1 In the vSphere Web Client inventory, select the host.
- 2 Click **Configure**.
- 3 Under System, click **Advanced System Settings**.
- 4 Filter for **syslog**.
- 5 To set up logging globally, select the setting to change and click **Edit**.

Option	Description
Syslog.global.defaultRotate	Maximum number of archives to keep. You can set this number globally and for individual subloggers.
Syslog.global.defaultSize	Default size of the log, in KB, before the system rotates logs. You can set this number globally and for individual subloggers.
Syslog.global.LogDir	Directory where logs are stored. The directory can be located on mounted NFS or VMFS volumes. Only the <code>/scratch</code> directory on the local file system is persistent across reboots. Specify the directory as <code>[datastorename] path_to_file</code> , where the path is relative to the root of the volume backing the datastore. For example, the path <code>[storage1] /systemlogs</code> maps to the path <code>/vmfs/volumes/storage1/systemlogs</code> .
Syslog.global.logDirUnique	Selecting this option creates a subdirectory with the name of the ESXi host under the directory specified by Syslog.global.LogDir . A unique directory is useful if the same NFS directory is used by multiple ESXi hosts.
Syslog.global.LogHost	Remote host to which syslog messages are forwarded and port on which the remote host receives syslog messages. You can include the protocol and the port, for example, <code>ssl://hostName1:1514</code> . UDP (default), TCP, and SSL are supported. The remote host must have syslog installed and correctly configured to receive the forwarded syslog messages. See the documentation for the syslog service installed on the remote host for information on configuration.

- 6 (Optional) To overwrite the default log size and log rotation for any of the logs.
 - a Click the name of the log that you want to customize.
 - b Click **Edit** and enter the number of rotations and the log size you want.
- 7 Click **OK**.

Changes to the syslog options take effect immediately.

Configure Log Filtering on ESXi Hosts

The log filtering capability lets you modify the logging policy of the syslog service that is running on an ESXi host. You can create log filters to reduce the number of repetitive entries in the ESXi logs and to blacklist specific log events entirely.

Log filters affect all log events that are processed by the ESXi host `vm syslogd` daemon, whether they are recorded to a log directory or to a remote syslog server.

When you create a log filter, you set a maximum number of log entries for the log messages. The logs messages are generated by one or more specified system components and that match a specified phrase. You must enable the log filtering capability and reload the syslog daemon to activate the log filters on the ESXi host.

Important Setting a limit to the amount of logging information, restricts your ability to troubleshoot potential system failures properly. If a log rotate occurs after the maximum number of log entries is reached, you might lose all instances of a filtered message.

Procedure

- 1 Log in to the ESXi Shell as root.
- 2 In the `/etc/vmware/logfilters` file, add the following entry to create a log filter.

```
numLogs | ident | logRegexp
```

where:

- *numLogs* sets the maximum number of log entries for the specified log messages. After reaching this number, the specified log messages are filtered and ignored. Use `0` to filter and ignore all the specified log messages.
- *ident* specifies one or more system components to apply the filter to the log messages that these components generate. For information about the system components that generate log messages, see the values of the `idents` parameters in the syslog configuration files. The files are located in the `/etc/vmsyslog.conf.d` directory. Use a comma-separated list to apply a filter to more than one system component. Use `*` to apply a filter to all system components.

- *logRegexp* specifies a case-sensitive phrase with Python regular expression syntax to filter the log messages by their content.

For example, to set a limit of maximum two-log entries from the `hostd` component for messages that resemble the `SOCKET connect failed, error 2: No such file or directory` phrase with any error number, add the following entry:

```
2 | hostd | SOCKET connect failed, error .*: No such file or directory
```

Note A line starting with `#` denotes a comment and the rest of the line is ignored.

- 3 In the `/etc/vmsyslog.conf` file, add the following entry to enable the log filtering capability.

```
enable_logfilters = true
```

- 4 Run the `esxcli system syslog reload` command to reload the syslog daemon and apply the configuration changes.

Set the Host Image Profile Acceptance Level

The Host Image Profile acceptance level determines which vSphere installation bundles (VIBs) are accepted for installation.

VIB signatures are checked and accepted for installation based on a combination of the VIB acceptance level and the host image profile acceptance level. VIBs are tagged with an acceptance level that depends on their signature status.

See [Acceptance Levels](#).

Prerequisites

Required privileges: **Host.Configuration.SecurityProfile** and **Host.Configuration.Firewall**

Procedure

- 1 From the vSphere Web Client, connect to the vCenter Server.
- 2 Select the host in the inventory.
- 3 Click the **Manage** tab.
- 4 Click **Settings**.
- 5 Under System, select **Security Profile**.
- 6 Scroll down to Host Image Profile Acceptance Level, and click **Edit**.

- 7 Select the acceptance level and click **OK**.

Table 2-22. Host Image Profile Acceptance Levels

Host Image Profile Acceptance Level	Accepted Levels of VIBs
VMware Certified	VMware Certified
VMware Accepted	VMware Certified, VMware Accepted
Partner Supported	VMware Certified, VMware Accepted, Partner Supported
Community Supported	VMware Certified, VMware Accepted, Partner Supported, Community Supported

Remove All Custom Packages on ESXi

After adding custom packages, you might decide to remove them.

Prerequisites

Before you remove custom packages, shut down or migrate running virtual machines off of the ESXi host.

Procedure

- 1 Reboot the ESXi host.
- 2 In the direct console, select **Remove Custom Extensions** and press F11 to confirm.
- 3 Reboot the host.

All custom packages are removed.

Disable Support for Non-ASCII Characters in Virtual Machine File and Directory Names

By default, ESXi supports the use of non-ASCII characters for virtual machine file and directory names. You can disable this support by modifying the `/etc/vmware/hostd/config.xml` file.

After you disable this support, you can still enter non-ASCII characters for virtual machine names. vSphere user interfaces display the virtual machine names in the non-ASCII characters, but ESXi converts the actual file and directory names to ASCII strings.

Procedure

- 1 Using a text editor, open the `/etc/vmware/hostd/config.xml` file for the ESXi host.
- 2 Within the `<config></config>` tag, add the following code.


```
<g11nSupport>false</g11nSupport>
```
- 3 Save and close the file.
- 4 Reboot the host.

Reset the System Configuration

If you are having trouble determining the source of a problem with your ESXi host, you can reset the system configuration.

Changes in the system configuration can be related to various problems, including problems with connectivity to the network and devices. Resetting the system configuration might solve such problems. If resetting the system configuration does not solve the problem, it can still rule out configuration changes made since the initial setup as the source of the problem.

When you reset the configuration, the software overrides all your configuration changes, deletes the password for the administrator account (root), and reboots the host. Configuration changes made by your hardware vendor, such as IP address settings and license configuration, might also be deleted.

Resetting the configuration does not remove virtual machines on the ESXi host. After you reset the configuration defaults, the virtual machines are not visible, but you make them visible again by reconfiguring storage and reregistering the virtual machines.

Caution When you reset the configuration defaults, users accessing the host lose connectivity.

Prerequisites

Before resetting the configuration, back up your ESXi configuration in case you want to restore your configuration.

Procedure

- 1 Back up the configuration using the vSphere CLI `vicfg-cfgbackup` command.
- 2 From the direct console, select **Reset System Configuration** and press Enter.
- 3 Press F11 to confirm.

The system reboots after all settings are reset to the default values.

After You Install and Set Up ESXi

After ESXi is installed and set up, you can manage the host by using the vSphere Web Client and vCenter Server, license the host, and back up your ESXi configuration.

You can also use the VMware Host Client to connect directly to the ESXi host and to manage it. For information about installing and using the VMware Host Client, see *vSphere Single Host Management*.

Managing the ESXi Host

The VMware Host Client provides the simplest way to manage your ESXi host and operate its virtual machines.

You can also use the vSphere Web Client to connect to and manage vCenter Server by using a Web browser. The vSphere Web Client is installed together with vCenter Server and the vCenter Server Appliance and you can use it to manage your ESXi hosts.

Licensing ESXi Hosts

After you install ESXi, it has a 60-day evaluation period during which you can explore the full set of vSphere features provided with a vSphere Enterprise Plus license. You must assign the host an appropriate license before the evaluation period expires.

ESXi hosts are licensed with vSphere licenses that have per-CPU capacity. To license hosts correctly, you must assign them a vSphere license that has enough CPU capacity to cover all CPUs in the hosts. The license must support all features that the hosts are using. For example, if the hosts are connected to a vSphere Distributed Switch, you must assign a license that has the vSphere Distributed Switch feature.

You can use one of following methods to license ESXi hosts:

- License multiple hosts at a time by using the license management function in the vSphere Web Client. The hosts must be connected to a vCenter Server system. For more information, see *vCenter Server and Host Management*.
- Set up bulk licensing by using PowerCLI commands. Bulk licensing works for all ESXi hosts, but is especially useful for hosts provisioned with Auto Deploy. See [Set Up Bulk Licensing](#)
- License individual ESXi hosts by using a direct connection with the VMware Host Client. For information about assigning a license key to an ESXi host, see *vSphere Single Host Management*.

About ESXi Evaluation and Licensed Modes

You can use evaluation mode to explore the entire set of features for ESXi hosts. The evaluation mode provides the set of features equal to a vSphere Enterprise Plus license. Before the evaluation mode expires, you must assign to your hosts a license that supports all the features in use.

For example, in evaluation mode, you can use vSphere vMotion technology, the vSphere HA feature, the vSphere DRS feature, and other features. If you want to continue using these features, you must assign a license that supports them.

The installable version of ESXi hosts is always installed in evaluation mode. ESXi Embedded is preinstalled on an internal storage device by your hardware vendor. It might be in evaluation mode or prelicensed.

The evaluation period is 60 days and begins when you turn on the ESXi host. At any time during the 60-day evaluation period, you can convert from licensed mode to evaluation mode. The time available in the evaluation period is decreased by the time already used.

For example, suppose that you use an ESXi host in evaluation mode for 20 days and then assign a vSphere Standard Edition license key to the host. If you set the host back in evaluation mode, you can explore the entire set of features for the host for the remaining evaluation period of 40 days.

For information about managing licensing for ESXi hosts, see the *vCenter Server and Host Management* documentation.

Recording the License Key of an ESXi Host

If a host becomes inaccessible or unbootable, you should have a record of its license key. You can write down the license key and tape it to the server, or put the license key in a secure location. You can access the license key from the direct console user interface or the vSphere Web Client.

View the License Keys of ESXi Hosts from the vSphere Web Client

You can view the license keys of the hosts that are connected to a vCenter Server system through the vSphere Web Client.

Procedure

- 1 In the vSphere Web Client, select **Administration**.
- 2 Under Licensing, select **Licenses**.
- 3 On the **Assets** tab, select **Hosts**.
- 4 In the License column, click a license.

You view information about the license, such as its usage and license key.

Access the ESXi License Key from the Direct Console

If you have physical access to the host or remote access to the direct console, you can use the direct console to access the ESXi license key.

Procedure

- ◆ From the direct console, select **View Support Information**.

The license key appears in the form XXXXX-XXXXX-XXXXX-XXXXX-XXXXX, labeled License Serial Number.

Note The physical machine serial number also appears, labeled Serial Number. Do not confuse the license key with the physical machine serial number.

View System Logs

System logs provide detailed information about system operational events.

Procedure

- 1 From the direct console, select **View System Logs**.
- 2 Press a corresponding number key to view a log.
vCenter Server Agent (vpxa) logs appear if you add the host to vCenter Server.
- 3 Press Enter or the spacebar to scroll through the messages.

4 Perform a regular expression search.

a Press the slash key (/).

b Type the text to find.

c Press Enter.

The found text is highlighted on the screen.

5 Press q to return to the direct console.

What to do next

See also [Configure Syslog on ESXi Hosts](#).

Deploying the vCenter Server Appliance and Platform Services Controller Appliance

3

You can deploy the vCenter Server Appliance with an embedded or external Platform Services Controller to manage your vSphere environment. You can deploy a Platform Services Controller appliance and register external deployments and Windows installations of vCenter Server Appliance with this Platform Services Controller appliance.

You can deploy the vCenter Server Appliance or Platform Services Controller appliance on an ESXi host 5.5 or later, or on an ESXi host or DRS cluster from the inventory of a vCenter Server instance 5.5 or later.

For information about the software included in the vCenter Server Appliance 6.5, see [Overview of the vCenter Server Appliance](#).

For information about the software and hardware requirements for deploying the vCenter Server Appliance and Platform Services Controller appliance, see [System Requirements for the vCenter Server Appliance and Platform Services Controller Appliance](#).

The vCenter Server Appliance installer contains executable files for GUI and CLI deployments, which you can use alternatively.

- The GUI deployment is a two stage process. The first stage is a deployment wizard that deploys the OVA file of the appliance on the target ESXi host or vCenter Server instance. After the OVA deployment finishes, you are redirected to the second stage of the process that sets up and starts the services of the newly deployed appliance.
- The CLI deployment method involves running a CLI command against a JSON file that you previously prepared. The CLI installer parses the configuration parameters and their values from the JSON file and generates an OVF Tool command that automatically deploys and sets up the appliance.

Important For topologies with external Platform Services Controller instances, you must deploy the replicating Platform Services Controller instances in a sequence. After the successful deployment of all Platform Services Controller instances in the domain, you can perform concurrent deployments of multiple vCenter Server appliances that point to a common external Platform Services Controller instance.

The vCenter Server Appliance and Platform Services Controller appliance have the following default user names:

User Name	Description
root	Use this user name to log in to the appliance operating system and the Appliance Management Interface. You set the password while deploying the virtual appliance.
administrator@your_domain_name	Use this user name for vCenter Single Sign-On login. You set the password while creating the vCenter Single Sign-On domain. You create a vCenter Single Sign-On domain during the deployment of a vCenter Server Appliance with an embedded Platform Services Controller or the first Platform Services Controller instance in a new vCenter Single Sign-On domain. After you create a vCenter Single Sign-On domain, only the administrator@your_domain_name user has the privileges required to log in to vCenter Single Sign-On and vCenter Server. The administrator@your_domain_name user can proceed as follows: <ul style="list-style-type: none"> ■ Add an identity source in which additional users and groups are defined to vCenter Single Sign-On. ■ Give permissions to the users and groups. For information about adding identity sources and giving permissions to the users and groups, see <i>Platform Services Controller Administration</i> .

For information about upgrading and patching the vCenter Server Appliance and Platform Services Controller appliance, see *vSphere Upgrade*.

For information about configuring the vCenter Server Appliance and Platform Services Controller appliance, see *vCenter Server Appliance Configuration*.

Starting with vSphere 6.5, vCenter Server supports mixed IPv4 and IPv6 environment. If you want to set up the vCenter Server Appliance to use an IPv6 address version, use the fully qualified domain name (FQDN) or host name of the appliance. To set up an IPv4 address, the best practice is to use the FQDN or host name of the appliance, because the IP address can change if assigned by DHCP.

This section includes the following topics:

- [System Requirements for the vCenter Server Appliance and Platform Services Controller Appliance](#)
- [Preparing for Deployment of the vCenter Server Appliance and Platform Services Controller Appliance](#)
- [Prerequisites for Deploying the vCenter Server Appliance or Platform Services Controller Appliance](#)
- [GUI Deployment of the vCenter Server Appliance and Platform Services Controller Appliance](#)
- [CLI Deployment of the vCenter Server Appliance and Platform Services Controller Appliance](#)

System Requirements for the vCenter Server Appliance and Platform Services Controller Appliance

You can deploy the vCenter Server Appliance or Platform Services Controller appliance on an ESXi host 5.5 or later, or on a vCenter Server instance 5.5 or later. Your system must also meet specific software and hardware requirements.

When you use Fully Qualified Domain Names, verify that the client machine from which you are deploying the appliance and the network on which you are deploying the appliance use the same DNS server.

Before you deploy the appliance, synchronize the clocks of the target server and all vCenter Server and Platform Services Controller instances on the vSphere network. Unsynchronized clocks might result in authentication problems and can cause the installation to fail or prevent the appliance services from starting. See [Synchronizing Clocks on the vSphere Network](#).

Hardware Requirements for the vCenter Server Appliance and Platform Services Controller Appliance

When you deploy the vCenter Server Appliance, you can select to deploy an appliance that is suitable for the size of your vSphere environment. The option that you select determines the number of CPUs and the amount of memory for the appliance. The size of the Platform Services Controller appliance is the same for all environment sizes.

Hardware Requirements for the vCenter Server Appliance

The hardware requirements for a vCenter Server Appliance depend on the size of your vSphere inventory.

Table 3-1. Hardware Requirements for a vCenter Server Appliance with an Embedded or External Platform Services Controller

	Number of vCPUs	Memory
Tiny environment (up to 10 hosts or 100 virtual machines)	2	10 GB
Small environment (up to 100 hosts or 1,000 virtual machines)	4	16 GB
Medium environment (up to 400 hosts or 4,000 virtual machine)	8	24 GB
Large environment (up to 1,000 hosts or 10,000 virtual machines)	16	32 GB
X-Large environment (up to 2,000 hosts or 35,000 virtual machines)	24	48 GB

Note If you want to add an ESXi host with more than 512 LUNs and 2,048 paths to the vCenter Server Appliance inventory, you must deploy a vCenter Server Appliance for a large or x-large environment.

Hardware Requirements for the Platform Services Controller Appliance

The hardware requirements for a Platform Services Controller appliance are 2 vCPUs and 4 GB memory.

Storage Requirements for the vCenter Server Appliance and Platform Services Controller Appliance

When you deploy the vCenter Server Appliance or Platform Services Controller appliance, the ESXi host or DRS cluster on which you deploy the appliance must meet minimum storage requirements. The required storage depends not only on the size of the vSphere environment and the storage size, but also on the disk provisioning mode.

Storage Requirements for the vCenter Server Appliance

The storage requirements are different for each vSphere environment size and depend on your database size requirements.

Table 3-2. Storage Requirements for a vCenter Server Appliance with an Embedded or External Platform Services Controller

	Default Storage Size	Large Storage Size	X-Large Storage Size
Tiny environment (up to 10 hosts or 100 virtual machines)	250 GB	775 GB	1650 GB
Small environment (up to 100 hosts or 1,000 virtual machines)	290 GB	820 GB	1700 GB
Medium environment (up to 400 hosts or 4,000 virtual machine)	425 GB	925 GB	1805 GB
Large environment (up to 1,000 hosts or 10,000 virtual machines)	640 GB	990 GB	1870 GB
X-Large environment (up to 2,000 hosts or 35,000 virtual machines)	980 GB	1030 GB	1910 GB

Note The storage requirements include the requirements for the VMware Update Manager that runs as a service in the vCenter Server Appliance.

Storage Requirements for the Platform Services Controller Appliance

The storage requirement for a Platform Services Controller appliance is 60 GB.

Software Requirements for the vCenter Server Appliance and Platform Services Controller Appliance

The VMware vCenter Server Appliance and Platform Services Controller appliance can be deployed on ESXi hosts 5.5 or later, or on vCenter Server instances 5.5 or later.

You can deploy the vCenter Server Appliance or Platform Services Controller appliance by using the GUI or CLI installer. You run the installer from a network client machine that you use to connect to the target server and deploy the appliance on the server. You can connect directly to an ESXi 5.5.x or 6.x host on which to deploy the appliance. You can also connect to a vCenter Server 5.5.x or 6.x instance to deploy the appliance on an ESXi host or DRS cluster that resides in the vCenter Server inventory.

For information about the requirements for network client machine, see [System Requirements for the vCenter Server Appliance Installer](#).

Required Ports for vCenter Server and Platform Services Controller

The vCenter Server system, both on Windows and in the appliance, must be able to send data to every managed host and receive data from the vSphere Web Client and the Platform Services Controller services. To enable migration and provisioning activities between managed hosts, the source and destination hosts must be able to receive data from each other.

If a port is in use or is blacklisted, the vCenter Server installer displays an error message. You must use another port number to proceed with the installation. There are internal ports that are used only for inter-process communication.

VMware uses designated ports for communication. Additionally, the managed hosts monitor designated ports for data from vCenter Server. If a built-in firewall exists between any of these elements, the installer opens the ports during the installation or upgrade process. For custom firewalls, you must manually open the required ports. If you have a firewall between two managed hosts and you want to perform source or target activities, such as migration or cloning, you must configure a means for the managed hosts to receive data.

Note In Microsoft Windows Server 2008 and later, firewall is enabled by default.

Table 3-3. Ports Required for Communication Between Components

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
22	TCP/UDP	System port for SSHD.	Appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No
53		DNS service	Windows installations and appliance deployments of Platform Services Controller	No
80	TCP	<p>vCenter Server requires port 80 for direct HTTP connections. Port 80 redirects requests to HTTPS port 443. This redirection is useful if you accidentally use http://server instead of https://server. WS-Management (also requires port 443 to be open).</p> <p>If you use a Microsoft SQL database that is stored on the same virtual machine or physical server as the vCenter Server, port 80 is used by the SQL Reporting Service. When you install or upgrade vCenter Server, the installer prompts you to change the HTTP port for vCenter Server. Change the vCenter Server HTTP port to a custom value to ensure a successful installation or upgrade.</p> <p>Important You can change this port number during the vCenter Server and Platform Services Controller installations on Windows.</p>	Windows installations and appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No
88	TCP	Active Directory server. This port must be open for host to join Active Directory. If you use native Active Directory, the port must be open on both vCenter Server and Platform Services Controller.	Windows installations and appliance deployments of Platform Services Controller	No

Table 3-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
389	TCP/UDP	<p>This port must be open on the local and all remote instances of vCenter Server. This is the LDAP port number for the Directory Services for the vCenter Server group. If another service is running on this port, it might be preferable to remove it or change its port to a different port. You can run the LDAP service on any port from 1025 through 65535.</p> <p>If this instance is serving as the Microsoft Windows Active Directory, change the port number from 389 to an available port from 1025 through 65535.</p>	Windows installations and appliance deployments of Platform Services Controller	<ul style="list-style-type: none"> ■ vCenter Server to Platform Services Controller ■ Platform Services Controller to Platform Services Controller
443	TCP	<p>The default port that the vCenter Server system uses to listen for connections from the vSphere Web Client. To enable the vCenter Server system to receive data from the vSphere Web Client, open port 443 in the firewall.</p> <p>The vCenter Server system also uses port 443 to monitor data transfer from SDK clients.</p> <p>This port is also used for the following services:</p> <ul style="list-style-type: none"> ■ WS-Management (also requires port 80 to be open) ■ Third-party network management client connections to vCenter Server ■ Third-party network management clients access to hosts <p>Important You can change this port number during the vCenter Server and Platform Services Controller installations on Windows.</p>	<p>Windows installations and appliance deployments of</p> <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	<ul style="list-style-type: none"> ■ vCenter Server to vCenter Server ■ vCenter Server to Platform Services Controller ■ Platform Services Controller to vCenter Server
514	TCP/UDP	<p>vSphere Syslog Collector port for vCenter Server on Windows and vSphere Syslog Service port for vCenter Server Appliance</p> <p>Important You can change this port number during the vCenter Server and Platform Services Controller installations on Windows.</p>	<p>Windows installations and appliance deployments of</p> <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No

Table 3-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
636	TCP	vCenter Single Sign-On LDAPS For backward compatibility with vSphere 6.0 only.	Windows installations and appliance deployments of Platform Services Controller	During upgrade from vSphere 6.0 only. vCenter Server 6.0 to Platform Services Controller 6.5
902	TCP/UDP	The default port that the vCenter Server system uses to send data to managed hosts. Managed hosts also send a regular heartbeat over UDP port 902 to the vCenter Server system. This port must not be blocked by firewalls between the server and the hosts or between hosts. Port 902 must not be blocked between the VMware Host Client and the hosts. The VMware Host Client uses this port to display virtual machine consoles Important You can change this port number during the vCenter Server installations on Windows.	Windows installations and appliance deployments of vCenter Server	No
1514	TCP/UDP	vSphere Syslog Collector TLS port for vCenter Server on Windows and vSphere Syslog Service TLS port for vCenter Server Appliance Important You can change this port number during the vCenter Server and Platform Services Controller installations on Windows.	Windows installations and appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No
2012	TCP	Control interface RPC for vCenter Single Sign-On	Windows installations and appliance deployments of Platform Services Controller	<ul style="list-style-type: none"> ■ vCenter Server to Platform Services Controller ■ Platform Services Controller to vCenter Server ■ Platform Services Controller to Platform Services Controller
2014	TCP	RPC port for all VMCA (VMware Certificate Authority) APIs Important You can change this port number during the Platform Services Controller installations on Windows.	Windows installations and appliance deployments of Platform Services Controller	<ul style="list-style-type: none"> ■ vCenter Server to Platform Services Controller ■ Platform Services Controller to vCenter Server

Table 3-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
2015	TCP	DNS management	Windows installations and appliance deployments of Platform Services Controller	Platform Services Controller to Platform Services Controller
2020	TCP/UDP	Authentication framework management Important You can change this port number during the vCenter Server and Platform Services Controller installations on Windows.	Windows installations and appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	<ul style="list-style-type: none"> ■ vCenter Server to Platform Services Controller ■ Platform Services Controller to vCenter Server
5480	TCP	Appliance Management Interface Open endpoint serving all HTTPS, XMLRPC and JSON-RPC requests over HTTPS.	Appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No
6500	TCP/UDP	ESXi Dump Collector port Important You can change this port number during the vCenter Server installations on Windows.	Windows installations and appliance deployments of vCenter Server	No
6501	TCP	Auto Deploy service Important You can change this port number during the vCenter Server installations on Windows.	Windows installations and appliance deployments of vCenter Server	No
6502	TCP	Auto Deploy management Important You can change this port number during the vCenter Server installations on Windows.	Windows installations and appliance deployments of vCenter Server	No
7080, 12721	TCP	Secure Token Service Note Internal ports	Windows installations and appliance deployments of Platform Services Controller	No
7081	TCP	VMware Platform Services Controller Web Client Note Internal port	Windows installations and appliance deployments of Platform Services Controller	No
8200, 8201, 8300, 8301	TCP	Appliance management Note Internal ports	Appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No

Table 3-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
7444	TCP	Secure Token Service For backward compatibility with vSphere 5.5 only.	Windows installations and appliance deployments of Platform Services Controller	During upgrade from vSphere 5.5 only. <ul style="list-style-type: none"> ■ vCenter Server 5.5 to Platform Services Controller 6.5 ■ Platform Services Controller 6.5 to vCenter Server 5.5
8084	TCP	vSphere Update Manager SOAP port The port used by vSphere Update Manager client plug-in to connect to the vSphere Update Manager SOAP server.	Appliance deployments of vCenter Server	No
9084	TCP	vSphere Update Manager Web Server Port The HTTP port used by ESXi hosts to access host patch files from vSphere Update Manager server.	Appliance deployments of vCenter Server	No
9087	TCP	vSphere Update Manager Web SSL Port The HTTPS port used by vSphere Update Manager client plug-in to upload host upgrade files to vSphere Update Manager server.	Appliance deployments of vCenter Server	No
9123	TCP	Migration Assistant port Only when you run the Migration Assistant on the source Windows installation. The Migration Assistant lets you migrate Windows installations of vCenter Server and Platform Services Controller to appliances.	Windows installations and appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	During migration only. <ul style="list-style-type: none"> ■ Source vCenter Server 5.5 or 6.5 to target vCenter Server Appliance 6.5 ■ Source vCenter Single Sign-On 5.5 to target Platform Services Controller appliance 6.5 ■ Source Platform Services Controller 5.5 to target Platform Services Controller appliance 6.5
9443	TCP	vSphere Web Client HTTPS	Windows installations and appliance deployments of vCenter Server	No

Table 3-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
11711	TCP	vCenter Single Sign-On LDAP For backward compatibility with vSphere 5.5 only.	Windows installations and appliance deployments of Platform Services Controller	During upgrade from vSphere 5.5 only. vCenter Single Sign-On 5.5 to Platform Services Controller 6.5
11712	TCP	vCenter Single Sign-On LDAPS For backward compatibility with vSphere 5.5 only.	Windows installations and appliance deployments of Platform Services Controller	During upgrade from vSphere 5.5 only. vCenter Single Sign-On 5.5 to Platform Services Controller 6.5

To configure the vCenter Server system to use a different port to receive vSphere Web Client data, see the *vCenter Server and Host Management* documentation.

For more information about firewall configuration, see the *vSphere Security* documentation.

DNS Requirements for the vCenter Server Appliance and Platform Services Controller Appliance

When you deploy the vCenter Server Appliance or Platform Services Controller appliance, similar to any network server, you can assign a fixed IP address and an FQDN that is resolvable by a DNS server so that clients can reliably access the service.

When you deploy the vCenter Server Appliance or Platform Services Controller appliance with a static IP address, you ensure that in case of system restart, the IP address of the appliance remains the same.

Before you deploy the vCenter Server Appliance or Platform Services Controller appliance with a static IP address, you must verify that this IP address has a valid internal domain name system (DNS) registration.

When you deploy the vCenter Server Appliance, the installation of the Web server component that supports the vSphere Web Client fails if the installer cannot look up the fully qualified domain name (FQDN) for the appliance from its IP address. Reverse lookup is implemented using PTR records.

If you plan to use an FQDN for the appliance system name, you must verify that the FQDN is resolvable by a DNS server.

You can use the `nslookup` command to verify that the DNS reverse lookup service returns an FQDN when queried with the IP address and to verify that the FQDN is resolvable.

```
nslookup -nosearch -nodefname FQDN_or_IP_address
```

If you use DHCP instead of a static IP address for the vCenter Server Appliance or Platform Services Controller appliance, verify that the appliance name is updated in the domain name service (DNS). If you can ping the appliance name, the name is updated in DNS.

Ensure that the ESXi host management interface has a valid DNS resolution from the vCenter Server and all vSphere Web Client instances. Ensure that the vCenter Server has a valid DNS resolution from all ESXi hosts and all vSphere Web Clients.

vSphere Web Client Software Requirements

Make sure that your browser supports the vSphere Web Client.

The vSphere Web Client 6.5 requires Adobe Flash Player v. 16 to 23. For best performance and security fixes, use Adobe Flash Player 23.

VMware has tested and supports the following guest operating systems and browser versions for the vSphere Web Client. For best performance, use Google Chrome.

Table 3-4. Supported Guest Operating Systems and Minimum Browser Versions for the vSphere Web Client

Operating system	Browser
Windows	Microsoft Internet Explorer v. 10.0.19 and later. Mozilla Firefox v. 39 and later. Google Chrome v. 34 and later.
Mac OS	Mozilla Firefox v. 39 and later. Google Chrome v. 34 and later.

Preparing for Deployment of the vCenter Server Appliance and Platform Services Controller Appliance

Before you deploy the vCenter Server Appliance or Platform Services Controller appliance, you must download the vCenter Server Appliance installer ISO file and mount it to a network virtual machine or physical server from which you want to perform the deployment.

The machine from which you deploy the appliance must run on a Windows, Linux, or Mac operating system that meets the operating system requirements. See [System Requirements for the vCenter Server Appliance Installer](#).

System Requirements for the vCenter Server Appliance Installer

You can run the vCenter Server Appliance GUI or CLI installer from a network client machine that is running on a Windows, Linux, or Mac operating system of a supported version.

To ensure optimal performance of the GUI and CLI installers, use a client machine that meets the minimum hardware requirements.

Table 3-5. System Requirements for the GUI and CLI Installers

Operating System	Supported Versions	Minimum Hardware Configuration for Optimal Performance
Windows	<ul style="list-style-type: none"> ■ Windows 7, 8, 8.1, 10 ■ Windows 2012 x64 bit ■ Windows 2012 R2 x64 bit ■ Windows 2016 x64 bit 	4 GB RAM, 2 CPU having 4 cores with 2.3 GHz, 32 GB hard disk, 1 NIC
Linux	<ul style="list-style-type: none"> ■ SUSE 12 ■ Ubuntu 14.04 	4 GB RAM, 1 CPU having 2 cores with 2.3 GHz, 16 GB hard disk, 1 NIC Note The CLI installer requires 64-bit OS.
Mac	<ul style="list-style-type: none"> ■ macOS v10.9, 10.10, 10.11 ■ macOS Sierra 	8 GB RAM, 1 CPU having 4 cores with 2.4 GHz, 150 GB hard disk, 1 NIC

Note For client machines that run on Mac 10.11, concurrent GUI deployments of multiple appliances are unsupported. You must deploy the appliances in a sequence.

Download and Mount the vCenter Server Appliance Installer

VMware releases the vCenter Server Appliance ISO image, which contains GUI and CLI installers for the vCenter Server Appliance and Platform Services Controller appliance.

With the GUI and CLI executable files that are included in the vCenter Server Appliance installer, you can:

- Deploy the vCenter Server Appliance and Platform Services Controller appliance.
- Upgrade the vCenter Server Appliance and Platform Services Controller appliance.
- Migrate Windows installations of vCenter Server, vCenter Single Sign-On, and Platform Services Controller to the vCenter Server Appliance and Platform Services Controller appliance.
- Restore a vCenter Server Appliance from a file-based backup.

Prerequisites

- Create a My VMware account at <https://my.vmware.com/web/vmware/>.
- Verify that your client machine meets the system requirements for the vCenter Server Appliance installer. See [System Requirements for the vCenter Server Appliance Installer](#).

Procedure

- 1 From the VMware Web site at <https://my.vmware.com/web/vmware/downloads>, download the vCenter Server Appliance ISO image.

`VMware-VCSA-all-version_number-build_number.iso`

- 2 Confirm that the md5sum is correct.

See the VMware Web site topic *Using MD5 Checksums* at <http://www.vmware.com/download/md5.html>.

- 3 Mount or extract the ISO image to the client machine from which you want to deploy, upgrade, migrate, or restore the appliance.

Note ISO mounting or extracting software that does not allow more than eight directory levels, for example, MagicISO Maker on Windows, is unsupported.

For Linux OS and Mac OS, Archive Manager is unsupported.

For Mac OS, you can use DiskImageMounter.

For Ubuntu 14.04, you can use Disk Image Mounter.

For SUSE 12 OS, you can use the terminal.

```
$ sudo mkdir mount_dir
$ sudo mount -o loop VMware-vCSA-all-version-number-build-number.iso mount_dir
```

What to do next

Open the `readme.txt` file and review the information about the other files and directories in the vCenter Server Appliance ISO image.

Synchronizing Clocks on the vSphere Network

Verify that all components on the vSphere network have their clocks synchronized. If the clocks on the machines in your vSphere network are not synchronized, SSL certificates, which are time-sensitive, might not be recognized as valid in communications between network machines.

Unsynchronized clocks can result in authentication problems, which can cause the installation to fail or prevent the vCenter Server Appliance `vpxd` service from starting.

Verify that any Windows host machine on which vCenter Server runs is synchronized with the Network Time Server (NTP) server. See the Knowledge Base article <http://kb.vmware.com/kb/1318>.

To synchronize ESXi clocks with an NTP server, you can use the VMware Host Client. For information about editing the time configuration of an ESXi host, see *vSphere Single Host Management*.

Prerequisites for Deploying the vCenter Server Appliance or Platform Services Controller Appliance

To ensure successful deployment of the vCenter Server Appliance or Platform Services Controller appliance, you must perform some required tasks and pre-checks before running the installer.

General Prerequisites

- [Download and Mount the vCenter Server Appliance Installer.](#)

- For topologies with external Platform Services Controller instances, verify that you deploy the different nodes with time synchronization between each other. All vCenter Server instances, Platform Services Controller instances, and third-party load balancers in the vCenter Single Sign-On domain must be time synchronized. See [Synchronizing Clocks on the vSphere Network](#).

Target System Prerequisites

- Verify that your system meets the minimum software and hardware requirements. See [System Requirements for the vCenter Server Appliance and Platform Services Controller Appliance](#).
- If you want to deploy the appliance on an ESXi host, verify that the ESXi host is not in lockdown or maintenance mode and not part of a fully automated DRS cluster.
- If you want to deploy the appliance on a DRS cluster of the inventory of a vCenter Server instance, verify that the cluster contains at least one ESXi host that is not in lockdown or maintenance mode.
- If you plan to use NTP servers for time synchronization, verify that the NTP servers are running and that the time between the NTP servers and the target server on which you want to deploy the appliance is synchronized.

Network Prerequisites

If you plan to assign a static IP address and an FQDN as a system name in the network settings of the appliance, verify that you have configured the forward and reverse DNS records for the IP address.

GUI Deployment of the vCenter Server Appliance and Platform Services Controller Appliance

You can use the GUI installer to perform an interactive deployment of a vCenter Server Appliance with an embedded Platform Services Controller, a Platform Services Controller appliance, or a vCenter Server Appliance with an external Platform Services Controller.

When you perform the GUI deployment, you download the vCenter Server Appliance installer on a network client machine, run the deployment wizard from the client machine, and provide the inputs that are required for the appliance deployment and setup.

Important For topologies with external Platform Services Controller instances, you must deploy the replicating Platform Services Controller instances in a sequence. After the successful deployment of all Platform Services Controller instances in the domain, you can perform concurrent deployments of multiple vCenter Server appliances that point to a common external Platform Services Controller instance.

The GUI deployment process includes a series of two stages.

Figure 3-1. Stage 1 - OVA Deployment

The first stage walks you through the deployment wizard to choose the deployment type and appliance settings. This stage completes the deployment of the OVA file on the target server with the deployment type and appliance settings that you provide.

As an alternative to performing the first stage of the deployment with the GUI installer, you can deploy the OVA file of the vCenter Server Appliance or Platform Services Controller appliance by using the vSphere Web Client or VMware Host Client. To deploy the OVA file on an ESXi host or vCenter Server instance 5.5 or 6.0, you can also use the vSphere Client. After the OVA deployment, you must log in to the appliance management interface of the newly deployed appliance to proceed with the second stage of the deployment process.

Figure 3-2. Stage 2 - Appliance Setup

The second stage walks you through the setup wizard to configure the appliance time synchronization and vCenter Single Sign-On. This stage completes the initial setup and starts the services of the newly deployed appliance.

As an alternative to performing the second stage of the deployment with the GUI installer, you can log in to the Appliance Management Interface of the newly deployed appliance, https://FQDN_or_IP_address:5480.

Required Information for Deploying a vCenter Server Appliance or Platform Services Controller Appliance

When you use the GUI method to deploy a vCenter Server Appliance with an embedded Platform Services Controller, a Platform Services Controller appliance, or a vCenter Server Appliance with an external Platform Services Controller, the wizard prompts you for deployment and setup information. It is a best practice to keep a record of the values that you enter in case you must reinstall the product.

You can use this worksheet to record the information that you need for deploying a vCenter Server Appliance with an embedded Platform Services Controller, a Platform Services Controller appliance, or a vCenter Server Appliance with an external Platform Services Controller.

Table 3-6. Required Information During Stage 1 of the GUI Deployment Process

Required for Deployment of	Required Information	Default	Your Entry
All deployment types	FQDN or IP address of the target server on which you want to deploy the appliance. The target server can be either an ESXi host or a vCenter Server instance.	-	
	HTTPS port of the target server	443	
	User name with administrative privileges on the target server <ul style="list-style-type: none"> ■ If your target server is an ESXi host, use root. ■ If your target server is a vCenter Server instance, use <i>user_name@your_domain_name</i>, for example, <i>administrator@vsphere.local</i>. 	-	
	Password of the user with administrative privileges on the target server	-	
All deployment types Only if your target server is a vCenter Server instance	Data center from the vCenter Server inventory on which you want to deploy the appliance Optionally you can provide a data center folder.	-	
	ESXi host or DRS cluster from the data center inventory on which you want to deploy the appliance	-	
All deployment types	VM name for the appliance <ul style="list-style-type: none"> ■ Must not contain a percent sign (%), backslash (\), or forward slash (/) ■ Must be no more than 80 characters in length 	VMware vCenter Server Appliance	

Table 3-6. Required Information During Stage 1 of the GUI Deployment Process (Continued)

Required for Deployment of	Required Information	Default	Your Entry
All deployment types	Password for the root user of the appliance operating system <ul style="list-style-type: none"> ■ Must contain only lower ASCII characters without spaces. ■ Must be at least 8 characters, but no more than 20 characters in length ■ Must contain at least one uppercase letter ■ Must contain at least one lowercase letter ■ Must contain at least one number ■ Must contain at least one special character, for example, a dollar sign (\$), hash key (#), at sign (@), period (.), or exclamation mark (!) 	-	

Table 3-6. Required Information During Stage 1 of the GUI Deployment Process (Continued)

Required for Deployment of	Required Information	Default	Your Entry
<ul style="list-style-type: none"> ■ vCenter Server Appliance with an embedded Platform Services Controller 	Deployment size of the vCenter Server Appliance for your vSphere environment	Tiny	
<ul style="list-style-type: none"> ■ vCenter Server Appliance with an external Platform Services Controller 	<ul style="list-style-type: none"> ■ Tiny <p>Deploys an appliance with 2 CPUs and 10 GB of memory.</p> <p>Suitable for environments with up to 10 hosts or 100 virtual machines.</p> 		
	<ul style="list-style-type: none"> ■ Small <p>Deploys an appliance with 4 CPUs and 16 GB of memory.</p> <p>Suitable for environments with up to 100 hosts or 1,000 virtual machines.</p> 		
	<ul style="list-style-type: none"> ■ Medium <p>Deploys an appliance with 8 CPUs and 24 GB of memory.</p> <p>Suitable for environments with up to 400 hosts or 4,000 virtual machines.</p> 		
	<ul style="list-style-type: none"> ■ Large <p>Deploys an appliance with 16 CPUs and 32 GB of memory.</p> <p>Suitable for environments with up to 1,000 hosts or 10,000 virtual machines.</p> 		
	<ul style="list-style-type: none"> ■ X-Large <p>Deploys an appliance with 24 CPUs and 48 GB of memory.</p> <p>Suitable for environments with up to 2,000 hosts or 35,000 virtual machines.</p> 		

Table 3-6. Required Information During Stage 1 of the GUI Deployment Process (Continued)

Required for Deployment of	Required Information	Default	Your Entry
<ul style="list-style-type: none"> ■ vCenter Server Appliance with an embedded Platform Services Controller 	Storage size of the vCenter Server Appliance for your vSphere environment	Default	
<ul style="list-style-type: none"> ■ vCenter Server Appliance with an external Platform Services Controller 	<p>Increase the default storage size if you want larger volume for SEAT data (stats, events, alarms, and tasks).</p> <ul style="list-style-type: none"> ■ Default <p>For tiny deployment size, deploys the appliance with 250 GB of storage.</p> <p>For small deployment size, deploys the appliance with 290 GB of storage.</p> <p>For medium deployment size, deploys the appliance with 425 GB of storage.</p> <p>For large deployment size, deploys the appliance with 640 GB of storage.</p> <p>For x-large deployment size, deploys the appliance with 980 GB of storage.</p> ■ Large <p>For tiny deployment size, deploys the appliance with 775 GB of storage.</p> <p>For small deployment size, deploys the appliance with 820 GB of storage.</p> <p>For medium deployment size, deploys the appliance with 925 GB of storage.</p> <p>For large deployment size, deploys the appliance with 990 GB of storage.</p> <p>For x-large deployment size, deploys the appliance with 1030 GB of storage.</p> ■ X-Large <p>For tiny deployment size, deploys the appliance with 1650 GB of storage.</p> <p>For small deployment size, deploys the appliance with 1700 GB of storage.</p> 		

Table 3-6. Required Information During Stage 1 of the GUI Deployment Process (Continued)

Required for Deployment of	Required Information	Default	Your Entry
	<p>For medium deployment size, deploys the appliance with 1805 GB of storage.</p> <p>For large deployment size, deploys the appliance with 1870 GB of storage.</p> <p>For x-large deployment size, deploys the appliance with 1910 GB of storage.</p>		
All deployment types	<p>Name of the datastore on which you want to store the configuration files and virtual disks of the appliance</p> <hr/> <p>Note The installer displays a list of datastores that are accessible from your target server.</p>	-	
	Enable or disable Thin Disk Mode	Disabled	
All deployment types	<p>Name of the network to which to connect the appliance</p> <hr/> <p>Note The installer displays a drop-down menu with networks that depend on the network settings of your target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are not supported and are not displayed in the drop-down menu.</p> <hr/> <p>The network must be accessible from the client machine from which you perform the deployment.</p>	-	
	IP version for the appliance address	IPv4	
	Can be either IPv4 or IPv6.		
	IP assignment for the appliance address	static	
	Can be either static or DHCP.		
All deployment types	System name (FQDN or IP address)	-	
Only if you use a static assignment	The system name is used for managing the local system. The system name must be FQDN. If a DNS server is not available, provide a static IP address.		
	IP address	-	

Table 3-6. Required Information During Stage 1 of the GUI Deployment Process (Continued)

Required for Deployment of	Required Information	Default	Your Entry
	For IPv4 version, a subnet mask as a dot decimal notation or a network prefix as an integer between 0 and 32 For IPv6 version, a network prefix as an integer between 0 and 128	-	
	Default gateway	-	
	DNS servers separated by commas	-	
All deployment types Only if you use a DHCP assignment with IPv4 version and a DDNS server is available in your environment.	System name (FQDN)	-	

Table 3-7. Required Information During Stage 2 of the GUI Deployment Process

Required for	Required Information	Default	Your Entry
All deployment types	Time synchronization settings You can synchronize the time of the appliance either with the time of the ESXi host or with one or more NTP servers. If you want to use more than one NTP servers, you must provide the IP addresses or FQDNs of the NTP servers as a comma-separated list.	Synchronize time with NTP servers	
	Enable or disable SSH access	Disabled	
	Note vCenter Server Appliance high availability requires remote SSH access to the appliance.		
<ul style="list-style-type: none"> ■ vCenter Server Appliance with an embedded Platform Services Controller ■ Platform Services Controller appliance as the first instance in a new domain 	Name for the new vCenter Single Sign-On domain For example, vsphere.local.	-	
	Password for the administrator account, <i>administrator@your_domain_name</i> <ul style="list-style-type: none"> ■ Must be at least 8 characters, but no more than 20 characters in length ■ Must contain at least one uppercase letter ■ Must contain at least one lowercase letter ■ Must contain at least one number ■ Must contain at least one special character, such as ampersand (&), hash key (#), and percent sign (%) 	-	

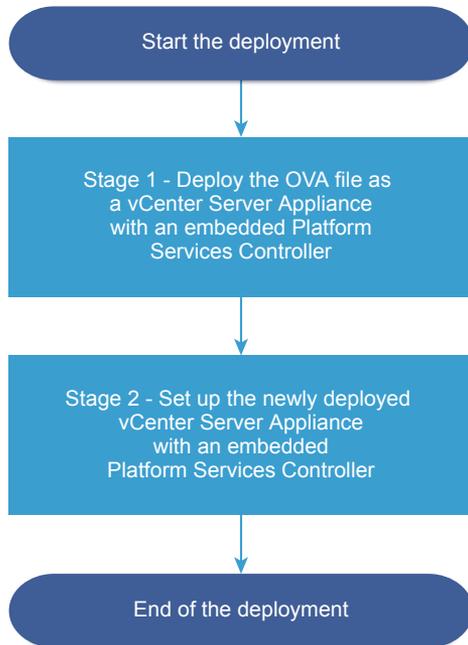
Table 3-7. Required Information During Stage 2 of the GUI Deployment Process (Continued)

Required for	Required Information	Default	Your Entry
	Site name	-	
<ul style="list-style-type: none"> ■ vCenter Server Appliance with an external Platform Services Controller ■ Platform Services Controller appliance as a subsequent instance in an existing domain 	<p>FQDN or IP address of the Platform Services Controller instance that you want to join</p> <p>You must join a Platform Services Controller instance of the same version.</p> <hr/> <p>HTTPS port of the Platform Services Controller instance</p> <hr/> <p>vCenter Single Sign On domain name of the Platform Services Controller instance For example, vsphere.local.</p> <hr/> <p>Password of the vCenter Single Sign On administrator user for the domain</p> <hr/> <p>vCenter Single Sign-On site name You can join an existing site or create a new site.</p>	<p>-</p> <p>443</p> <p>-</p> <p>-</p> <p>-</p>	
<ul style="list-style-type: none"> ■ vCenter Server Appliance with an embedded Platform Services Controller ■ Platform Services Controller appliance 	<p>Join or do not participate in the VMware Customer Experience Improvement Program (CEIP)</p> <p>For information about the CEIP, see the Configuring Customer Experience Improvement Program section in <i>vCenter Server and Host Management</i>.</p>	Join the CEIP	

Deploy the vCenter Server Appliance with an Embedded Platform Services Controller by Using the GUI

You can use the GUI installer to perform an interactive deployment of a vCenter Server Appliance with an embedded Platform Services Controller. You must run the GUI deployment from a Windows, Linux, or Mac machine that is in the network on which you want to deploy the appliance.

Figure 3-3. Deployment Workflow of a vCenter Server Appliance with an Embedded Platform Services Controller



Prerequisites

- See [Prerequisites for Deploying the vCenter Server Appliance or Platform Services Controller Appliance](#).
- See [Required Information for Deploying a vCenter Server Appliance or Platform Services Controller Appliance](#).

Procedure

1 [Stage 1 - Deploy the OVA File as a vCenter Server Appliance with an Embedded Platform Services Controller](#)

With stage 1 of the deployment process, you deploy the OVA file, which is included in the vCenter Server Appliance installer, as a vCenter Server Appliance with an embedded Platform Services Controller.

2 Stage 2 - Set up the Newly Deployed vCenter Server Appliance with an Embedded Platform Services Controller

When the OVA deployment finishes, you are redirected to stage 2 of the deployment process to set up and start the services of the newly deployed vCenter Server Appliance with an embedded Platform Services Controller.

Stage 1 - Deploy the OVA File as a vCenter Server Appliance with an Embedded Platform Services Controller

With stage 1 of the deployment process, you deploy the OVA file, which is included in the vCenter Server Appliance installer, as a vCenter Server Appliance with an embedded Platform Services Controller.

Procedure

- 1 In the vCenter Server Appliance installer, navigate to the `vcsa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.
- 2 On the Home page, click **Install** to start the deployment wizard.
- 3 Review the Introduction page to understand the deployment process and click **Next**.
- 4 Read and accept the license agreement, and click **Next**.
- 5 On the Select deployment type page, select **vCenter Server with an Embedded Platform Services Controller** and click **Next**.

This option deploys an appliance in which both the Platform Services Controller and vCenter Server are installed.

6 Connect to the target server on which you want to deploy the vCenter Server Appliance.

Option	Steps
You can connect to an ESXi host on which to deploy the appliance.	<ol style="list-style-type: none"> 1 Enter the FQDN or IP address of the ESXi host. 2 Enter the HTTPS port of the ESXi host. 3 Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. 4 Click Next. 5 Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target ESXi host, and click Yes to accept the certificate thumbprint.
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the appliance.	<ol style="list-style-type: none"> 1 Enter the FQDN or IP address of the vCenter Server instance. 2 Enter the HTTPS port of the vCenter Server instance. 3 Enter the user name and password of user with vCenter Single Sign-On administrative privileges on the vCenter Server instance, for example, the administrator@<i>your_domain_name</i> user. 4 Click Next. 5 Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target vCenter Server instance, and click Yes to accept the certificate thumbprint. 6 Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next <p>Note You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode.</p> <ol style="list-style-type: none"> 7 Select the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next

7 On the Set up appliance VM page, enter a name for the vCenter Server Appliance, set the password for the root user, and click **Next**.

The appliance name must not contain a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.

The password must contain only lower ASCII characters without spaces, at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

8 Select the deployment size for the vCenter Server Appliance for your vSphere inventory.

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 CPUs and 10 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 CPUs and 16 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 CPUs and 24 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines
Large	Deploys an appliance with 16 CPUs and 32 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 CPUs and 48 GB of memory. Suitable for environments with up to 2,000 hosts or 35,000 virtual machines

- 9 Select the storage size for the vCenter Server Appliance, and click **Next**.

Storage Size Option	Description for Tiny Deployment Size	Description for Small Deployment Size	Description for Medium Deployment Size	Description for Large Deployment Size	Description for X-Large Deployment Size
Default	Deploys an appliance with 250 GB of storage.	Deploys an appliance with 290 GB of storage.	Deploys an appliance with 425 GB of storage.	Deploys an appliance with 640 GB of storage.	Deploys an appliance with 980 GB of storage.
Large	Deploys an appliance with 775 GB of storage.	Deploys an appliance with 820 GB of storage.	Deploys an appliance with 925 GB of storage.	Deploys an appliance with 990 GB of storage.	Deploys an appliance with 1030 GB of storage.
X-Large	Deploys an appliance with 1650 GB of storage.	Deploys an appliance with 1700 GB of storage.	Deploys an appliance with 1805 GB of storage.	Deploys an appliance with 1870 GB of storage.	Deploys an appliance with 1910 GB of storage.

- 10 From the list of available datastores, select the location where all the virtual machine configuration files and virtual disks will be stored and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**.
- 11 On the Configure network settings page, set up the network settings.

The IP address or the FQDN of the appliance is used as a system name. It is recommended to use an FQDN. However, if you want to use an IP address, use static IP address allocation for the appliance, because IP addresses allocated by DHCP might change.

Option	Action
Network	<p>Select the network to which to connect the appliance.</p> <p>The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are not supported and are not displayed in the drop-down menu.</p>
IP version	<p>Select the version for the appliance IP address.</p> <p>You can select either IPv4 or IPv6.</p>
IP assignment	<p>Select how to allocate the IP address of the appliance.</p> <ul style="list-style-type: none"> ■ static <p>The wizard prompts you to enter the IP address and network settings.</p> <p>Note Avoid using an IP address as a system name. If you use an IP address as a system name, you cannot change the IP address and update the DNS settings after deployment, and you cannot join the appliance to an Active Directory domain.</p> ■ DHCP <p>A DHCP server is used to allocate the IP address. Select this option only if a DHCP server is available in your environment.</p> <p>If there is an enabled DDNS in your environment, you can enter a preferred fully qualified domain name (FQDN) for the appliance.</p>

- 12 On the Ready to complete stage 1 page, review the deployment settings for the vCenter Server Appliance and click **Finish** to start the OVA deployment process.

- 13 Wait for the OVA deployment to finish, and click **Continue** to proceed with stage 2 of the deployment process to set up and start the services of the newly deployed appliance.

Note If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Appliance Management Interface to set up and start the services.

The newly deployed vCenter Server Appliance with an embedded Platform Services Controller is running on the target server but the services are not started.

Stage 2 - Set up the Newly Deployed vCenter Server Appliance with an Embedded Platform Services Controller

When the OVA deployment finishes, you are redirected to stage 2 of the deployment process to set up and start the services of the newly deployed vCenter Server Appliance with an embedded Platform Services Controller.

Procedure

- 1 Review the introduction to stage 2 of the deployment process and click **Next**.
- 2 Configure the time settings in the appliance, optionally enable remote SSH access to the appliance, and click **Next**.

Option	Description
Synchronize time with the ESXi host	Enables periodic time synchronization, and VMware Tools sets the time of the guest operating system to be the same as the time of the ESXi host.
Synchronize time with NTP servers	Uses a Network Time Protocol server for synchronizing the time. If you select this option, you must enter the names or IP addresses of the NTP servers separated by commas.

- 3 On the SSO configuration page, create the vCenter Single Sign-On domain, and click **Next**.

- a Enter the domain name, for example, **vsphere.local**
- b Set the password for the vCenter Single Sign-On administrator account.

This is the password for the user `administrator@your_domain_name`.

After the deployment, you can log in to vCenter Single Sign-On and to vCenter Server as `administrator@your_domain_name`.

- c Enter the site name for vCenter Single Sign-On

The site name must contain alphanumeric characters. Choose your own name for the vCenter Single Sign-On site. You cannot change the name after installation.

Extended ASCII or non-ASCII characters are unsupported in site names. Your site name must include alphanumeric characters and a comma (,), period (.), question mark (?), dash (-), underscore (_), plus sign (+) or equals sign (=).

- 4 Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.

For information about the CEIP, see the Configuring Customer Experience Improvement Program section in *vCenter Server and Host Management*.

- 5 On the Ready to complete page, review the configuration settings for the vCenter Server Appliance, click **Finish**, and click **OK** to complete stage 2 of the deployment process and set up the appliance.
- 6 (Optional) After the initial setup finishes, click the **https://vcenter_server_appliance_fqdn/vsphere-client** to go to the vSphere Web Client and log in to the vCenter Server instance in the vCenter Server Appliance, or click the **https://vcenter_server_appliance_fqdn:443** to go the vCenter Server Appliance Getting Started page.
- 7 Click **Close** to exit the wizard.

You are redirected to the vCenter Server Appliance Getting Started page.

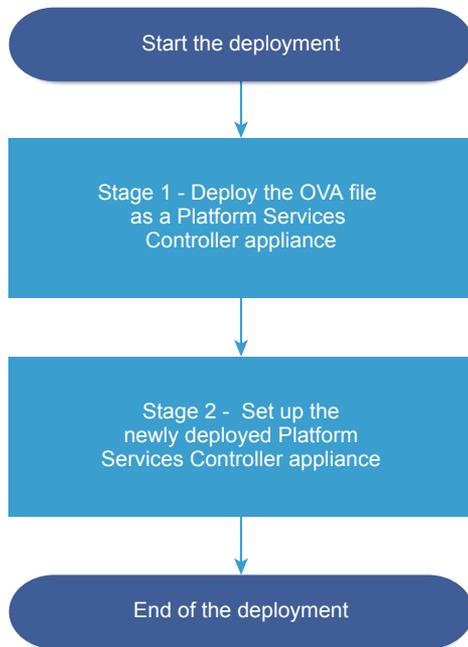
What to do next

You can configure high availability for the vCenter Server Appliance. For information about providing vCenter Server Appliance high availability, see *vSphere Availability*.

Deploy a Platform Services Controller Appliance by Using the GUI

You can use the GUI installer to perform an interactive deployment of a Platform Services Controller appliance as the first instance in a new vCenter Single Sign-On domain or as a replication partner in an existing vCenter Single Sign-On domain. You must run the GUI deployment from a Windows, Linux, or Mac machine that is in the network on which you want to deploy the appliance.

Important You must deploy the replicating Platform Services Controller instances in a sequence.

Figure 3-4. Deployment Workflow of a Platform Services Controller Appliance**Prerequisites**

- See [Prerequisites for Deploying the vCenter Server Appliance or Platform Services Controller Appliance](#).
- See [Required Information for Deploying a vCenter Server Appliance or Platform Services Controller Appliance](#).

Procedure**1 Stage 1 - Deploy the OVA File as a Platform Services Controller Appliance**

With stage 1 of the deployment process, you deploy the OVA file, which is included in the vCenter Server Appliance installer, as a Platform Services Controller appliance.

2 Stage 2 - Set up the Newly Deployed Platform Services Controller Appliance

When the OVA deployment finishes, you are redirected to stage 2 of the deployment process to set up and start the services of the newly deployed Platform Services Controller appliance.

Stage 1 - Deploy the OVA File as a Platform Services Controller Appliance

With stage 1 of the deployment process, you deploy the OVA file, which is included in the vCenter Server Appliance installer, as a Platform Services Controller appliance.

Procedure

- 1 In the vCenter Server Appliance installer, navigate to the `vcasa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.

- For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.
- 2 On the Home page, click **Install** to start the deployment wizard.
 - 3 Review the Introduction page to understand the deployment process and click **Next**.
 - 4 Read and accept the license agreement, and click **Next**.
 - 5 On the Select a deployment type page, select **Platform Services Controller** and click **Next**.
 - 6 Connect to the target server on which you want to deploy the Platform Services Controller appliance and click **Next**.

Option	Steps
You can connect to an ESXi host on which to deploy the appliance.	<ol style="list-style-type: none"> 1 Enter the FQDN or IP address of the ESXi host. 2 Enter the HTTPS port of the ESXi host. 3 Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. 4 Click Next. 5 Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target ESXi host, and click Yes to accept the certificate thumbprint.
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the appliance.	<ol style="list-style-type: none"> 1 Enter the FQDN or IP address of the vCenter Server instance. 2 Enter the HTTPS port of the vCenter Server instance. 3 Enter the user name and password of user with vCenter Single Sign-On administrative privileges on the vCenter Server instance, for example, the <code>administrator@your_domain_name</code> user. 4 Click Next. 5 Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target vCenter Server instance, and click Yes to accept the certificate thumbprint. 6 Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next <p>Note You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode.</p> <ol style="list-style-type: none"> 7 Select the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next

- 7 On the Set up appliance VM page, enter a name for the Platform Services Controller appliance, set the password for the root user, and click **Next**.

The appliance name must not contain a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.

The password must contain only lower ASCII characters without spaces, at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

- 8 From the list of available datastores, select the location where all the virtual machine configuration files and virtual disks will be stored and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**.

- 9 On the Configure network settings page, set up the network settings.

The IP address or the FQDN of the appliance is used as a system name. It is recommended to use an FQDN. However, if you want to use an IP address, use static IP address allocation for the appliance, because IP addresses allocated by DHCP might change.

Option	Action
Network	<p>Select the network to which to connect the appliance.</p> <p>The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are not supported and are not displayed in the drop-down menu.</p>
IP version	<p>Select the version for the appliance IP address.</p> <p>You can select either IPv4 or IPv6.</p>
IP assignment	<p>Select how to allocate the IP address of the appliance.</p> <ul style="list-style-type: none"> ▪ static <p>The wizard prompts you to enter the IP address and network settings.</p> <p>Note Avoid using an IP address as a system name. If you use an IP address as a system name, you cannot change the IP address and update the DNS settings after deployment, and you cannot join the appliance to an Active Directory domain.</p> ▪ DHCP <p>A DHCP server is used to allocate the IP address. Select this option only if a DHCP server is available in your environment.</p> <p>If there is an enabled DDNS in your environment, you can enter a preferred fully qualified domain name (FQDN) for the appliance.</p>

- 10 On the Ready to complete stage 1 page, review the deployment settings for the Platform Services Controller appliance and click **Finish** to start the OVA deployment process.
- 11 Wait for the OVA deployment to finish, and click **Continue** to proceed with stage 2 of the deployment process to set up and start the services of the newly deployed appliance.

Note If you exit the wizard by clicking **Close**, you must log in to the Platform Services Controller Appliance Management Interface to set up and start the services.

The newly deployed Platform Services Controller appliance is running on the target server but the services are not started.

Stage 2 - Set up the Newly Deployed Platform Services Controller Appliance

When the OVA deployment finishes, you are redirected to stage 2 of the deployment process to set up and start the services of the newly deployed Platform Services Controller appliance.

Procedure

- 1 Review the introduction to stage 2 of the deployment process and click **Next**.

- 2 Configure the time settings in the appliance, optionally enable remote SSH access to the appliance, and click **Next**.

Option	Description
Synchronize time with the ESXi host	Enables periodic time synchronization, and VMware Tools sets the time of the guest operating system to be the same as the time of the ESXi host.
Synchronize time with NTP servers	Uses a Network Time Protocol server for synchronizing the time. If you select this option, you must enter the names or IP addresses of the NTP servers separated by commas.

- 3 Create a new vCenter Single Sign-On domain or join an existing domain.

Option	Description
Create a new Single Sign-On domain	<p>Creates a vCenter Single Sign-On domain.</p> <ol style="list-style-type: none"> Enter the domain name, for example vsphere.local. Set the password for the vCenter Single Sign-On administrator account. This is the password for the user <code>administrator@your_domain_name</code>. Enter the site name for vCenter Single Sign-On. The site name is important if you are using vCenter Single Sign-On in multiple locations. The site name must contain alphanumeric characters. Choose your own name for the vCenter Single Sign-On site. You cannot change the name after installation. Extended ASCII and non-ASCII characters are unsupported in site names. Your site name must include alphanumeric characters and a comma (,), period (.), question mark (?), dash (-), underscore (_), plus sign (+) or equals sign (=). Click Next.
Join an existing vCenter Single Sign-On domain	<p>Joins the Platform Services Controller appliance to an existing vCenter Single Sign-On domain as a replication partner of an existing Platform Services Controller instance. You must provide the information about the partner Platform Services Controller instance that you want to join.</p> <ol style="list-style-type: none"> Enter the fully qualified domain name (FQDN) or IP address of the partner Platform Services Controller instance. Enter the HTTPS port of the partner Platform Services Controller instance. Enter the vCenter Single Sign-On domain name of the partner Platform Services Controller instance. Enter the password of the vCenter Single Sign-On administrator user. Click Next. Select whether to create or join an existing vCenter Single Sign-On site.

- 4 Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.

For information about the CEIP, see the *Configuring Customer Experience Improvement Program* section in *vCenter Server and Host Management*.

- 5 On the Ready to complete page, review the configuration settings for the Platform Services Controller appliance, click **Finish**, and click **OK** to complete stage 2 of the deployment process and set up the appliance.

6 (Optional) After the initial setup finishes, click the https://platform_services_controller_fqdn/psc to go to the Platform Services Controller Web interface, or click the https://platform_services_controller_fqdn:443 to go the Platform Services Controller Getting Started page.

7 Click **Close** to exit the wizard.

You are redirected to the Platform Services Controller Getting Started page.

If you joined the new Platform Services Controller appliance to an existing vCenter Single Sign-On domain, the appliance replicates infrastructure data with the other Platform Services Controller instances within the domain.

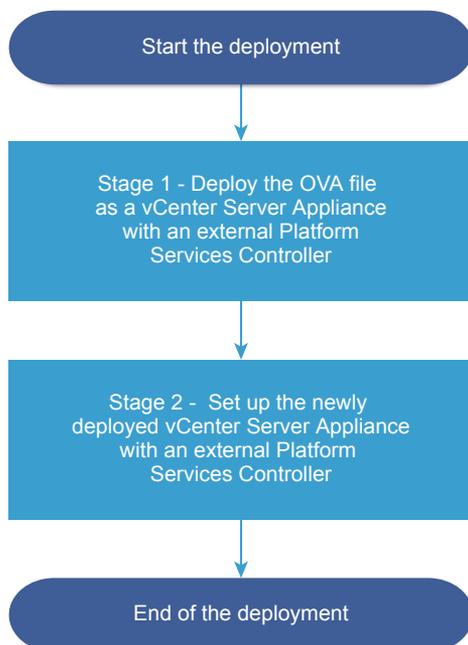
What to do next

- You can deploy a vCenter Server Appliance with an external Platform Services Controller and register it with the newly deployed Platform Services Controller appliance.
- You can deploy one or more Platform Services Controller instances joining the same vCenter Single Sign-On domain to replicate infrastructure data and distribute the load.

Deploy the vCenter Server Appliance with an External Platform Services Controller by Using the GUI

You can use the GUI installer to perform an interactive deployment of a vCenter Server Appliance and register it with an existing external Platform Services Controller instance. You must run the GUI deployment from a Windows, Linux, or Mac machine that is in the network on which you want to deploy the appliance.

Figure 3-5. Deployment Workflow of a vCenter Server Appliance with an External Platform Services Controller



Prerequisites

- See [Prerequisites for Deploying the vCenter Server Appliance or Platform Services Controller Appliance](#).
- See [Required Information for Deploying a vCenter Server Appliance or Platform Services Controller Appliance](#).
- Verify that you have installed or deployed the Platform Services Controller instance with which you plan to register the vCenter Server Appliance.

Procedure

1 [Stage 1 - Deploy the OVA File as a vCenter Server Appliance With an External Platform Services Controller](#)

With stage 1 of the deployment process, you deploy the OVA file, which is included in the vCenter Server Appliance installer, as a vCenter Server Appliance with an external Platform Services Controller.

2 [Stage 2 - Set up the Newly Deployed vCenter Server Appliance With an External Platform Services Controller](#)

When the OVA deployment finishes, you are redirected to stage 2 of the deployment process to set up and start the services of the newly deployed vCenter Server Appliance with an external Platform Services Controller.

Stage 1 - Deploy the OVA File as a vCenter Server Appliance With an External Platform Services Controller

With stage 1 of the deployment process, you deploy the OVA file, which is included in the vCenter Server Appliance installer, as a vCenter Server Appliance with an external Platform Services Controller.

Procedure

- 1 In the vCenter Server Appliance installer, navigate to the `vcasa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.
- 2 On the Home page, click **Install** to start the deployment wizard.
- 3 Review the Introduction page to understand the deployment process and click **Next**.
- 4 Read and accept the license agreement, and click **Next**.
- 5 On the Select deployment type page, select **vCenter Server (Requires External Platform Services Controller)** and click **Next**.

6 Connect to the target server on which you want to deploy the vCenter Server Appliance.

Option	Steps
You can connect to an ESXi host on which to deploy the appliance.	<ol style="list-style-type: none"> 1 Enter the FQDN or IP address of the ESXi host. 2 Enter the HTTPS port of the ESXi host. 3 Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. 4 Click Next. 5 Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target ESXi host, and click Yes to accept the certificate thumbprint.
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the appliance.	<ol style="list-style-type: none"> 1 Enter the FQDN or IP address of the vCenter Server instance. 2 Enter the HTTPS port of the vCenter Server instance. 3 Enter the user name and password of user with vCenter Single Sign-On administrative privileges on the vCenter Server instance, for example, the administrator@<i>your_domain_name</i> user. 4 Click Next. 5 Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target vCenter Server instance, and click Yes to accept the certificate thumbprint. 6 Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next <p>Note You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode.</p> <ol style="list-style-type: none"> 7 Select the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next

7 On the Set up appliance VM page, enter a name for the vCenter Server Appliance, set the password for the root user, and click **Next**.

The appliance name must not contain a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.

The password must contain only lower ASCII characters without spaces, at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

8 Select the deployment size for the vCenter Server Appliance for your vSphere inventory.

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 CPUs and 10 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 CPUs and 16 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 CPUs and 24 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines
Large	Deploys an appliance with 16 CPUs and 32 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 CPUs and 48 GB of memory. Suitable for environments with up to 2,000 hosts or 35,000 virtual machines

- 9 Select the storage size for the vCenter Server Appliance, and click **Next**.

Storage Size Option	Description for Tiny Deployment Size	Description for Small Deployment Size	Description for Medium Deployment Size	Description for Large Deployment Size	Description for X-Large Deployment Size
Default	Deploys an appliance with 250 GB of storage.	Deploys an appliance with 290 GB of storage.	Deploys an appliance with 425 GB of storage.	Deploys an appliance with 640 GB of storage.	Deploys an appliance with 980 GB of storage.
Large	Deploys an appliance with 775 GB of storage.	Deploys an appliance with 820 GB of storage.	Deploys an appliance with 925 GB of storage.	Deploys an appliance with 990 GB of storage.	Deploys an appliance with 1030 GB of storage.
X-Large	Deploys an appliance with 1650 GB of storage.	Deploys an appliance with 1700 GB of storage.	Deploys an appliance with 1805 GB of storage.	Deploys an appliance with 1870 GB of storage.	Deploys an appliance with 1910 GB of storage.

- 10 From the list of available datastores, select the location where all the virtual machine configuration files and virtual disks will be stored and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**.
- 11 On the Configure network settings page, set up the network settings.

The IP address or the FQDN of the appliance is used as a system name. It is recommended to use an FQDN. However, if you want to use an IP address, use static IP address allocation for the appliance, because IP addresses allocated by DHCP might change.

Option	Action
Network	<p>Select the network to which to connect the appliance.</p> <p>The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are not supported and are not displayed in the drop-down menu.</p>
IP version	<p>Select the version for the appliance IP address.</p> <p>You can select either IPv4 or IPv6.</p>
IP assignment	<p>Select how to allocate the IP address of the appliance.</p> <ul style="list-style-type: none"> ■ static <p>The wizard prompts you to enter the IP address and network settings.</p> <p>Note Avoid using an IP address as a system name. If you use an IP address as a system name, you cannot change the IP address and update the DNS settings after deployment.</p> ■ DHCP <p>A DHCP server is used to allocate the IP address. Select this option only if a DHCP server is available in your environment.</p> <p>If there is an enabled DDNS in your environment, you can enter a preferred fully qualified domain name (FQDN) for the appliance.</p>

- 12 On the Ready to complete stage 1 page, review the deployment settings for the vCenter Server Appliance and click **Finish** to start the OVA deployment process.

- 13 Wait for the OVA deployment to finish, and click **Continue** to proceed with stage 2 of the deployment process to set up and start the services of the newly deployed appliance.

Note If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Appliance Management Interface to set up and start the services.

The newly deployed vCenter Server Appliance with an external Platform Services Controller is running on the target server but the services are not started.

Stage 2 - Set up the Newly Deployed vCenter Server Appliance With an External Platform Services Controller

When the OVA deployment finishes, you are redirected to stage 2 of the deployment process to set up and start the services of the newly deployed vCenter Server Appliance with an external Platform Services Controller.

Procedure

- 1 Review the introduction to stage 2 of the deployment process and click **Next**.
- 2 Configure the time settings in the appliance, optionally enable remote SSH access to the appliance, and click **Next**.

Option	Description
Synchronize time with the ESXi host	Enables periodic time synchronization, and VMware Tools sets the time of the guest operating system to be the same as the time of the ESXi host.
Synchronize time with NTP servers	Uses a Network Time Protocol server for synchronizing the time. If you select this option, you must enter the names or IP addresses of the NTP servers separated by commas.

- 3 Provide the FQDN or IP address of the Platform Services Controller instance with which you want to register the vCenter Server Appliance, enter the vCenter Single Sign-On HTTPS port, domain name, and administrator password, and click **Next**.

If the Platform Services Controller instance is a Windows installation, provide the system name of the host machine on which the Platform Services Controller is running.

- 4 On the Ready to complete page, review the configuration settings for the vCenter Server Appliance, click **Finish**, and click **OK** to complete stage 2 of the deployment process and set up the appliance.
- 5 (Optional) After the initial setup finishes, click the **https://vcenter_server_appliance_fqdn/vsphere-client** to go to the vSphere Web Client and log in to the vCenter Server instance in the vCenter Server Appliance, or click the **https://vcenter_server_appliance_fqdn:443** to go the vCenter Server Appliance Getting Started page.
- 6 Click **Close** to exit the wizard.

You are redirected to the vCenter Server Appliance Getting Started page.

The newly deployed vCenter Server Appliance joined the vCenter Single Sign-On domain and site of the Platform Services Controller instance with which you registered the appliance.

What to do next

You can configure high availability for the vCenter Server Appliance. For information about providing vCenter Server Appliance high availability, see *vSphere Availability*.

CLI Deployment of the vCenter Server Appliance and Platform Services Controller Appliance

You can use the CLI installer to perform a silent deployment of a vCenter Server Appliance or Platform Services Controller appliance on an ESXi host or vCenter Server instance.

The CLI deployment process includes downloading the vCenter Server Appliance installer on a network virtual machine or physical server from which you want to perform the deployment, preparing a JSON configuration file with the deployment information, and running the deployment command.

Important The user name that you use to log in to the machine from which you want to run the CLI installer, the path to the vCenter Server Appliance installer, the path to your JSON configuration file, and the string values in your JSON configuration file, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

The vCenter Server Appliance ISO file contains templates of JSON files that contain the minimum configuration parameters that are required for deploying the vCenter Server Appliance or Platform Services Controller appliance.

The vCenter Server Appliance ISO file contains templates of JSON files that contain the minimum configuration parameters that are required for deploying the vCenter Server Appliance or Platform Services Controller appliance. For information about preparing JSON templates for CLI deployment, see [Prepare Your JSON Configuration File for CLI Deployment](#).

Important For topologies with external Platform Services Controller instances, you must deploy the replicating Platform Services Controller instances in a sequence. After the successful deployment of all Platform Services Controller instances in the domain, you can perform concurrent deployments of multiple vCenter Server appliances that point to a common external Platform Services Controller instance.

Prepare Your JSON Configuration File for CLI Deployment

Before you run the CLI installer to deploy a vCenter Server Appliance or Platform Services Controller appliance, you must prepare a JSON file with configuration parameters and their values for your deployment specification.

The vCenter Server Appliance installer contains JSON templates for all deployment types. For information about the templates, see [JSON Templates for CLI Deployment of the vCenter Server Appliance and Platform Services Controller Appliance](#).

You can deploy an appliance with minimum configurations by setting values to the configuration parameters in the JSON template for your specification. You can edit the preset values, remove configuration parameters, and add configuration parameters for custom configurations.

For a complete list of the configuration parameters and their descriptions, navigate to the installer subdirectory for your operating system and run the `vcsa-deploy install --template-help` command or see [Deployment Configuration Parameters](#).

Prerequisites

- You must be familiar with the JSON syntax.
- [Download and Mount the vCenter Server Appliance Installer](#).

Procedure

- 1 In the vCenter Server Appliance installer, navigate to the `vcsa-cli-installer` directory, and open the `templates` subfolder.
- 2 Copy the deployment templates from the `install` subfolder to your workspace.

Important The path to the JSON configuration files must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

- 3 In a text editor, open the template file for your specification.

To ensure the correct syntax of your JSON configuration file, use a JSON editor.

- 4 Fill in the values for the required configuration parameters and, optionally, enter additional parameters and their values.

For example, if you want to use an IPv4 DHCP assignment for the network of the appliance, in the `network` subsection of the template, change the value of the `mode` parameter to `dhcp` and remove the default configuration parameters that are for a static assignment.

```
"network": {
  "ip.family": "ipv4",
  "mode": "dhcp"
},
```

Important The string values, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

To set a value that contains a backslash (`\`) or quotation mark (`"`) character, you must precede the character with the backslash (`\`) character. For example, `"password": "my\"password"` sets the password `my"password`, `"image": "G:\\vcsa\\VMware-vCenter-Server-Appliance-6.5.0.XXXX-YYYYYYY_OVF10.ova"` sets the path `G:\vcsa\VMware-vCenter-Server-Appliance-6.5.0.XXXX-YYYYYYY_OVF10.ova`.

The Boolean values must contain only lowercase characters, that is, a value can be either `true` or `false`. For example, `"ssh.enable": false`.

- 5 (Optional) Use a JSON editor of your choice to validate the JSON file.
- 6 Save in UTF-8 format and close the file.

What to do next

You can create and save additional templates if needed for your deployment specification.

JSON Templates for CLI Deployment of the vCenter Server Appliance and Platform Services Controller Appliance

The vCenter Server Appliance installer contains JSON templates that are located in the `vcsa-cli-installer/templates` directory. In the `install` subfolder, you can find eight JSON templates with the minimum configuration parameters for all deployment types.

For each deployment type, there is one template for deploying the appliance on an ESXi host and another template for deploying the appliance on a vCenter Server instance.

Table 3-8. Deployment JSON Templates Included in the vCenter Server Appliance Installer

Location	Template	Description
vcsa-cli-installer\templates\install	embedded_vCSA_on_ESXi.json	Contains the minimum configuration parameters that are required for deployment of a vCenter Server Appliance with an embedded Platform Services Controller on an ESXi host.
	embedded_vCSA_on_VC.json	Contains the minimum configuration parameters that are required for deployment of a vCenter Server Appliance with an embedded Platform Services Controller on a vCenter Server instance.
	PSC_first_instance_on_ESXi.json	Contains the minimum configuration parameters that are required for deployment of a Platform Services Controller appliance as the first instance in a new vCenter Single Sign-On domain on an ESXi host.
	PSC_first_instance_on_VC.json	Contains the minimum configuration parameters that are required for deployment of a Platform Services Controller appliance as the first instance in a new vCenter Single Sign-On domain on a vCenter Server instance.
	PSC_replication_on_ESXi.json	Contains the minimum configuration parameters that are required for deployment of a Platform Services Controller appliance joining an existing vCenter Single Sign-On domain on an ESXi host.

Table 3-8. Deployment JSON Templates Included in the vCenter Server Appliance Installer (Continued)

Location	Template	Description
	PSC_replication_on_VC.json	Contains the minimum configuration parameters that are required for deployment of a Platform Services Controller appliance joining an existing vCenter Single Sign-On domain on a vCenter Server instance.
	vCSA_on_ESXi.json	Contains the minimum configuration parameters that are required for deployment of a vCenter Server Appliance with an external Platform Services Controller on an ESXi host.
	vCSA_on_VC.json	Contains the minimum configuration parameters that are required for deployment of a vCenter Server Appliance with an external Platform Services Controller on a vCenter Server instance.

Deployment Configuration Parameters

When you prepare your JSON configuration files for CLI deployment, you must set parameters and values to provide input data for the deployment of a vCenter Server Appliance or Platform Services Controller appliance.

Sections and Subsections of Configuration Parameters in the JSON Deployment Files

The configuration parameters in the JSON configuration files for CLI upgrade are organized in sections and subsections.

Table 3-9. Sections and Subsections of Configuration Parameters in the JSON Deployment Files

Section	Subsection	Description
new.vcsa - describes the appliance that you want to deploy	esxi	Use only if you want to deploy the appliance directly on an ESXi host. Contains the configuration parameters that describe the target ESXi host. See Table 3-10 . Note You must fill in either the esxi or the vc subsection.
	vc	Use only if you want to deploy the appliance on the inventory of a vCenter Server instance. Contains the configuration parameters that describe the target ESXi host or DRS cluster from the vCenter Server inventory. See Table 3-11 . Note You must fill in either the vc or the esxi subsection.

Table 3-9. Sections and Subsections of Configuration Parameters in the JSON Deployment Files (Continued)

Section	Subsection	Description
	appliance	Contains the configuration parameters that describe the appliance. See Table 3-12 .
	network	Contains the configuration parameters that describe the network settings for the appliance. See Table 3-13 .
	os	Contains the configuration parameters that describe the operating system settings for the appliance. See Table 3-14 .
	sso	Contains the configuration parameters that describe the vCenter Single Sign-On settings for the appliance. See Table 3-15 .
	ovftool.arguments	Optional subsection for adding arbitrary arguments and their values to the OVF Tool command that the installer generates. Important The vCenter Server Appliance installer does not validate the configuration parameters in the <code>ovftool.arguments</code> subsection. If you set arguments that the OVF Tool does not recognize, the deployment might fail.
ceip - describes joining the VMware Customer Experience Improvement Program (CEIP)	settings	Contains only the <code>ceip.enabled</code> configuration parameter to join or not to join the VMware Customer Experience Improvement Program (CEIP). See Table 3-16 . Required only if you are deploying a vCenter Server Appliance with an embedded Platform Services Controller or a Platform Services Controller appliance. Note If set to <code>true</code> , you must run the CLI deployment command with the <code>--acknowledge-ceip</code> argument. For information about the CEIP, see the <i>Configuring Customer Experience Improvement Program</i> section in <i>vCenter Server and Host Management</i> .

Important The string values, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

To set a value that contains a backslash (\) or quotation mark (") character, you must precede the character with the backslash (\) character. For example, `"password": "my\"password"` sets the password `my"password`, `"image": "G:\vcsa\VMware-vCenter-Server-Appliance-6.5.0.XXX-YYYYYYY_OVF10.ova"` sets the path `G:\vcsa\VMware-vCenter-Server-Appliance-6.5.0.XXX-YYYYYYY_OVF10.ova`.

The Boolean values must contain only lowercase characters. Can be either `true` or `false`. For example, `"ssh.enable": false`.

Configuration Parameters in the `new.vcsa` Section

Table 3-10. Configuration Parameters in the `new.vcsa` Section, `esxi` Subsection

Name	Type	Description
hostname	string	The IP address or FQDN of the target ESXi host on which you want to deploy the appliance.
username	string	A user name with administrative privileges on the target ESXi host, for example, <code>root</code> .

Table 3-10. Configuration Parameters in the new.vcsa Section, esxi Subsection (Continued)

Name	Type	Description
password	string	The password of the user with administrative privileges on the target ESXi host.
deployment.network	string	The name of the network to which to connect the appliance. Note The network must be accessible from the target ESXi host. Ignored if the target ESXi host has only one network.
datastore	string	The name of the datastore that you want to store all virtual machine configuration files and virtual disks of the appliance. Note The datastore must be accessible from the ESXi host. The datastore must have enough free space.
port	integer	The HTTPS reverse proxy port of the target ESXi host. The default port is 443. Use only if the target ESXi host uses a custom HTTPS reverse proxy port.

Table 3-11. Configuration Parameters in the new.vcsa Section, vc Subsection

Name	Type	Description
hostname	string	The IP address or FQDN of the target vCenter Server instance on which you want to deploy the appliance.
username	string	vCenter Single Sign-On administrator user name on the target vCenter Server instance, for example, administrator@vsphere.local.
password	string	The password of the vCenter Single Sign-On administrator user on the target vCenter Server instance.
deployment.network	string	The name of the network to which to connect the appliance. Note The network must be accessible from the target ESXi host or DRS cluster on which you want to deploy the appliance. Ignored if the target ESXi host or DRS cluster has only one network.
datacenter	string or array	The vCenter Server datacenter that contains the target ESXi host or DRS cluster on which you want to deploy the appliance. If the datacenter is located in a folder or a structure of folders, the value must be either a comma-separated list of strings or a comma-separated list as a single string. For example, <pre>["parent_folder", "child_folder", "datacenter_name"]</pre> or <pre>"parent_folder, child_folder, datacenter_name"</pre> Note The value is case-sensitive.
datastore	string	The name of the datastore that you want to store all virtual machine configuration files and virtual disks of the appliance. Note The datastore must be accessible from the target ESXi host or DRS cluster. The datastore must have at least 15 GB of free space.

Table 3-11. Configuration Parameters in the new.vcsa Section, vc Subsection (Continued)

Name	Type	Description
port	integer	<p>The HTTPS reverse proxy port of the target vCenter Server instance.</p> <p>The default port is 443. Use only if the target vCenter Server instance uses a custom HTTPS reverse proxy port.</p>
target	string or array	<p>The target ESXi host or DRS cluster on which you want to deploy the appliance.</p> <p>Important You must provide the name that is displayed in the vCenter Server inventory. For example, if the name of the target ESXi host is an IP address in the vCenter Server inventory, you cannot provide an FQDN.</p> <p>If the target ESXi host or DRS cluster is located in a folder or a structure of folders, the value must be a comma-separated list of strings or a comma-separated list as a single string. For example,</p> <pre data-bbox="590 655 1426 724">["parent_folder", "child_folder", "esxi-host.domain.com"]</pre> <p>or</p> <pre data-bbox="590 777 1426 846">"parent_folder, child_folder, esxi-host.domain.com"</pre> <p>If the target ESXi host is part of a cluster, use a comma-separated list of strings or a comma-separated list as a single string to provide the path. For example,</p> <pre data-bbox="590 934 1426 1024">["cluster_name", "esxi-host.domain.com"]</pre> <p>or</p> <pre data-bbox="590 1077 1426 1146">"cluster_name, esxi-host.domain.com"</pre> <p>Note The value is case-sensitive.</p>
vm.folder	string	Optional. The name of the VM folder to which to add the appliance.

Table 3-12. Configuration Parameters in the new.vcsa Section, appliance Subsection

Name	Type	Description
thin.disk.mode	Boolean	Set to true to deploy the appliance with thin virtual disks.
deployment.option	string	<p>The size of the appliance.</p> <ul style="list-style-type: none"> ■ Set to <code>tiny</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 10 hosts and 100 virtual machines with the default storage size. <p>Deploys an appliance with 2 CPUs, 10 GB of memory, and 250 GB of storage.</p> <ul style="list-style-type: none"> ■ Set to <code>tiny-lstorage</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 10 hosts and 100 virtual machines with the large storage size. <p>Deploys an appliance with 2 CPUs, 10 GB of memory, and 775 GB of storage.</p> <ul style="list-style-type: none"> ■ Set to <code>tiny-xlstorage</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 10 hosts and 100 virtual machines with the x-large storage size. <p>Deploys an appliance with 2 CPUs, 10 GB of memory, and 1650 GB of storage.</p> <ul style="list-style-type: none"> ■ Set to <code>small</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 100 hosts and 1,000 virtual machines with the default storage size. <p>Deploys an appliance with 4 CPUs, 16 GB of memory, and 290 GB of storage.</p> <ul style="list-style-type: none"> ■ Set to <code>small-lstorage</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 100 hosts and 1,000 virtual machines with the large storage size. <p>Deploys an appliance with 4 CPUs, 16 GB of memory, and 820 GB of storage.</p> <ul style="list-style-type: none"> ■ Set to <code>small-xlstorage</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 100 hosts and 1,000 virtual machines with the x-large storage size. <p>Deploys an appliance with 4 CPUs, 16 GB of memory, and 1700 GB of storage.</p> <ul style="list-style-type: none"> ■ Set to <code>medium</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 400 hosts and 4,000 virtual machines with the default storage size. <p>Deploys an appliance with 8 CPUs, 24 GB of memory, and 425 GB of storage.</p> <ul style="list-style-type: none"> ■ Set to <code>medium-lstorage</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 400 hosts and 4,000 virtual machines with the large storage size. <p>Deploys an appliance with 8 CPUs, 24 GB of memory, and 925 GB of storage.</p> <ul style="list-style-type: none"> ■ Set to <code>medium-xlstorage</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 400 hosts and 4,000 virtual machines with the x-large storage size. <p>Deploys an appliance with 8 CPUs, 24 GB of memory, and 1805 GB of storage.</p> <ul style="list-style-type: none"> ■ Set to <code>large</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 1,000 hosts and 10,000 virtual machines with the default storage size. <p>Deploys an appliance with 16 CPUs, 32 GB of memory, and 640 GB of storage.</p>

Table 3-12. Configuration Parameters in the new.vcsa Section, appliance Subsection (Continued)

Name	Type	Description
		<ul style="list-style-type: none"> ■ Set to <code>large-lstorage</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 1,000 hosts and 10,000 virtual machines with the large storage size. <p>Deploys an appliance with 16 CPUs, 32 GB of memory, and 990 GB of storage.</p>
		<ul style="list-style-type: none"> ■ Set to <code>large-xlstorage</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 1,000 hosts and 10,000 virtual machines with the x-large storage size. <p>Deploys an appliance with 16 CPUs, 32 GB of memory, and 1870 GB of storage.</p>
		<ul style="list-style-type: none"> ■ Set to <code>xlarge</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 2,000 hosts and 35,000 virtual machines with the default storage size. <p>Deploys an appliance with 24 CPUs, 48 GB of memory, and 980 GB of storage.</p>
		<ul style="list-style-type: none"> ■ Set to <code>xlarge-lstorage</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 2,000 hosts and 35,000 virtual machines with the large storage size. <p>Deploys an appliance with 24 CPUs, 48 GB of memory, and 1030 GB of storage.</p>
		<ul style="list-style-type: none"> ■ Set to <code>xlarge-xlstorage</code> if you want to deploy a vCenter Server Appliance with an embedded Platform Services Controller for up to 2,000 hosts and 35,000 virtual machines with the x-large storage size. <p>Deploys an appliance with 24 CPUs, 48 GB of memory, and 1910 GB of storage.</p>
		<ul style="list-style-type: none"> ■ Set to <code>management-tiny</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 10 hosts and 100 virtual machines with the default storage size. <p>Deploys an appliance with 2 CPUs, 10 GB of memory, and 250 GB of storage.</p>
		<ul style="list-style-type: none"> ■ Set to <code>management-tiny-lstorage</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 10 hosts and 100 virtual machines with the large storage size. <p>Deploys an appliance with 2 CPUs, 10 GB of memory, and 775 GB of storage.</p>
		<ul style="list-style-type: none"> ■ Set to <code>management-tiny-xlstorage</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 10 hosts and 100 virtual machines with the x-large storage size. <p>Deploys an appliance with 2 CPUs, 10 GB of memory, and 1650 GB of storage.</p>
		<ul style="list-style-type: none"> ■ Set to <code>management-small</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 100 hosts and 1,000 virtual machines with the default storage size. <p>Deploys an appliance with 4 CPUs, 16 GB of memory, and 290 GB of storage.</p>
		<ul style="list-style-type: none"> ■ Set to <code>management-small-lstorage</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 100 hosts and 1,000 virtual machines with the large storage size. <p>Deploys an appliance with 4 CPUs, 16 GB of memory, and 820 GB of storage.</p>

Table 3-12. Configuration Parameters in the new.vcsa Section, appliance Subsection (Continued)

Name	Type	Description
		<ul style="list-style-type: none"> ■ Set to <code>management-small-xlstorage</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 100 hosts and 1,000 virtual machines with the x-large storage size. Deploys an appliance with 4 CPUs, 16 GB of memory, and 1700 GB of storage. ■ Set to <code>management-medium</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 400 hosts and 4,000 virtual machines with the default storage size. Deploys an appliance with 8 CPUs, 24 GB of memory, and 425 GB of storage. ■ Set to <code>management-medium-lstorage</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 400 hosts and 4,000 virtual machines with the large storage size. Deploys an appliance with 8 CPUs, 24 GB of memory, and 925 GB of storage. ■ Set to <code>management-medium-xlstorage</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 400 hosts and 4,000 virtual machines with the x-large storage size. Deploys an appliance with 8 CPUs, 24 GB of memory, and 1805 GB of storage. ■ Set to <code>management-large</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 1,000 hosts and 10,000 virtual machines with the default storage size. Deploys an appliance with 16 CPUs, 32 GB of memory, and 640 GB of storage. ■ Set to <code>management-large-lstorage</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 1,000 hosts and 10,000 virtual machines with the large storage size. Deploys an appliance with 16 CPUs, 32 GB of memory, and 990 GB of storage. ■ Set to <code>management-large-xlstorage</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 1,000 hosts and 10,000 virtual machines with the x-large storage size. Deploys an appliance with 16 CPUs, 32 GB of memory, and 1870 GB of storage. ■ Set to <code>management-xlarge</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 2,000 hosts and 35,000 virtual machines with the default storage size. Deploys an appliance with 24 CPUs, 48 GB of memory, and 980 GB of storage. ■ Set to <code>management-xlarge-lstorage</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 2,000 hosts and 35,000 virtual machines with the large storage size. Deploys an appliance with 24 CPUs, 48 GB of memory, and 1030 GB of storage. ■ Set to <code>management-xlarge-xlstorage</code> if you want to deploy a vCenter Server Appliance with an external Platform Services Controller for up to 2,000 hosts and 35,000 virtual machines with the x-large storage size. Deploys an appliance with 24 CPUs, 48 GB of memory, and 1910 GB of storage. ■ Set to <code>infrastructureif</code> if you want to deploy a Platform Services Controller appliance. Deploys an appliance with 2 CPUs, 4 GB of memory, and 60 GB of storage.

Table 3-12. Configuration Parameters in the new.vcsa Section, appliance Subsection (Continued)

Name	Type	Description
image	string	Optional. A local file path or URL to the vCenter Server Appliance installation package. By default the installer uses the installation package that is included in the ISO file, in the vcsa folder.
name	string	The VM name for the appliance. Must contain only ASCII characters except a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.
ovftool.path	string	Optional. A local file path to the OVF Tool executable file. By default the installer uses the OVF Tool instance that is included in the ISO file, in the vcsa/ovftool folder.

Table 3-13. Configuration Parameters in the new.vcsa Section, network Subsection

Name	Type	Description
ip.family	string	IP version for the network of the appliance. Set to ipv4 or ipv6.
mode	string	IP assignment for the network of the appliance. Set to static or dhcp.
ip	string	IP address for the appliance. Required only if you use static assignment, that is, if you set the mode parameter to static. You must set an IPv4 or IPv6 address that corresponds to the network IP version, that is, to the value of the ip.family parameter. An IPv4 address must comply with the RFC 790 guidelines. An IPv6 address must comply with the RFC 2373 guidelines.
dns.servers	string or array	IP addresses of one or more DNS servers. To set more than one DNS server, use a comma-separated list of strings or a comma-separated list as a single string to provide the path. For example, <pre>["x.y.z.a", "x.y.z.b"]</pre> or <pre>"x.y.z.a, x.y.z.b"</pre> Required only if you use static assignment, that is, if you set the mode parameter to static.
prefix	string	Network prefix length. Use only if the mode parameter is set to static. Remove if the mode parameter is set to dhcp. The network prefix length is the number of bits that are set in the subnet mask. For example, if the subnet mask is 255.255.255.0, there are 24 bits in the binary version of the prefix length, so the network prefix length is 24. For IPv4 version, the value must be between 0 and 32. For IPv6 version, the value must be between 0 and 128.

Table 3-13. Configuration Parameters in the new.vcsa Section, network Subsection (Continued)

Name	Type	Description
gateway	string	IP address of the default gateway. For IPv6 version, the value can be default.
system.name	string	Primary network identity. Can be an IP address or FQDN, preferably FQDN. You cannot change the value of this parameter after the deployment. The FQDN and dotted-decimal numbers must comply with the RFC 1123 guidelines.

Table 3-14. Configuration Parameters in the new.vcsa Section, os Subsection

Name	Type	Description
password	string	The password for the root user of the appliance operating system. The password must contain between 8 and 20 characters, at least one uppercase letter, at least one lowercase letter, at least one number, and at least one special character, for example, a dollar sign (\$), hash key (#), at sign (@), period (.), or exclamation mark (!). All characters must be lower ASCII characters without spaces.
ntp.servers	string or array	Optional. Host names or IP addresses of one or more NTP servers for time synchronization. To set more than one NTP server, use a comma-separated list of strings or a comma-separated list as a single string to provide the path. For example, <pre>["x.y.z.a", "x.y.z.b"]</pre> or <pre>"x.y.z.a, x.y.z.b"</pre>
ssh.enable	Boolean	Set to true to enable SSH administrator login to the appliance. Note vCenter Server Appliance high availability requires remote SSH access to the appliance.
time.tools-sync	Boolean	Optional. Set to true to deploy the appliance with the VMware Tools time synchronization. VMware Tools synchronizes the time of the appliance with the time of the ESxi host. Ignored if you set NTP servers for time synchronization, that is, if you set the ntp.servers parameter.

Table 3-15. Configuration Parameters in the new.vcsa Section, sso Subsection

Name	Type	Description
password	string	<p>Password of the vCenter Single Sign-On administrator user, <code>administrator@your_domain_name</code>.</p> <ul style="list-style-type: none"> If you are deploying a vCenter Server Appliance with an embedded Platform Services Controller or a Platform Services Controller appliance as the first instance in a new vCenter Single Sign-On domain, you must set the password for the vCenter Single Sign-On administrator user. <p>The password must contain between 8 and 20 characters, at least one uppercase letter, at least one lowercase letter, at least one number, and at least one special character, for example, a dollar sign (\$), hash key (#), at sign (@), period (.), or exclamation mark (!). All characters must be ASCII characters.</p> <ul style="list-style-type: none"> If you are deploying a Platform Services Controller appliance as a replication partner in an existing vCenter Single Sign-On domain, you must provide the password of the vCenter Single Sign-On administrator user of the partner Platform Services Controller. If you are deploying a vCenter Server Appliance with an external Platform Services Controller, you must provide the password of the vCenter Single Sign-On administrator user of the external Platform Services Controller.
domain-name	string	<p>vCenter Single Sign-On domain name, for example, <code>vsphere.local</code>.</p> <ul style="list-style-type: none"> If you are deploying a vCenter Server Appliance with an embedded Platform Services Controller or a Platform Services Controller appliance as the first instance in a new vCenter Single Sign-On domain, you must set the name for the new vCenter Single Sign-On domain. If you are deploying a vCenter Server Appliance with an external Platform Services Controller or a Platform Services Controller appliance as a replication partner in an existing vCenter Single Sign-On domain, you must provide the name of the existing vCenter Single Sign-On domain.
first-instance	Boolean	<p>Required only if you are deploying a Platform Services Controller appliance.</p> <p>The default value is <code>true</code>.</p> <p>Set to <code>false</code> if you want to join the Platform Services Controller appliance to an existing vCenter Single Sign-On domain.</p> <p>Joined Platform Services Controller instances replicate their infrastructure data and enable Enhanced Linked Mode. For information about managing the Platform Services Controller services, see <i>Platform Services Controller Administration</i>.</p>
platform.services.controller	string	<p>The system name of the external Platform Services Controller.</p> <p>Required only if you are deploying a vCenter Server Appliance with an external Platform Services Controller.</p>
replication-partner-hostname	string	<p>The system name of the partner Platform Services Controller.</p> <p>Required only if you are deploying a Platform Services Controller appliance as a replication partner in an existing vCenter Single Sign-On domain.</p>
sso.port	integer	<p>The HTTPS reverse proxy port of the partner Platform Services Controller.</p> <p>The default port is 443. Use only if you the partner Platform Services Controller uses a custom HTTPS reverse proxy port.</p>
site-name	string	<p>vCenter Single Sign-On site name.</p> <p>Required only if you are deploying a vCenter Server Appliance with an embedded Platform Services Controller or a Platform Services Controller appliance.</p>

Configuration Parameters in the `ceip` Section

Table 3-16. Configuration Parameters in the `ceip` Section, `settings` Subsection

Name	Type	Description
<code>ceip.enabled</code>	Boolean	Set to <code>true</code> to join the CEIP for this appliance.

Deploy a vCenter Server Appliance or Platform Services Controller Appliance by Using the CLI

You can use the CLI installer to perform an unattended deployment of a vCenter Server Appliance or Platform Services Controller appliance. You must run the CLI deployment from a Windows, Linux, or Mac machine that is in the network on which you want to deploy the appliance.

Prerequisites

- See [Prerequisites for Deploying the vCenter Server Appliance or Platform Services Controller Appliance](#).
- [Prepare Your JSON Configuration File for CLI Deployment](#).
- Review [Syntax of the CLI Deployment Command](#).
- Verify that the user name with which you are logged in to your client machine, the path to the vCenter Server Appliance installer, the path to your JSON configuration file, and the string values in your JSON configuration file contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

Procedure

- 1 Navigate to the `vcsa-cli-installer` subdirectory for your operating system.
 - If you are running the deployment on Windows OS, navigate to the `vcsa-cli-installer\win32` directory.
 - If you are running the deployment on Linux OS, navigate to the `vcsa-cli-installer/lin64` directory.
 - If you are running the deployment on Mac OS, navigate to the `vcsa-cli-installer/mac` directory.
- 2 (Optional) Run a pre-deployment check without deploying the appliance to verify that you prepared the deployment template correctly.

```
vcsa-deploy install --verify-only path_to_the_json_file
```

3 Run the deployment command.

```
vcsa-deploy install --accept-eula --acknowledge-ceip optional_arguments path_to_the_json_file
```

Use *optional_arguments* to enter space-separated arguments to set additional execution parameters of the deployment command.

For example, you can set the location of the log and other output files that the installer generates.

```
vcsa-deploy install --accept-eula --acknowledge-ceip --log-dir=path_to_the_location
path_to_the_json_file
```

Syntax of the CLI Deployment Command

You can use command arguments to set the execution parameters of the deployment command.

You can add a space-separated list of arguments to the CLI deployment command.

```
vcsa-deploy install path_to_the_json_file list_of_arguments
```

Argument	Description
--accept-eula	Accepts the end-user license agreement. Required for executing the deployment command.
--acknowledge-ceip	Confirms your acknowledgement of your VMware Customer Experience Improvement Program (CEIP) participation. Required if the <code>ceip.enabled</code> parameter is set to <code>true</code> in the JSON deployment template.
-v, --verbose	Adds debug information to the console output.
-t, --terse	Hides the console output. Displays only warning and error messages.
--log-dir <i>LOG_DIR</i>	Sets the location of the log and other output files.
--skip-ovftool-verification	Performs basic verification of the configuration parameters in the JSON file and deploys the appliance. Does not perform verification of the OVF Tool parameters.
--no-esx-ssl-verify	Skips the SSL verification for ESXi connections. Important Avoid using this option because it might cause problems during deployment or after deployment because of not validated identity of the target ESXi host.
--deployment-target-ssl-thumbprint <i>TARGET_THUMBPRINT</i>	Thumbprint to pass to the OVF Tool for verifying the target ESXi host or vCenter Server instance on which you want to deploy the appliance.
--pause-on-warnings	Pauses and waits for acknowledgment of warnings.
--verify-only	Performs basic verification of the configuration parameters in the JSON file and verification of the OVF Tool parameters. Does not deploy the appliance.

Argument	Description
<code>--sso-ssl-thumbprint SSL-SHA1-THUMBPRINT</code>	Validates server certificate against the supplied SHA1 thumbprint.
<code>-h, --help</code>	Displays the help message for the <code>vcsa-deploy install</code> command.
<code>--template-help</code>	Displays the help message for the use of configuration parameters in the JSON deployment file.

After the execution finishes, you can get the exit code of the command.

Exit Code	Description
0	Command ran successfully
1	Runtime error
2	Validation error
3	Template error

Installing vCenter Server and Platform Services Controller on Windows

4

You can install vCenter Server with an embedded or external Platform Services Controller on a Microsoft Windows virtual machine or physical server to manage your vSphere environment.

Before you install vCenter Server, download the installer ISO file and mount it to the Windows host machine on which you want to perform the installation, and then start the installation wizard.

Windows installations of vCenter Server can use either the embedded PostgreSQL database or an external database. Before installing vCenter Server that uses an external database, you must prepare your database. See [Preparing vCenter Server Databases for Install](#).

For information about the vCenter Server requirements, see [vCenter Server for Windows Requirements](#).

For information about the inputs that are required during the installation of vCenter Server, see [Required Information for Installing vCenter Server or Platform Services Controller on Windows](#).

Important For topologies with external Platform Services Controller instances, you must install the replicating Platform Services Controller instances in a sequence. After the successful deployment of all Platform Services Controller instances in the domain, you can perform concurrent installations of multiple vCenter Server instances that point to a common external Platform Services Controller instance.

After you install vCenter Server, only the user `administrator@your_domain_name` has the privileges to log in to the vCenter Server system.

The `administrator@your_domain_name` user can perform the following tasks:

- Add an identity source in which additional users and groups are defined in vCenter Single Sign-On.
- Assign roles to users and groups to give them privileges.

For information about adding identity sources and giving permissions to the users and groups, see *Platform Services Controller Administration*.

Starting with vSphere 6.5, vCenter Server supports mixed IPv4 and IPv6 environment. If you want to set up the vCenter Server instance to use an IPv6 address version, use the fully qualified domain name (FQDN) or host name of the host machine. To set up an IPv4 address, the best practice is to use the FQDN or host name of the host machine, because the IP address can change if assigned by DHCP.

This section includes the following topics:

- [vCenter Server for Windows Requirements](#)
- [Preparing for Installing vCenter Server and Platform Services Controller on Windows](#)
- [Required Information for Installing vCenter Server or Platform Services Controller on Windows](#)
- [Installing vCenter Server and Platform Services Controller on Windows](#)

vCenter Server for Windows Requirements

To install vCenter Server on a Windows virtual machine or physical server, your system must meet specific hardware and software requirements.

- Synchronize the clocks of the virtual machines on which you plan to install vCenter Server and the Platform Services Controller. See [Synchronizing Clocks on the vSphere Network](#).
- Verify that the DNS name of the virtual machine or physical server matches the actual full computer name.
- Verify that the host name of the virtual machine or physical server on which you are installing or upgrading vCenter Server complies with RFC 1123 guidelines.
- Verify that the system on which you are installing vCenter Server is not an Active Directory domain controller.
- If you plan to use a user account other than the Local System account in which to run your vCenter Server service, verify that the user account has the following permissions:
 - **Member of the Administrators group**
 - **Log on as a service**
 - **Act as part of the operating system (if the user is a domain user)**

Note Starting with vSphere 6.5, the vCenter Server services run as child processes of the VMware Service Lifecycle Manager service.

- Verify that the local policy of the virtual machine or physical server on which you are installing or upgrading vCenter Server allows assigning **Log on as a batch job** rights to new local users.

Note Starting with vSphere 6.5, some vCenter Server processes use separate local users that are automatically created and added to the local security policy **Log on as a batch job**. Such new local users are cm, content-library, eam, imagebuilder, mbc, netdumper, perfcharts, rbd, vapiEndpoint, vmware-vpostgres, vsan-health, vsm, vsphere-client, and vsphere-ui.

- If the system that you use for your vCenter Server installation belongs to a workgroup rather than a domain, not all functionality is available to vCenter Server. If assigned to a workgroup, the vCenter Server system is not able to discover all domains and systems available on the network when using some features. Your host machine must be connected to a domain if you want to add Active Directory identity sources after the installation.

- Verify that the LOCAL SERVICE account has read permission on the folder in which vCenter Server is installed and on the HKLM registry.
- Verify that the connection between the virtual machine or physical server and the domain controller is working.

Pre-Install Checks for vCenter Server and Platform Services Controller on Windows

When you install or upgrade vCenter Server and Platform Services Controller on Windows, the installer does a pre-check, for example, to verify that enough space is available on the virtual machine or physical server where you are installing or upgrading vCenter Server, and verifies that the external database, if any, can be successfully accessed.

When you install Platform Services Controller as an embedded or external instance, vCenter Single Sign-On is installed as part of the Platform Services Controller. During the installation of an external Platform Services Controller, the installer provides you with the option to join an existing vCenter Single Sign-On server domain. During the installation of vCenter Server with an external Platform Services Controller, the installer prompts you to join an existing vCenter Single Sign-On server domain. When you provide the information about the vCenter Single Sign-On service, the installer uses the administrator account to check the host name and password, to verify that the details of the vCenter Single Sign-On server you provided can be authenticated before proceeding with the installation process.

The pre-install checker performs checks for the following aspects of the environment:

- Windows version
- Minimum processor requirements
- Minimum memory requirements
- Minimum disk space requirements
- Permissions on the selected install and data directory
- Internal and external port availability
- External database version
- External database connectivity
- Administrator privileges on the Windows machine
- Any credentials that you enter

For information about the minimum storage requirements, see [Storage Requirements for vCenter Server and Platform Services Controller on Windows](#). For information about the minimum hardware requirements, see [Hardware Requirements for vCenter Server and Platform Services Controller on Windows](#).

Hardware Requirements for vCenter Server and Platform Services Controller on Windows

When you install vCenter Server or Platform Services Controller on a virtual machine or physical server running Microsoft Windows, your system must meet specific hardware requirements.

You can install vCenter Server and the Platform Services Controller on the same virtual machine or physical server or on different virtual machines or physical servers. When you install vCenter Server with an embedded Platform Services Controller, you install vCenter Server and the Platform Services Controller on the same virtual machine or physical server. When you install the vCenter Server with an external Platform Services Controller, first install the Platform Services Controller that contains all of the required services on one virtual machine or physical server, and then install vCenter Server and the vCenter Server components on another virtual machine or physical server.

Note Installing vCenter Server on a network drive or USB flash drive is not supported.

Table 4-1. Minimum Recommended Hardware Requirements for Installing vCenter Server and Platform Services Controller on Windows

	Platform Services Controller	vCenter Server with an Embedded or External Platform Services Controller for a Tiny Environment (up to 10 Hosts, 100 Virtual Machines)	vCenter Server with an Embedded or External Platform Services Controller for a Small Environment (up to 100 Hosts, 1000 Virtual Machines)	vCenter Server with an Embedded or External Platform Services Controller for a Medium Environment (up to 400 Hosts, 4,000 Virtual Machines)	vCenter Server with an Embedded or External Platform Services Controller for a Large Environment (up to 1,000 Hosts, 10,000 Virtual Machines)	vCenter Server with an Embedded or External Platform Services Controller for an X-Large Environment (up to 2,000 Hosts, 35,000 Virtual Machines)
Number of CPUs	2	2	4	8	16	24
Memory	4 GB RAM	10 GB RAM	16 GB RAM	24 GB RAM	32 GB RAM	48 GB RAM

Note If you want to add an ESXi host with more than 512 LUNs and 2,048 paths to the vCenter Server inventory, your vCenter Server instance must be suitable for a large or x-large environment.

For the hardware requirements of your database, see the database documentation. The database requirements are in addition to the vCenter Server requirements if the database and vCenter Server run on the same machine.

Storage Requirements for vCenter Server and Platform Services Controller on Windows

When you install vCenter Server, your system must meet minimum storage requirements.

The storage requirements per folder depend on the deployment model that you decide to install. During installation, you can select a folder other than the default C:\Program Files\VMware folder to install vCenter Server and the Platform Services Controller. You can also select a folder other than the default C:\ProgramData\VMware\vCenterServer\ in which to store data.

Table 4-2. vCenter Server Minimum Storage Requirements Depending On the Deployment Model

Default Folder	vCenter Server with an Embedded Platform Services Controller	vCenter Server with an External Platform Services Controller	External Platform Services Controller
Program Files	6 GB	6 GB	1 GB
ProgramData	8 GB	8 GB	2 GB
System folder (to cache the MSI installer)	3 GB	3 GB	1 GB

Software Requirements for vCenter Server and Platform Services Controller on Windows

Verify that your operating system supports vCenter Server.

vCenter Server requires a 64-bit operating system, and the 64-bit system DSN is required for vCenter Server to connect to the external database.

The earliest Windows Server version that vCenter Server supports is Windows Server 2008 SP2. Your Windows Server must have the latest updates and patches installed. For a full list of supported operating systems, see <http://kb.vmware.com/kb/2091273>.

Database Requirements for vCenter Server on Windows

vCenter Server requires a database to store and organize server data.

Each vCenter Server instance must have its own database. For environments with up to 20 hosts and 200 virtual machines, you can use the bundled PostgreSQL database that the vCenter Server installer can install and set up for you during the vCenter Server installation. A larger installation requires a supported external database for the size of the environment.

During vCenter Server installation you must select to install the embedded database or point the vCenter Server system to any existing supported database. vCenter Server supports Oracle and Microsoft SQL Server databases.

For information about supported database server versions, see the VMware Product Interoperability Matrix at http://www.vmware.com/resources/compatibility/sim/interop_matrix.php.

Required Ports for vCenter Server and Platform Services Controller

The vCenter Server system, both on Windows and in the appliance, must be able to send data to every managed host and receive data from the vSphere Web Client and the Platform Services Controller services. To enable migration and provisioning activities between managed hosts, the source and destination hosts must be able to receive data from each other.

If a port is in use or is blacklisted, the vCenter Server installer displays an error message. You must use another port number to proceed with the installation. There are internal ports that are used only for inter-process communication.

VMware uses designated ports for communication. Additionally, the managed hosts monitor designated ports for data from vCenter Server. If a built-in firewall exists between any of these elements, the installer opens the ports during the installation or upgrade process. For custom firewalls, you must manually open the required ports. If you have a firewall between two managed hosts and you want to perform source or target activities, such as migration or cloning, you must configure a means for the managed hosts to receive data.

Note In Microsoft Windows Server 2008 and later, firewall is enabled by default.

Table 4-3. Ports Required for Communication Between Components

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
22	TCP/UDP	System port for SSHD.	Appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No
53		DNS service	Windows installations and appliance deployments of Platform Services Controller	No

Table 4-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
80	TCP	<p>vCenter Server requires port 80 for direct HTTP connections. Port 80 redirects requests to HTTPS port 443. This redirection is useful if you accidentally use http://server instead of https://server. WS-Management (also requires port 443 to be open).</p> <p>If you use a Microsoft SQL database that is stored on the same virtual machine or physical server as the vCenter Server, port 80 is used by the SQL Reporting Service. When you install or upgrade vCenter Server, the installer prompts you to change the HTTP port for vCenter Server. Change the vCenter Server HTTP port to a custom value to ensure a successful installation or upgrade.</p> <p>Important You can change this port number during the vCenter Server and Platform Services Controller installations on Windows.</p>	<p>Windows installations and appliance deployments of</p> <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No
88	TCP	Active Directory server. This port must be open for host to join Active Directory. If you use native Active Directory, the port must be open on both vCenter Server and Platform Services Controller.	Windows installations and appliance deployments of Platform Services Controller	No
389	TCP/UDP	<p>This port must be open on the local and all remote instances of vCenter Server. This is the LDAP port number for the Directory Services for the vCenter Server group. If another service is running on this port, it might be preferable to remove it or change its port to a different port. You can run the LDAP service on any port from 1025 through 65535.</p> <p>If this instance is serving as the Microsoft Windows Active Directory, change the port number from 389 to an available port from 1025 through 65535.</p>	Windows installations and appliance deployments of Platform Services Controller	<ul style="list-style-type: none"> ■ vCenter Server to Platform Services Controller ■ Platform Services Controller to Platform Services Controller

Table 4-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
443	TCP	<p>The default port that the vCenter Server system uses to listen for connections from the vSphere Web Client. To enable the vCenter Server system to receive data from the vSphere Web Client, open port 443 in the firewall.</p> <p>The vCenter Server system also uses port 443 to monitor data transfer from SDK clients.</p> <p>This port is also used for the following services:</p> <ul style="list-style-type: none"> ■ WS-Management (also requires port 80 to be open) ■ Third-party network management client connections to vCenter Server ■ Third-party network management clients access to hosts <p>Important You can change this port number during the vCenter Server and Platform Services Controller installations on Windows.</p>	<p>Windows installations and appliance deployments of</p> <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	<ul style="list-style-type: none"> ■ vCenter Server to vCenter Server ■ vCenter Server to Platform Services Controller ■ Platform Services Controller to vCenter Server
514	TCP/UDP	<p>vSphere Syslog Collector port for vCenter Server on Windows and vSphere Syslog Service port for vCenter Server Appliance</p> <p>Important You can change this port number during the vCenter Server and Platform Services Controller installations on Windows.</p>	<p>Windows installations and appliance deployments of</p> <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No
636	TCP	<p>vCenter Single Sign-On LDAPS</p> <p>For backward compatibility with vSphere 6.0 only.</p>	<p>Windows installations and appliance deployments of Platform Services Controller</p>	<p>During upgrade from vSphere 6.0 only. vCenter Server 6.0 to Platform Services Controller 6.5</p>

Table 4-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
902	TCP/UDP	<p>The default port that the vCenter Server system uses to send data to managed hosts. Managed hosts also send a regular heartbeat over UDP port 902 to the vCenter Server system. This port must not be blocked by firewalls between the server and the hosts or between hosts.</p> <p>Port 902 must not be blocked between the VMware Host Client and the hosts. The VMware Host Client uses this port to display virtual machine consoles</p> <p>Important You can change this port number during the vCenter Server installations on Windows.</p>	Windows installations and appliance deployments of vCenter Server	No
1514	TCP/UDP	<p>vSphere Syslog Collector TLS port for vCenter Server on Windows and vSphere Syslog Service TLS port for vCenter Server Appliance</p> <p>Important You can change this port number during the vCenter Server and Platform Services Controller installations on Windows.</p>	<p>Windows installations and appliance deployments of</p> <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No
2012	TCP	Control interface RPC for vCenter Single Sign-On	Windows installations and appliance deployments of Platform Services Controller	<ul style="list-style-type: none"> ■ vCenter Server to Platform Services Controller ■ Platform Services Controller to vCenter Server ■ Platform Services Controller to Platform Services Controller
2014	TCP	<p>RPC port for all VMCA (VMware Certificate Authority) APIs</p> <p>Important You can change this port number during the Platform Services Controller installations on Windows.</p>	Windows installations and appliance deployments of Platform Services Controller	<ul style="list-style-type: none"> ■ vCenter Server to Platform Services Controller ■ Platform Services Controller to vCenter Server
2015	TCP	DNS management	Windows installations and appliance deployments of Platform Services Controller	Platform Services Controller to Platform Services Controller

Table 4-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
2020	TCP/UDP	Authentication framework management Important You can change this port number during the vCenter Server and Platform Services Controller installations on Windows.	Windows installations and appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	<ul style="list-style-type: none"> ■ vCenter Server to Platform Services Controller ■ Platform Services Controller to vCenter Server
5480	TCP	Appliance Management Interface Open endpoint serving all HTTPS, XMLRPS and JSON-RPC requests over HTTPS.	Appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No
6500	TCP/UDP	ESXi Dump Collector port Important You can change this port number during the vCenter Server installations on Windows.	Windows installations and appliance deployments of vCenter Server	No
6501	TCP	Auto Deploy service Important You can change this port number during the vCenter Server installations on Windows.	Windows installations and appliance deployments of vCenter Server	No
6502	TCP	Auto Deploy management Important You can change this port number during the vCenter Server installations on Windows.	Windows installations and appliance deployments of vCenter Server	No
7080, 12721	TCP	Secure Token Service Note Internal ports	Windows installations and appliance deployments of Platform Services Controller	No
7081	TCP	VMware Platform Services Controller Web Client Note Internal port	Windows installations and appliance deployments of Platform Services Controller	No
8200, 8201, 8300, 8301	TCP	Appliance management Note Internal ports	Appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	No

Table 4-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
7444	TCP	Secure Token Service For backward compatibility with vSphere 5.5 only.	Windows installations and appliance deployments of Platform Services Controller	During upgrade from vSphere 5.5 only. <ul style="list-style-type: none"> ■ vCenter Server 5.5 to Platform Services Controller 6.5 ■ Platform Services Controller 6.5 to vCenter Server 5.5
8084	TCP	vSphere Update Manager SOAP port The port used by vSphere Update Manager client plug-in to connect to the vSphere Update Manager SOAP server.	Appliance deployments of vCenter Server	No
9084	TCP	vSphere Update Manager Web Server Port The HTTP port used by ESXi hosts to access host patch files from vSphere Update Manager server.	Appliance deployments of vCenter Server	No
9087	TCP	vSphere Update Manager Web SSL Port The HTTPS port used by vSphere Update Manager client plug-in to upload host upgrade files to vSphere Update Manager server.	Appliance deployments of vCenter Server	No
9123	TCP	Migration Assistant port Only when you run the Migration Assistant on the source Windows installation. The Migration Assistant lets you migrate Windows installations of vCenter Server and Platform Services Controller to appliances.	Windows installations and appliance deployments of <ul style="list-style-type: none"> ■ vCenter Server ■ Platform Services Controller 	During migration only. <ul style="list-style-type: none"> ■ Source vCenter Server 5.5 or 6.5 to target vCenter Server Appliance 6.5 ■ Source vCenter Single Sign-On 5.5 to target Platform Services Controller appliance 6.5 ■ Source Platform Services Controller 5.5 to target Platform Services Controller appliance 6.5
9443	TCP	vSphere Web Client HTTPS	Windows installations and appliance deployments of vCenter Server	No

Table 4-3. Ports Required for Communication Between Components (Continued)

Port	Protocol	Description	Required for	Used for Node-to-Node Communication
11711	TCP	vCenter Single Sign-On LDAP For backward compatibility with vSphere 5.5 only.	Windows installations and appliance deployments of Platform Services Controller	During upgrade from vSphere 5.5 only. vCenter Single Sign-On 5.5 to Platform Services Controller 6.5
11712	TCP	vCenter Single Sign-On LDAPS For backward compatibility with vSphere 5.5 only.	Windows installations and appliance deployments of Platform Services Controller	During upgrade from vSphere 5.5 only. vCenter Single Sign-On 5.5 to Platform Services Controller 6.5

To configure the vCenter Server system to use a different port to receive vSphere Web Client data, see the *vCenter Server and Host Management* documentation.

For more information about firewall configuration, see the *vSphere Security* documentation.

DNS Requirements for vCenter Server and Platform Services Controller on Windows

You install or upgrade vCenter Server, like any other network server, on a host machine with a fixed IP address and well-known DNS name, so that clients can reliably access the service.

Assign a static IP address and host name to the Windows server that will host the vCenter Server system. This IP address must have a valid (internal) domain name system (DNS) registration. When you install vCenter Server and the Platform Services Controller, you must provide the fully qualified domain name (FQDN) or the static IP of the host machine on which you are performing the install or upgrade. The recommendation is to use the FQDN.

Ensure that DNS reverse lookup returns an FQDN when queried with the IP address of the host machine on which vCenter Server is installed. When you install or upgrade vCenter Server, the installation or upgrade of the Web server component that supports the vSphere Web Client fails if the installer cannot look up the fully qualified domain name of the vCenter Server host machine from its IP address. Reverse lookup is implemented using PTR records.

If you plan to use an FQDN for the virtual machine or physical server, you must verify that the FQDN is resolvable.

You can use the `nslookup` command to verify that the DNS reverse lookup service returns an FQDN when queried with the IP address and to verify that the FQDN is resolvable.

```
nslookup -nosearch -nodefname FQDN_or_IP_address
```

If you use DHCP instead of a static IP address for vCenter Server, make sure that the vCenter Server computer name is updated in the domain name service (DNS). If you can ping the computer name, the name is updated in DNS.

Ensure that the ESXi host management interface has a valid DNS resolution from the vCenter Server and all vSphere Web Client instances. Ensure that the vCenter Server has a valid DNS resolution from all ESXi hosts and all vSphere Web Clients.

vSphere Web Client Software Requirements

Make sure that your browser supports the vSphere Web Client.

The vSphere Web Client 6.5 requires Adobe Flash Player v. 16 to 23. For best performance and security fixes, use Adobe Flash Player 23.

VMware has tested and supports the following guest operating systems and browser versions for the vSphere Web Client. For best performance, use Google Chrome.

Table 4-4. Supported Guest Operating Systems and Minimum Browser Versions for the vSphere Web Client

Operating system	Browser
Windows	Microsoft Internet Explorer v. 10.0.19 and later.
	Mozilla Firefox v. 39 and later.
	Google Chrome v. 34 and later.
Mac OS	Mozilla Firefox v. 39 and later.
	Google Chrome v. 34 and later.

Preparing for Installing vCenter Server and Platform Services Controller on Windows

Before you install vCenter Server or Platform Services Controller, you must download the vCenter Server installer ISO file and mount it to the Windows virtual machine or physical server on which you want to install vCenter Server or Platform Services Controller.

If you plan to use an external vCenter Server database, before you install vCenter Server, you must set up the database.

Download the vCenter Server Installer for Windows

Download the .iso installer for vCenter Server for Windows and the associated vCenter Server components and support tools.

Prerequisites

Create a My VMware account at <https://my.vmware.com/web/vmware/>.

Procedure

- 1 Download the vCenter Server installer from the VMware Web site at <https://my.vmware.com/web/vmware/downloads>.
vCenter Server is part of VMware vCloud Suite and VMware vSphere, listed under Datacenter & Cloud Infrastructure.
- 2 Confirm that the md5sum is correct.
See the VMware Web site topic Using MD5 Checksums at <http://www.vmware.com/download/md5.html>.
- 3 Mount the ISO image to the Windows virtual machine or physical server on which you want to install vCenter Server for Windows.

Preparing vCenter Server Databases for Install

vCenter Server requires a database to store and organize server data. For vCenter Server on Windows, you can either use the bundled PostgreSQL database that can be installed and configured together with vCenter Server, or you can set up an external database prior to installing vCenter Server.

vCenter Server for Windows supports Oracle and Microsoft SQL Server as external databases.

You can configure an external database manually or by using a script. In addition, the data source name user must have a specific list of permissions.

The database passwords are stored in clear text on the Windows virtual machine or physical host on which you install vCenter Server and in the vCenter Server Appliance. The files containing the passwords are protected by using the operating system protection, that is, you must be a Windows local administrator or a Linux root user to access and read these files.

vCenter Server instances cannot share the same database schema. Multiple vCenter Server databases can reside on the same database server, or they can be separated across multiple database servers. For Oracle databases, which have the concept of schema objects, you can run multiple vCenter Server instances in a single database server if you have a different schema owner for each vCenter Server instance. You can also use a dedicated Oracle database server for each vCenter Server instance.

You cannot install vCenter Server and point to an older external vCenter Server database. You can upgrade the old vCenter Server database to the latest version only by upgrading the vCenter Server instance connected to that database. For information about upgrading vCenter Server, see *vSphere Upgrade*.

vCenter Server Database Configuration Notes

After you select a supported database type, make sure you understand any special configuration requirements.

[Table 4-5](#) is not a complete list of databases supported with vCenter Server for Windows. For information about specific database versions and service pack configurations supported with vCenter Server, see the [VMware Product Interoperability Matrixes](#). Only special database configuration notes not listed in the Product Interoperability Matrixes are provided in [Table 4-5](#).

vCenter Server databases require a UTF code set.

Contact your DBA for the appropriate database credentials.

Table 4-5. Configuration Notes for Databases Supported with vCenter Server

Database Type	Configuration Notes
Embedded PostgreSQL	For vCenter Server 6.5, the bundled PostgreSQL database is suitable for environments with up to 20 hosts and 200 virtual machines. Important If you use the embedded PostgreSQL database, uninstalling vCenter Server on Windows, uninstalls the embedded database, and all data is lost.
Microsoft SQL Server 2008 R2 SP2 or higher	Ensure that the machine has a valid ODBC DSN entry.
Microsoft SQL Server 2012	Ensure that the machine has a valid ODBC DSN entry.
Microsoft SQL Server 2014	Ensure that the machine has a valid ODBC DSN entry.
Oracle 11g and Oracle 12c	Ensure that the machine has a valid ODBC DSN entry. After you complete the vCenter Server installation, apply the latest patch to the Oracle client and server.

Configure Microsoft SQL Server Databases

To use a Microsoft SQL database for your vCenter Server repository, configure your database to work with vCenter Server.

You can install and configure the Microsoft SQL Server database on the same machine on which you plan to install vCenter Server. You can install and configure the Microsoft SQL Server database on a separate machine.

Procedure

1 Prepare the vCenter Server SQL Server Database

You first create a database and user for vCenter Server. Then you assign permissions to the vCenter Server database user either by using the existing dbo schema and dbo_owner role or by creating custom database schema and roles.

2 (Optional) Use a Script to Create Microsoft SQL Server Database Objects Manually

This topic describes how to create database objects manually instead of letting the vCenter Server installer create the data objects automatically.

3 Configure a SQL Server ODBC Connection

After you create and configure a SQL Server database and user for vCenter Server, you must create a 64-bit DSN on the machine on which you plan to install vCenter Server. During the vCenter Server installation, you use the DSN to establish a connection between vCenter Server and the database.

4 Configure Microsoft SQL Server TCP/IP for JDBC

If the Microsoft SQL Server database has TCP/IP disabled and the dynamic ports are not set, the JDBC connection remains closed. The closed connection causes the vCenter Server statistics to malfunction. You can configure the server TCP/IP for JDBC.

Prepare the vCenter Server SQL Server Database

You first create a database and user for vCenter Server. Then you assign permissions to the vCenter Server database user either by using the existing dbo schema and dbo_owner role or by creating custom database schema and roles.

Prerequisites

Log in to the Microsoft SQL Server Management Studio as the sysadmin (SA) or a user account with sysadmin privileges.

Prepare the vCenter Server Database by Using the dbo Schema and the db_owner Database Role

The simplest way to assign permissions for a vCenter Server database user is through the database role db_owner.

You must first create a database and user for vCenter Server. Then you can use the existing db_owner database role and let the vCenter Server installer create the default dbo schema that assigns database user permissions to that role. You must also enable database monitoring for the user before you install vCenter Server. See [Database Permission Requirements for vCenter Server](#).

To perform the following procedure, you can either use the graphical user interface or run scripts. The vCenter Server installer package contains example scripts in the vCenter-Server\dbschema\DB_and_schema_creation_scripts_PostgreSQL.txt file.

Procedure

- 1 Create a database and user for vCenter Server.
 - a In the master database, create a database for vCenter Server.
 - b Create a database user for vCenter Server and map it to the vCenter Server and msdb databases.

For example, to create the database VCDB and user vpxuser, you can run the following script:

```
use master
go
CREATE DATABASE VCDB ON PRIMARY
(NAME = N'vcdb', FILENAME = N'C:\database_path\VCDB.mdf', SIZE = 10MB, FILEGROWTH = 10% )
LOG ON
(NAME = N'vcdb_log', FILENAME = N'C:\database_path\VCDB.ldf', SIZE = 1000KB, FILEGROWTH = 10%)
COLLATE SQL_Latin1_General_CP1_CI_AS
go
use VCDB
go
CREATE LOGIN vpxuser WITH PASSWORD=N'vpxuser!0', DEFAULT_DATABASE=VCDB,
DEFAULT_LANGUAGE=us_english, CHECK_POLICY=OFF
go
CREATE USER vpxuser for LOGIN vpxuser
go
```

```
use MSDB
go
CREATE USER vpxuser for LOGIN vpxuser
go
```

You now have a Microsoft SQL Server database that you can use with vCenter Server.

- 2 Assign the db_owner role to the vCenter Server database user on both the vCenter Server and msdb databases.

For example, to assign the db_owner role to the vpxuser user, you can run the following script:

```
use VCDB
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
use MSDB
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
```

- 3 Enable database monitoring for the vCenter Server database user.

For example, to grant database disk size monitoring permissions to the vpxuser user, you can run the following script:

```
use master
go
grant VIEW SERVER STATE to vpxuser
go
GRANT VIEW ANY DEFINITION TO vpxuser
go
```

When you install vCenter Server, the installer uses the default dbo schema to assign permissions to the db_owner role.

Prepare the vCenter Server Database by Creating Custom Database Schema and Roles

As an alternative to using the db_owner database role, experienced database administrators can set permissions by creating database schema and roles manually, which ensures greater control over database permissions.

You must first create a database and user for vCenter Server. Then you can create a custom schema and new database roles for the database user. You must also enable database monitoring for the user before you install vCenter Server. See [Database Permission Requirements for vCenter Server](#).

To perform the following procedure, you can either use the graphical user interface or run scripts. The vCenter Server installer package contains example scripts in the vCenter-Server\dbschema\DB_and_schema_creation_scripts_PostgreSQL.txt file.

Procedure

- 1 Create a database and user for vCenter Server.
 - a In the master database, create a database for vCenter Server.
 - b Create a database user for vCenter Server and map it to the vCenter Server and msdb databases.

For example, to create the database VCDB and user vpxuser, you can run the following script:

```
use master
go
CREATE DATABASE VCDB ON PRIMARY
(NAME = N'vcdb', FILENAME = N'C:\database_path\VCDB.mdf', SIZE = 10MB, FILEGROWTH = 10% )
LOG ON
(NAME = N'vcdb_log', FILENAME = N'C:\database_path\VCDB.ldf', SIZE = 1000KB, FILEGROWTH = 10%)
COLLATE SQL_Latin1_General_CP1_CI_AS
go
use VCDB
go
CREATE LOGIN vpxuser WITH PASSWORD=N'vpxuser!0', DEFAULT_DATABASE=VCDB,
DEFAULT_LANGUAGE=us_english, CHECK_POLICY=OFF
go
CREATE USER vpxuser for LOGIN vpxuser
go
use MSDB
go
CREATE USER vpxuser for LOGIN vpxuser
go
```

You now have a Microsoft SQL Server database that you can use with vCenter Server.

- 2 In the vCenter Server database, create a database schema and assign it to the vCenter Server database user.

For example, to create the schema VMW in VCDB and assign it to the vpxuser user, you can run the following script:

```
use VCDB
CREATE SCHEMA VMW
go
ALTER USER vpxuser WITH DEFAULT_SCHEMA =VMW
```

- 3 In the vCenter Server database, create and grant privileges to the VC_ADMIN_ROLE and VC_USER_ROLE database roles and assign them to the vCenter Server database user.

For example, to create the roles in VCDB and assign them to the vpxuser user, you can run the following script:

```
use VCDB
go
if not exists (SELECT name FROM sysusers WHERE issqlrole=1 AND name = 'VC_ADMIN_ROLE')
CREATE ROLE VC_ADMIN_ROLE;
```

```

GRANT ALTER ON SCHEMA :: VMW to VC_ADMIN_ROLE;
GRANT REFERENCES ON SCHEMA :: VMW to VC_ADMIN_ROLE;
GRANT INSERT ON SCHEMA :: VMW to VC_ADMIN_ROLE;

GRANT CREATE TABLE to VC_ADMIN_ROLE;
GRANT CREATE VIEW to VC_ADMIN_ROLE;
GRANT CREATE Procedure to VC_ADMIN_ROLE;

if not exists (SELECT name FROM sysusers WHERE issqlrole=1 AND name = 'VC_USER_ROLE')
CREATE ROLE VC_USER_ROLE
go
GRANT SELECT ON SCHEMA :: VMW to VC_USER_ROLE
go
GRANT INSERT ON SCHEMA :: VMW to VC_USER_ROLE
go
GRANT DELETE ON SCHEMA :: VMW to VC_USER_ROLE
go
GRANT UPDATE ON SCHEMA :: VMW to VC_USER_ROLE
go
GRANT EXECUTE ON SCHEMA :: VMW to VC_USER_ROLE
go
sp_addrolemember VC_USER_ROLE , vpxuser
go
sp_addrolemember VC_ADMIN_ROLE , vpxuser
go

```

- 4 In the msdb database, create and grant privileges to the VC_ADMIN_ROLE database role and assign it to the vCenter Server database user.

For example, to create the roles and assign them to the vpxuser user, you can run the following script:

```

use MSDB
go
if not exists (SELECT name FROM sysusers WHERE issqlrole=1 AND name = 'VC_ADMIN_ROLE')
CREATE ROLE VC_ADMIN_ROLE;
go
GRANT SELECT on msdb.dbo.syscategories to VC_ADMIN_ROLE
go
GRANT SELECT on msdb.dbo.sysjobsteps to VC_ADMIN_ROLE
go
GRANT SELECT ON msdb.dbo.sysjobs to VC_ADMIN_ROLE
go
GRANT SELECT ON msdb.dbo.sysjobs_view to VC_ADMIN_ROLE
go
GRANT EXECUTE ON msdb.dbo.sp_add_job TO VC_ADMIN_ROLE
go
GRANT EXECUTE ON msdb.dbo.sp_delete_job TO VC_ADMIN_ROLE
go
GRANT EXECUTE ON msdb.dbo.sp_add_jobstep TO VC_ADMIN_ROLE
go
GRANT EXECUTE ON msdb.dbo.sp_update_job TO VC_ADMIN_ROLE
go
GRANT EXECUTE ON msdb.dbo.sp_add_jobserver TO VC_ADMIN_ROLE

```

```

go
GRANT EXECUTE ON msdb.dbo.sp_add_jobschedule TO VC_ADMIN_ROLE
go
GRANT EXECUTE ON msdb.dbo.sp_add_category TO VC_ADMIN_ROLE
go
sp_addrolemember VC_ADMIN_ROLE , vpxuser
go

```

Note The VC_ADMIN_ROLE role in the msdb database is required only during installation and upgrade of vCenter Server. After the installation or upgrade, you can revoke the role and leave it as inactive for future upgrades, or you can remove it for increased security.

5 Enable database monitoring for the vCenter Server database user.

For example, to grant database disk size monitoring permissions to the vpxuser user, you can run the following script:

```

use master
go
grant VIEW SERVER STATE to vpxuser
go
GRANT VIEW ANY DEFINITION TO vpxuser
go

```

(Optional) Use a Script to Create Microsoft SQL Server Database Objects Manually

This topic describes how to create database objects manually instead of letting the vCenter Server installer create the data objects automatically.

Procedure

- 1 Log in to a Microsoft SQL Server Management Studio session with the vCenter Server database user account that you created on the vCenter Server and msdb databases.
- 2 In the vCenter Server installation package, locate the dbschema scripts in the vCenter-Server/dbschema directory.
- 3 Open the VCDB_mssql.SQL and the TopN_DB_mssql.sql files by using Microsoft SQL Server Management Studio and replace all occurrences of \$schema with your schema name.
- 4 Open the VCDB_views_mssql.sql file by using Microsoft SQL Server Management Studio and after each occurrence of ;, insert a new line and write go.
- 5 Run the scripts in a sequence on the database.

The DBO user must own the objects created by these scripts. Open the scripts one at a time in Microsoft SQL Server Management Studio and press F5 to execute each script in the following order:

- a VCDB_mssql.SQL
- b insert_stats_proc_mssql.sql
- c load_stats_proc_mssql.sql

- d purge_stat2_proc_mssql.sql
- e purge_stat3_proc_mssql.sql
- f purge_usage_stats_proc_mssql.sql
- g stats_rollup1_proc_mssql.sql
- h stats_rollup2_proc_mssql.sql
- i stats_rollup3_proc_mssql.sql
- j cleanup_events_mssql.sql
- k delete_stats_proc_mssql.sql
- l upsert_last_event_proc_mssql.sql
- m load_usage_stats_proc_mssql.sql
- n TopN_DB_mssql.sql
- o calc_topn1_proc_mssql.sql
- p calc_topn2_proc_mssql.sql
- q calc_topn3_proc_mssql.sql
- r calc_topn4_proc_mssql.sql
- s clear_topn1_proc_mssql.sql
- t clear_topn2_proc_mssql.sql
- u clear_topn3_proc_mssql.sql
- v clear_topn4_proc_mssql.sql
- w rule_topn1_proc_mssql.sql
- x rule_topn2_proc_mssql.sql
- y rule_topn3_proc_mssql.sql
- z rule_topn4_proc_mssql.sql
- aa process_license_snapshot_mssql.sql
- ab l_stats_rollup3_proc_mssql.sql
- ac l_purge_stat2_proc_mssql.sql
- ad l_purge_stat3_proc_mssql.sql
- ae l_stats_rollup1_proc_mssql.sql
- af l_stats_rollup2_proc_mssql.sql
- ag VCDB_views_mssql.sql

6 (Optional) Run the scripts to enable database health monitoring.

- a job_dbm_performance_data_mssql.sql

b process_performance_data_mssql.sql

- 7 For all supported editions of Microsoft SQL Server except Microsoft SQL Server Express, run the scripts to set up scheduled jobs on the database.

These scripts ensure that the SQL Server Agent service is running.

a job_schedule1_mssql.sql

b job_schedule2_mssql.sql

c job_schedule3_mssql.sql

d job_cleanup_events_mssql.sql

e job_topn_past_day_mssql.sql

f job_topn_past_week_mssql.sql

g job_topn_past_month_mssql.sql

h job_topn_past_year_mssql.sql

- 8 For all the procedures you created in [Step 5](#), grant the execute privilege to the vCenter Server database user in the vCenter Server database.

For example, to grant execute privilege for the procedures to the vpxuser user, you can run the following script.

```
grant execute on insert_stats_proc to vpxuser
grant execute on purge_stat2_proc to vpxuser
grant execute on purge_stat3_proc to vpxuser
grant execute on purge_usage_stat_proc to vpxuser
grant execute on stats_rollup1_proc to vpxuser
grant execute on stats_rollup2_proc to vpxuser
grant execute on stats_rollup3_proc to vpxuser
grant execute on cleanup_events_tasks_proc to vpxuser
grant execute on delete_stats_proc to vpxuser
grant execute on upsert_last_event_proc to vpxuser
grant execute on load_usage_stats_proc to vpxuser
grant execute on load_stats_proc to vpxuser
grant execute on calc_topn1_proc to v
grant execute on calc_topn2_proc to vpxuser
grant execute on calc_topn3_proc to vpxuser
grant execute on calc_topn4_proc to vpxuser
grant execute on clear_topn1_proc to vpxuser
grant execute on clear_topn2_proc to vpxuser
grant execute on clear_topn3_proc to vpxuser
grant execute on clear_topn4_proc to vpxuser
grant execute on rule_topn1_proc to vpxuser
grant execute on rule_topn2_proc to vpxuser
grant execute on rule_topn3_proc to vpxuser
grant execute on rule_topn4_proc to vpxuser
grant execute on process_license_snapshot_proc to vpxuser
grant execute on l_stats_rollup3_proc to vpxuser
```

```
grant execute on l_purge_stat2_proc to vpxuser
grant execute on l_purge_stat3_proc to vpxuser
grant execute on l_stats_rollup1_proc to vpxuser
grant execute on l_stats_rollup2_proc to vpxuser
```

If you ran the script `process_performance_data_mssql.sql` in [Step 5](#), grant the following execute privilege to the vCenter Server database.

```
grant execute on process_performance_data_proc to vpxuser
```

You created the vCenter Server tables manually.

Note During the vCenter Server installation, when a database reinitialization warning message appears, select **Do not overwrite, leave my existing database in place** and continue the installation.

Configure a SQL Server ODBC Connection

After you create and configure a SQL Server database and user for vCenter Server, you must create a 64-bit DSN on the machine on which you plan to install vCenter Server. During the vCenter Server installation, you use the DSN to establish a connection between vCenter Server and the database.

If you use SQL Server for vCenter Server, do not use the master or any other system database.

See your Microsoft SQL ODBC documentation for specific instructions for configuring the SQL Server ODBC connection.

Caution If you are using a named instance of Microsoft SQL Server 2008 Standard Edition with vCenter Server, do not name the instance MSSQLSERVER. If you do, the JDBC connection does not work, and certain features, such as Performance Charts, are not available.

Prerequisites

Deploy SQL Native Client version 10 or 11.

Procedure

- 1 On the machine on which you plan to install vCenter Server, select **Start > Administrative Tools > Data Sources (ODBC)**.
- 2 On the **System DSN** tab, modify an existing or create a new SQL Server ODBC connection.
 - To modify an existing SQL Server ODBC connection, select the connection from the System Data Source list and click **Configure**.

Important The existing DSN must use SQL Native Client version 10 or 11.

- To create a new SQL Server ODBC connection, click **Add**, select **SQL Native Client**, and click **Finish**.
- 3 In the **Name** text box, enter an ODBC data source name (DSN).

For example, **VMware vCenter Server**.

- 4 (Optional) In the **Description** text box, enter an ODBC DSN description.
- 5 In the **Server** text box, enter the IP address or FQDN of the SQL Server and, if you want to use a non-default port to access the SQL Server, enter a custom port separated by a comma.

For example, if the IP address of your SQL Server is 10.160.10.160 and you want to access the server by using custom port 8347, enter **10.160.10.160,8347**.

Note You cannot use a database server alias to create a DSN.

- 6 Select an authentication method.

- **Integrate Windows authentication.**

Additionally, you can also enter the Service Principal Name (SPN).

Important You cannot use this option if the vCenter Server service is running under the Microsoft Windows built-in system account.

- **SQL Server authentication.**

Enter your SQL Server login name and password.

- 7 Select the database created for the vCenter Server system from the **Change the default database to** menu.
- 8 Click **Finish**.
- 9 Test the data source by selecting **Test Data Source** and clicking **OK** from the **ODBC Microsoft SQL Server Setup** menu.
- 10 Verify that the SQL Agent is running on your database server.

Configure Microsoft SQL Server TCP/IP for JDBC

If the Microsoft SQL Server database has TCP/IP disabled and the dynamic ports are not set, the JDBC connection remains closed. The closed connection causes the vCenter Server statistics to malfunction. You can configure the server TCP/IP for JDBC.

This task applies to remote Microsoft SQL Server database servers. You can skip this task if your database is located on the same machine as vCenter Server.

Procedure

- 1 Select **Start > All Programs > Microsoft SQL Server > Configuration Tool > SQL Server Configuration Manager**.
- 2 Select **SQL Server Network Configuration > Protocols for *Instance name***.
- 3 Enable TCP/IP.
- 4 Open TCP/IP Properties.

- 5 On the **Protocol** tab, make the following entries.

Enabled	Yes
Listen All	Yes
Keep Alive	30000

- 6 On the **IP Addresses** tab, make the following selections.

Active	Yes
TCP Dynamic Ports	0

- 7 Restart the SQL Server service from **SQL Server Configuration Manager > SQL Server Services**.
- 8 Start the SQL Server Browser service from **SQL Server Configuration Manager > SQL Server Services**.

Configure Oracle Databases

To use an Oracle database for your vCenter Server repository, configure your database to work with vCenter Server.

You can install and configure the Oracle database on the same machine on which you plan to install vCenter Server. You can install and configure the Oracle database on a separate machine.

Procedure

1 [Prepare the vCenter Server Oracle Database](#)

To use an Oracle database with vCenter Server, you must create the database with certain tablespaces and privileges, and the database user with certain permissions.

2 [\(Optional\) Use a Script to Create the Oracle Database Schema](#)

The vCenter Server installer creates the schema during installation. For experienced database administrators who need more control over schema creation because of environmental constraints, you can optionally use a script to create your database schema.

3 [Create a Net Service Name](#)

To configure an Oracle ODBC DSN, you must have a net service name for your database. On the machine on which your Oracle database is running, you must create a net service name for the vCenter Server tablespace.

4 [Configure an Oracle ODBC Connection](#)

After you create and configure an Oracle database and user for vCenter Server, you must create a 64-bit DSN on the machine on which you plan to install vCenter Server. During the vCenter Server installation, you use the DSN to establish a connection between vCenter Server and the database.

Prepare the vCenter Server Oracle Database

To use an Oracle database with vCenter Server, you must create the database with certain tablespaces and privileges, and the database user with certain permissions.

You must first create a tablespace and user for vCenter Server. Then you grant permissions to the database user. You must also enable database monitoring for the user before you install vCenter Server. See [Database Permission Requirements for vCenter Server](#).

To perform the following procedure, you can either use the graphical user interface or run scripts. The vCenter Server installer package contains example scripts in the vCenter-Server\dbschema\DB_and_schema_creation_scripts_PostgreSQL.txt file.

Prerequisites

Log in to a SQL*Plus session with the system account.

Procedure

- 1 Create a tablespace for vCenter Server.

For example, to create the tablespace VPX, you can run the following script:

```
CREATE SMALLFILE TABLESPACE "VPX" DATAFILE 'C:\database_path\vpx01.dbf'
SIZE 1G AUTOEXTEND ON NEXT 10M MAXSIZE UNLIMITED LOGGING EXTENT MANAGEMENT LOCAL SEGMENT
SPACE MANAGEMENT AUTO;
```

- 2 Create a database user with the correct permissions for vCenter Server.

For example, to create the VPXADMIN user, you can run the following script:

```
CREATE USER "VPXADMIN" PROFILE "DEFAULT" IDENTIFIED BY "oracle" DEFAULT TABLESPACE "VPX" ACCOUNT
UNLOCK;
grant connect to VPXADMIN;
grant resource to VPXADMIN;
grant create view to VPXADMIN;
grant create sequence to VPXADMIN;
grant create table to VPXADMIN;
grant create materialized view to VPXADMIN;
grant execute on dbms_lock to VPXADMIN;
grant execute on dbms_job to VPXADMIN;
grant select on dba_lock to VPXADMIN;
grant select on dba_tablespaces to VPXADMIN;
```

```
grant select on dba_temp_files to VPXADMIN;
grant select on dba_data_files to VPXADMIN;
grant select on v_$session to VPXADMIN;
grant unlimited tablespace to VPXADMIN;
```

By default, the RESOURCE role has the **CREATE PROCEDURE**, **CREATE TABLE**, and **CREATE SEQUENCE** privileges assigned. If the RESOURCE role lacks these privileges, grant them to the vCenter Server database user.

Note Instead of granting unlimited tablespace, you can set a specific tablespace quota. The recommended quota is unlimited with a minimum of at least 500MB. To set an unlimited quota, use the following command.

```
alter user "VPXADMIN" quota unlimited on "VPX";
```

If you set a limited quota, monitor the remaining available tablespace to avoid the following error.

```
ORA-01536: space quota exceeded for tablespace 'tablespace'
```

You now have an Oracle database user for vCenter Server.

3 Enable database monitoring for the vCenter Server database user.

For example, to grant database disk size monitoring permissions to the VPXADMIN user, you can run the following script:

```
grant select on v_$system_event to VPXADMIN;
grant select on v_$sysmetric_history to VPXADMIN;
grant select on v_$sysstat to VPXADMIN;
grant select on dba_data_files to VPXADMIN;
grant select on v_$loghist to VPXADMIN;
```

(Optional) Use a Script to Create the Oracle Database Schema

The vCenter Server installer creates the schema during installation. For experienced database administrators who need more control over schema creation because of environmental constraints, you can optionally use a script to create your database schema.

Procedure

- 1 Open a SQL*Plus window with a user that has schema owner rights on the vCenter Server database.
- 2 Locate the dbschema scripts in the vCenter Server installation package */installation directory/vCenter-Server/dbschema* directory.
- 3 In SQL*Plus, run the scripts in a sequence on the database.
 - a VCDB_oracle.SQL
 - b VCDB_views_oracle.SQL
 - c insert_stats_proc_oracle.sql

- d load_stats_proc_oracle.sql
- e purge_stat2_proc_oracle.sql
- f purge_stat3_proc_oracle.sql
- g purge_usage_stats_proc_oracle.sql
- h stats_rollup1_proc_oracle.sql
- i stats_rollup2_proc_oracle.sql
- j stats_rollup3_proc_oracle.sql
- k cleanup_events_oracle.sql
- l delete_stats_proc_oracle.sql
- m load_usage_stats_proc_oracle.sql
- n TopN_DB_oracle.sql
- o calc_topn1_proc_oracle.sql
- p calc_topn2_proc_oracle.sql
- q calc_topn3_proc_oracle.sql
- r calc_topn4_proc_oracle.sql
- s clear_topn1_proc_oracle.sql
- t clear_topn2_proc_oracle.sql
- u clear_topn3_proc_oracle.sql
- v clear_topn4_proc_oracle.sql
- w rule_topn1_proc_oracle.sql
- x rule_topn2_proc_oracle.sql
- y rule_topn3_proc_oracle.sql
- z rule_topn4_proc_oracle.sql
- aa process_license_snapshot_oracle.sql
- ab l_purge_stat2_proc_oracle.sql
- ac l_purge_stat3_proc_oracle.sql
- ad l_stats_rollup1_proc_oracle.sql
- ae l_stats_rollup2_proc_oracle.sql
- af l_stats_rollup3_proc_oracle.sql

4 (Optional) You can also run the following scripts to enable database health monitoring.

- a job_dbm_performance_data_oracle.sql
- b process_performance_data_oracle.sql

- 5 For all supported editions of Oracle Server, run the scripts to set up scheduled jobs on the database.
 - a `job_schedule1_oracle.sql`
 - b `job_schedule2_oracle.sql`
 - c `job_schedule3_oracle.sql`
 - d `job_cleanup_events_oracle.sql`
 - e `job_topn_past_day_oracle.sql`
 - f `job_topn_past_week_oracle.sql`
 - g `job_topn_past_month_oracle.sql`
 - h `job_topn_past_year_oracle.sql`

You created the vCenter Server tables manually.

Note During the vCenter Server installation, when a database reinitialization warning message appears, select **Do not overwrite, leave my existing database in place** and continue the installation.

Create a Net Service Name

To configure an Oracle ODBC DSN, you must have a net service name for your database. On the machine on which your Oracle database is running, you must create a net service name for the vCenter Server tablespace.

Procedure

- 1 Use a text editor or the Net8 Configuration Assistant to open the `tnsnames.ora` file located in the directory `C:\Oracle\Oraxx\NETWORK\ADMIN`, where `xx` is either **10g** or **11g**.
- 2 Add the following entry, where `HOST` is the managed host to which the client must connect.

```
VPX_TNS =
(DESCRIPTION =
(AADDRESS_LIST =
(AADDRESS=(PROTOCOL=TCP)(HOST=vpxd-oracle)(PORT=1521))
)
(CONNECT_DATA =
(SERVICE_NAME = ORCL)
)
)
```

Configure an Oracle ODBC Connection

After you create and configure an Oracle database and user for vCenter Server, you must create a 64-bit DSN on the machine on which you plan to install vCenter Server. During the vCenter Server installation, you use the DSN to establish a connection between vCenter Server and the database.

Prerequisites

Install the Oracle Client 11.2.0.3 p16656151 (Patch 19) or later, 11.2.0.4, 12.1.0.1.12 or later, or 12.1.0.2.

Procedure

- 1 On the machine on which you plan to install vCenter Server, select **Start > Administrative Tools > Data Sources (ODBC)**.
- 2 On the **System DSN** tab, modify an existing or create a new Oracle ODBC connection.
 - To modify an existing Oracle ODBC connection, select the connection from the System Data Source list and click **Configure**.
 - To create an Oracle ODBC connection, click **Add**, select the Oracle client, and click **Finish**.
- 3 In the **Data Source Name** text box, enter an ODBC data source name (DSN).
For example, **VMware vCenter Server**.
- 4 (Optional) In the **Description** text box, enter an ODBC DSN description.
- 5 In the **TNS Service Name** text box, enter the net service name for the database to which you want to connect.
For example, **VPX_TNS**.

This is the net service name that you previously configured in the `tnsnames.ora` file that is located in the `NETWORK\ADMIN` folder in the Oracle database installation location.
- 6 In the **User ID** text box, enter the database user name for vCenter Server.
For example, **VPXADMIN**.
- 7 Click **Test Connection**.
- 8 In the **Password** text box, enter the password of the database user and click **OK**.
If you configured the DNS correctly, the `Connection successful` message appears.
- 9 Click **OK**.

Database Permission Requirements for vCenter Server

vCenter Server requires a database. If you decide to use an external Oracle or Microsoft SQL Server database, when you create the database, you must grant certain permissions to the database user.

Table 4-6. Microsoft SQL Database Permissions for vCenter Server

Permission	Description
<code>GRANT ALTER ON SCHEMA :: [VMW] TO VC_ADMIN_ROLE</code>	Mandatory when you work with SQL Server custom schema.
<code>GRANT REFERENCES ON SCHEMA :: [VMW] TO VC_ADMIN_ROLE</code>	Mandatory when you work with SQL Server custom schema.
<code>GRANT INSERT ON SCHEMA :: [VMW] TO VC_ADMIN_ROLE</code>	Mandatory when you work with SQL Server custom schema.
<code>GRANT CREATE TABLE TO VC_ADMIN_ROLE</code>	Necessary for creating a table.
<code>GRANT CREATE VIEW TO VC_ADMIN_ROLE</code>	Necessary for creating a view.
<code>GRANT CREATE PROCEDURE TO VC_ADMIN_ROLE</code>	Necessary for creating a stored procedure.

Table 4-6. Microsoft SQL Database Permissions for vCenter Server (Continued)

Permission	Description
GRANT SELECT ON SCHEMA :: [VMW] TO VC_USER_ROLE	Permissions that let you run SELECT, INSERT, DELETE, UPDATE operations on tables which are part of the VMW schema.
GRANT INSERT ON SCHEMA :: [VMW] TO VC_USER_ROLE	
GRANT DELETE ON SCHEMA :: [VMW] TO VC_USER_ROLE	
GRANT UPDATE ON SCHEMA :: [VMW] TO VC_USER_ROLE	
GRANT EXECUTE ON SCHEMA :: [VMW] TO VC_USER_ROLE	Necessary for running a stored procedure in the db schema.
GRANT SELECT ON msdb.dbo.syscategories TO VC_ADMIN_ROLE	Necessary for deploying SQL Server jobs. These permissions are mandatory only during installation and upgrade and not required after deployment.
GRANT SELECT ON msdb.dbo.sysjobsteps TO VC_ADMIN_ROLE	
GRANT SELECT ON msdb.dbo.sysjobs TO VC_ADMIN_ROLE	
GRANT EXECUTE ON msdb.dbo.sp_add_job TO VC_ADMIN_ROLE	
GRANT EXECUTE ON msdb.dbo.sp_delete_job TO VC_ADMIN_ROLE	
GRANT EXECUTE ON msdb.dbo.sp_add_jobstep TO VC_ADMIN_ROLE	
GRANT EXECUTE ON msdb.dbo.sp_update_job TO VC_ADMIN_ROLE	
GRANT EXECUTE ON msdb.dbo.sp_add_jobserver TO VC_ADMIN_ROLE	
GRANT EXECUTE ON msdb.dbo.sp_add_jobschedule TO VC_ADMIN_ROLE	
GRANT EXECUTE ON msdb.dbo.sp_add_category TO VC_ADMIN_ROLE	
GRANT VIEW SERVER STATE TO [vpxuser]	Provides access to SQL Server DMV views and sp_lock execution.
GRANT VIEW ANY DEFINITION TO [vpxuser]	Necessary for providing the user with the privileges to see metadata for SQL Server objects.

Table 4-7. Oracle Database Permissions for vCenter Server

Permission	Description
GRANT CONNECT TO VPXADMIN	Necessary for connecting to the Oracle database.
GRANT RESOURCE TO VPXADMIN	Necessary for creating a trigger, sequence, type, procedure, and so on. By default, the RESOURCE role has the CREATE PROCEDURE, CREATE TABLE, and CREATE SEQUENCE privileges assigned. If the RESOURCE role lacks these privileges, grant them to the vCenter Server database user.
GRANT CREATE VIEW TO VPXADMIN	Necessary for creating a view.

Table 4-7. Oracle Database Permissions for vCenter Server (Continued)

Permission	Description
GRANT CREATE SEQUENCE TO VPXADMIN	Necessary for creating a sequence.
GRANT CREATE TABLE TO VPXADMIN	Necessary for creating a table.
GRANT CREATE MATERIALIZED VIEW TO VPXADMIN	Necessary for creating a materialized view.
GRANT EXECUTE ON dbms_lock TO VPXADMIN	Necessary for guaranteeing that the vCenter Server database is used by a single vCenter Server instance.
GRANT EXECUTE ON dbms_job TO VPXADMIN	Necessary during installation or upgrade for scheduling and managing the SQL jobs. This permission is not required after deployment.
GRANT SELECT ON dba_lock TO VPXADMIN	Necessary for determining existing locks on the vCenter Server database.
GRANT SELECT ON dba_tablespaces TO VPXADMIN	Necessary during upgrade for determining the required disk space. This permission is not required after deployment.
GRANT SELECT ON dba_temp_files TO VPXADMIN	Necessary during upgrade for determining the required disk space. This permission is not required after deployment.
GRANT SELECT ON dba_data_files TO VPXADMIN	Necessary for monitoring the free space while vCenter Server is working.
GRANT SELECT ON v_\$session TO VPXADMIN	View used to determine existing locks on the vCenter Server database.
GRANT UNLIMITED TABLESPACE TO VPXADMIN	Necessary for granting unlimited tablespace permissions to the vCenter Server database user.
GRANT SELECT ON v_\$system_event TO VPXADMIN	Necessary for checking log file switches.
GRANT SELECT ON v_\$sysmetric_history TO VPXADMIN	Necessary for checking the CPU utilization.
GRANT SELECT ON v_\$sysstat TO VPXADMIN	Necessary for determining the Buffer Cache Hit Ratio.
GRANT SELECT ON dba_data_files TO VPXADMIN	Necessary for determining the tablespace utilization.
GRANT SELECT ON v_\$loghist TO VPXADMIN	Necessary for checking the checkpoint frequency.

The privileges on the master database are used to monitor the vCenter Server database. so that, for example, if a certain threshold is reached, you can see an alert.

Verify That vCenter Server Can Communicate with the Local Database

If your database is on the same machine on which vCenter Server is to be installed, and you changed the machine name, verify the configuration. Make sure that the vCenter Server DSN is configured to communicate with the new name of the machine.

Changing the vCenter Server computer name impacts database communication if the database server is on the same computer with vCenter Server. If you changed the machine name, you can verify that communication remains intact.

If your database is remote, you can skip this procedure. The name change has no effect on communication with remote databases.

After you rename the server, verify with your database administrator or the database vendor that all components of the database are working.

Prerequisites

- Make sure that the database server is running.
- Make sure that the vCenter Server computer name is updated in the domain name service (DNS).

Procedure

- 1 Update the data source information, as needed.
- 2 To test this condition, ping the computer name.

For example, if the computer name is `host-1.company.com`, run the following command at the Windows command prompt:

```
ping host-1.company.com
```

If you can ping the computer name, the name is updated in DNS.

vCenter Server communication is confirmed. You can continue to prepare other components of your environment.

Maintaining a vCenter Server Database

After your vCenter Server database instance and vCenter Server are installed and operational, perform standard database maintenance processes.

The standard database maintenance processes include the following:

- Monitoring the growth of the log file and compacting the database log file, as needed.
- Scheduling regular backups of the database.
- Backing up the database before any vCenter Server upgrade.

See your database vendor's documentation for specific maintenance procedures and support.

Synchronizing Clocks on the vSphere Network

Verify that all components on the vSphere network have their clocks synchronized. If the clocks on the machines in your vSphere network are not synchronized, SSL certificates, which are time-sensitive, might not be recognized as valid in communications between network machines.

Unsynchronized clocks can result in authentication problems, which can cause the installation to fail or prevent the vCenter Server Appliance vpxd service from starting.

Verify that any Windows host machine on which vCenter Server runs is synchronized with the Network Time Server (NTP) server. See the Knowledge Base article <http://kb.vmware.com/kb/1318>.

To synchronize ESXi clocks with an NTP server, you can use the VMware Host Client. For information about editing the time configuration of an ESXi host, see *vSphere Single Host Management*.

Using a User Account for Running vCenter Server

You can use the Microsoft Windows built-in system account or a user account to run vCenter Server. With a user account, you can enable Windows authentication for SQL Server, and it provides more security.

The user account must be an administrator on the local machine. In the installation wizard, you specify the account name as *DomainName\Username*. You must configure the SQL Server database to allow the domain account access to SQL Server.

The Microsoft Windows built-in system account has more permissions and rights on the server than the vCenter Server system needs, which can contribute to security problems.

Important If the vCenter Server service is running under the Microsoft Windows built-in system account, when using Microsoft SQL Server, vCenter Server supports only DSNs with SQL Server authentication.

For SQL Server DSNs configured with Windows authentication, use the same user account for the VMware VirtualCenter Management Webservices service and the DSN user.

If you do not plan to use Microsoft Windows authentication for SQL Server or you are using an Oracle database, you might still want to set up a local user account for the vCenter Server system. The only requirement is that the user account is an administrator on the local machine and the account must be granted the **Log on as a service** privilege.

Note Starting with vSphere 6.5, the vCenter Server services are not standalone services under Windows SCM, instead they run as child processes of the VMware Service Lifecycle Manager service.

Installing vCenter Server on IPv6 Machines

Starting with vSphere 6.5, vCenter Server supports mixed IPv4 and IPv6 environments.

You can connect vCenter Server with an IPv4 address to vCenter Server with an IPv6 address. When you install vCenter Server with an IPv6 address, use the fully qualified domain name (FQDN) or host name of the machine on which you install vCenter Server. When you install vCenter Server with an IPv4 address, the best practice is to use the fully qualified domain name (FQDN) or host name of the machine on which you install vCenter Server, because the IP address can change if assigned by DHCP.

Running the vCenter Server Installer from a Network Drive

You can run the vCenter Server installer from a network drive, but you cannot install the software on a network drive.

In Windows, you can run the installers from the network drive and install the software on the local machine.

Required Information for Installing vCenter Server or Platform Services Controller on Windows

When you install vCenter Server with an embedded Platform Services Controller, Platform Services Controller, or vCenter Server with an external Platform Services Controller, the wizard prompts you for the installation information. It is a best practice to keep a record of the values that you entered in case you must reinstall the product.

You can use this worksheet to record the information that you need for the installation of vCenter Server with an embedded Platform Services Controller, Platform Services Controller, or vCenter Server with an external Platform Services Controller.

Table 4-8. Required Information for Installing vCenter Server or Platform Services Controller on Windows

Required for	Required Information	Default	Your Entry
All deployment types	System name of the local system A system name to use for managing the local system. The system name must be an FQDN. If a DNS is not available, provide a static IP address.	-	
<ul style="list-style-type: none"> ■ vCenter Server with an embedded Platform Services Controller ■ Platform Services Controller as the first instance in a new domain 	Name for the new vCenter Single Sign-On domain User name Password for the vCenter Single Sign-On administrator account The password must be at least 8 characters, but no more than 20 characters in length. The password must conform to the following requirements: <ul style="list-style-type: none"> ■ Must contain at least one uppercase letter. ■ Must contain at least one lowercase letter. ■ Must contain at least one number. ■ Must contain at least one special character, such as ampersand (&), hash key (#), and percent sign (%). Site name A name for the vCenter Single Sign-On site.	vsphere.local administrator - - Default-First-Site	 You cannot change the default user name during installation.
<ul style="list-style-type: none"> ■ vCenter Server with an external Platform Services Controller ■ Platform Services Controller as a subsequent instance in an existing domain 	FQDN or IP address of the Platform Services Controller instance that you want to join You must join a Platform Services Controller instance of the same version. HTTPS port of the Platform Services Controller instance Password of the vCenter Single Sign On administrator user for the domain vCenter Single Sign-On site name You can join an existing site or create a new site.	- 443 - -	

Table 4-8. Required Information for Installing vCenter Server or Platform Services Controller on Windows (Continued)

Required for	Required Information	Default	Your Entry
■ vCenter Server with an embedded Platform Services Controller	vCenter Server service account information Can be the Windows local system account or a user service account.	Windows local system account	
	Note Starting with vSphere 6.5, the vCenter Server services run as child processes of the VMware Service Lifecycle Manager service.		
■ vCenter Server with an external Platform Services Controller	Account user name Only if you use a user service account	-	
	Account password Only if you use a user service account	-	
■ vCenter Server with an embedded Platform Services Controller	vCenter Server database Can be the embedded VMware Postgres database or an existing external database	embedded Postgres database	
	Data source name (DSN) Only if you use an existing external database. Leading and trailing spaces are not supported. Remove spaces from the beginning or end of the DSN.	-	
■ vCenter Server with an external Platform Services Controller	Database user name Only if you use an existing external database. Non-ASCII characters are not supported.	-	
	Database password Only if you use an existing external database.	-	
All deployment types	HTTP port	80	
	HTTPS port	443	
	Syslog Service port	514	
	Syslog Service TLS port	1514	
■ vCenter Server with an embedded Platform Services Controller	Secure Token Service port	7444	
■ Platform Services Controller			
■ vCenter Server with an embedded Platform Services Controller	Auto Deploy Management port	6502	
	Auto Deploy Service port	6501	
	ESXi Dump Collector port	6500	
■ vCenter Server with an external Platform Services Controller	ESXi Heartbeat port	902	
	vSphere Web Client port	9443	

Table 4-8. Required Information for Installing vCenter Server or Platform Services Controller on Windows (Continued)

Required for	Required Information	Default	Your Entry
All deployment types	Destination folder <ul style="list-style-type: none"> ■ The folder in which to install vCenter Server or Platform Services Controller ■ The folder in which to store data for vCenter Server or Platform Services Controller The installation paths cannot contain non-ASCII characters, commas (,), periods (.), exclamation points (!), pound signs (#), at signs (@), or percentage signs (%).	<ul style="list-style-type: none"> ■ The default installation folder is C:\Program Files\VMware. ■ The default folder for data storage is C:\ProgramData\VMware. 	
<ul style="list-style-type: none"> ■ vCenter Server with an embedded Platform Services Controller ■ Platform Services Controller 	Join or do not participate in the VMware Customer Experience Improvement Program (CEIP) For information about the CEIP, see the Configuring Customer Experience Improvement Program section in <i>vCenter Server and Host Management</i> .	Join the CEIP	

Installing vCenter Server and Platform Services Controller on Windows

You can install vCenter Server with an embedded Platform Services Controller, Platform Services Controller, or vCenter Server with an external Platform Services Controller on a Windows virtual or physical machine.

You download the vCenter Server installer ISO file, mount it to the Windows host machine on which you want to perform the installation, start the installation wizard, and provide the inputs that required for the installation and setup.

Before installing vCenter Server that uses an external database, you must prepare your database. See [Preparing vCenter Server Databases for Install](#).

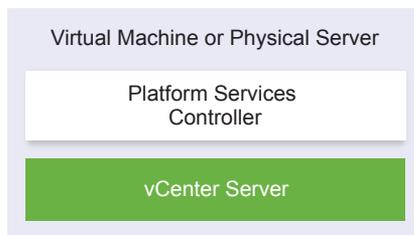
Important For topologies with external Platform Services Controller instances, you must install the replicating Platform Services Controller instances in a sequence. After the successful deployment of all Platform Services Controller instances in the domain, you can perform concurrent installations of multiple vCenter Server instances that point to a common external Platform Services Controller instance.

Install vCenter Server with an Embedded Platform Services Controller

You can deploy vCenter Server, the vCenter Server components, and the Platform Services Controller on one virtual machine or physical server.

After you deploy vCenter Server with an embedded Platform Services Controller, you can reconfigure your topology and switch to vCenter Server with an external Platform Services Controller. This is a one-way process after which you cannot switch back to vCenter Server with an embedded Platform Services Controller. You can repoint the vCenter Server instance only to an external Platform Services Controller that is configured to replicate the infrastructure data within the same domain.

Figure 4-1. vCenter Server with an Embedded Platform Services Controller



Prerequisites

- Verify that your system meets the minimum software and hardware requirements. See [vCenter Server for Windows Requirements](#).
- [Download the vCenter Server Installer for Windows](#).
- If you want to use the vSphere Web Client on the host machine on which you install vCenter Server, verify that Adobe Flash Player version 11.9 or later is installed on the system.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vCenter Server for Windows** and click **Install**.
- 3 Follow the prompts of the installation wizard to review the welcome page and accept the license agreement.
- 4 Select **vCenter Server and Embedded Platform Services Controller**, and click **Next**.
- 5 Enter the system network name, preferably an FQDN, and click **Next**.

You can also enter an IP address. If you enter an IP address, provide a static IP address.

Important Make sure the FQDN or IP address that you provide does not change. The system name cannot be changed after deployment. If the system name changes, you must uninstall vCenter Server and install it again.

6 Set up the new vCenter Single Sign-On domain and click **Next**.

a Enter the domain name, for example **vsphere.local**.

b Set the password for the vCenter Single Sign-On administrator account.

This is the password for the user `administrator@your_domain_name`. After installation, you can log in to vCenter Single Sign-On and to vCenter Server as `administrator@your_domain_name`.

c Enter the site name for vCenter Single Sign-On.

The site name is important if you are using vCenter Single Sign-On in multiple locations. The site name must contain alphanumeric characters. Choose your own name for the vCenter Single Sign-On site. You cannot change the name after installation.

Extended ASCII and non-ASCII characters are unsupported in site names. Your site name must include alphanumeric characters and a comma (,), period (.), question mark (?), dash (-), underscore (_), plus sign (+) or equals sign (=).

7 Select the vCenter Server service account and click **Next**.

Note Starting with vSphere 6.5, the vCenter Server services are not standalone services under Windows SCM, instead they run as child processes of the VMware Service Lifecycle Manager service.

Option	Description
Use Windows Local System Account	The vCenter Server service runs in the Windows Local System account. This option prevents you from connecting to an external database by using Windows integrated authentication.
Specify a user service account	The vCenter Server service runs in an administrative user account with a user name and password that you provide. Important The user credentials that you provide must be of a user who is in the local administrator group and who has the Log on as a service privilege.

8 Select the type of database that you want to use and click **Next**.

Option	Description
Use an embedded database (PostgreSQL)	vCenter Server uses the embedded PostgreSQL database. This database is suitable for small scale deployments.
Use an external database	vCenter Server uses an existing external database. <ul style="list-style-type: none"> a Select your database from the list of available DSNs. b Type the user name and the password for the DSN. If your database uses Windows NT authentication, the user name and password text boxes are disabled.

9 For each component, accept the default port numbers, or if another service is using the defaults, enter alternative ports, and click **Next**.

Make sure that ports 80 and 443 are free and dedicated, so that vCenter Single Sign-On can use these ports. Otherwise, use custom ports during installation.

10 (Optional) Change the default destination folders and click **Next**.

Important Do not use folders that end with an exclamation mark (!).

11 Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.

For information about the CEIP, see the Configuring Customer Experience Improvement Program section in *vCenter Server and Host Management*.

12 Click **Next**.

13 Review the summary of the installation settings and click **Install** to start the installation.

14 (Optional) After the installation finishes, click **Launch vSphere Web Client** to start the vSphere Web Client and log in to vCenter Server.

15 Click **Finish** to close the installer.

vCenter Server, the vCenter Server components, and the Platform Services Controller are installed.

Install a Platform Services Controller on Windows

Before installing vCenter Server with an external Platform Services Controller, you install a Platform Services Controller. The Platform Services Controller contains the common services, such as vCenter Single Sign-On and the License service, which can be shared across several vCenter Server instances.

You can install many Platform Services Controllers of the same version and join them as replicating partners in the same vCenter Single Sign-On domain. Concurrent installations of replicating Platform Services Controllers are not supported. You must install the Platform Services Controllers in the domain in a sequence.

Important If you want to replace the VMCA-signed certificate with a CA-signed certificate, install the Platform Services Controller first, and then include VMCA in the certificate chain and generate new certificates from VMCA that are signed by the whole chain. You can then install vCenter Server. For information about managing vCenter Server certificates, see *Platform Services Controller Administration*.

Prerequisites

- Verify that your system meets the minimum software and hardware requirements. See [vCenter Server for Windows Requirements](#).
- [Download the vCenter Server Installer for Windows](#).

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vCenter Server for Windows** and click **Install**.
- 3 Follow the prompts of the installation wizard to review the welcome page and accept the license agreement.

- 4 Select **Platform Services Controller** and click **Next**.
- 5 Enter the system name, preferably an FQDN, and click **Next**.

You can also enter an IP address. If you enter an IP address, provide a static IP address.

Important When you provide an FQDN or an IP address as the system name of the Platform Services Controller, make sure that the FQDN or IP address does not change. If the FQDN or IP address of the host machine changes, you have to reinstall the Platform Services Controller and the vCenter Server instances registered with it. The FQDN or IP address of the Platform Services Controller is used to generate an SSL certificate for the Platform Services Controller host machine.

6 Create a new vCenter Single Sign-On domain or join an existing domain.

Option	Description
Create a new Single Sign-On domain	<p>Creates a new vCenter Single Sign-On domain.</p> <ol style="list-style-type: none"> Enter the domain name, for example vsphere.local. Set the user name for the vCenter Single Sign-On administrator account, for example, administrator. <p>After the deployment, you can log in to vCenter Single Sign-On and to vCenter Server as <i>administrator_user_name@your_domain_name</i>.</p> <ol style="list-style-type: none"> Set the password for the vCenter Single Sign-On administrator account. <p>This is the password for the user <i>administrator_user_name@your_domain_name</i>.</p> <ol style="list-style-type: none"> Enter the site name for vCenter Single Sign-On. <p>The site name is important if you are using vCenter Single Sign-On in multiple locations. The site name must contain alphanumeric characters. Choose your own name for the vCenter Single Sign-On site. You cannot change the name after installation.</p> <p>Extended ASCII and non-ASCII characters are unsupported in site names. Your site name must include alphanumeric characters and a comma (,), period (.), question mark (?), dash (-), underscore (_), plus sign (+) or equals sign (=).</p> <ol style="list-style-type: none"> Click Next.
Join an existing vCenter Single Sign-On domain	<p>Joins a new vCenter Single Sign-On server to a vCenter Single Sign-On domain in an existing Platform Services Controller. You must provide the information about the vCenter Single Sign-On server to which you join the new vCenter Single Sign-On server.</p> <ol style="list-style-type: none"> Enter the fully qualified domain name (FQDN) or IP address of the Platform Services Controller that contains the vCenter Single Sign-On server to join. Enter the HTTPS port to use for communication with the Platform Services Controller. Enter the user name and password of the vCenter Single Sign-On administrator account. Click Next. Approve the certificate provided by the remote machine, and you must select whether to create or join an existing vCenter Single Sign-On site. Select whether to create or join an existing vCenter Single Sign-On site.

7 Click **Next**.

8 For each component, accept the default port numbers, or if another service is using the defaults, enter alternative ports, and click **Next**.

Make sure that ports 80 and 443 are free and dedicated, so that vCenter Single Sign-On can use these ports. Otherwise, use custom ports during installation.

9 (Optional) Change the default destination folders and click **Next**.

Important Do not use folders that end with an exclamation mark (!).

- 10 Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.

For information about the CEIP, see the *Configuring Customer Experience Improvement Program* section in *vCenter Server and Host Management*.

- 11 Review the summary of the installation settings and click **Install** to start the installation.
- 12 After the installation completes, click **Finish** to close the installer.

The Platform Services Controller is installed.

What to do next

Install vCenter Server on another Windows virtual machine or physical server and register vCenter Server and the vCenter Server components to the Platform Services Controller.

Install vCenter Server with an External Platform Services Controller on Windows

After you install a Platform Services Controller on a Windows host machine, or deploy a Platform Services Controller appliance, you can install vCenter Server and the vCenter Server components and connect the vCenter Server instance to the deployed Platform Services Controller.

Prerequisites

- Verify that your system meets the minimum software and hardware requirements. See [vCenter Server for Windows Requirements](#).
- [Download the vCenter Server Installer for Windows](#).
- If you want to use the vSphere Web Client on the host machine on which you install vCenter Server, verify that Adobe Flash Player version 11.9 or later is installed on the system.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vCenter Server for Windows** and click **Install**.
- 3 Follow the prompts of the installation wizard to review the welcome page and accept the license agreement.
- 4 Select **vCenter Server** and click **Next**.
- 5 Enter the system network name, preferably a static IP address, and click **Next**.

Important The name that you type is encoded in the SSL certificate of the system. The components communicate with each other by using this name. The system name must be either a static IP address or a fully qualified domain name (FQDN). Make sure that the system name does not change. You cannot change the system name after the installation completes.

- 6 Provide the system name of the Platform Services Controller that you already installed or deployed, the HTTPS port to use for communication with the vCenter Single Sign-On server, as well as the vCenter Single Sign-On password, and click **Next**.

Important Make sure that you use either the IP address or the FQDN that you provided during the installation of the Platform Services Controller. If you provided the FQDN as a system name of the Platform Services Controller, you cannot use an IP address, and the reverse. When a service from vCenter Server connects to a service running in the Platform Services Controller, the certificate is verified. If the IP address or FQDN changes, the verification fails and vCenter Server cannot connect to the Platform Services Controller.

- 7 Approve the certificate provided by the remote machine.
- 8 Select the vCenter Server service account and click **Next**.

Note Starting with vSphere 6.5, the vCenter Server services are not standalone services under Windows SCM, instead they run as child processes of the VMware Service Lifecycle Manager service.

Option	Description
Use Windows Local System Account	The vCenter Server service runs in the Windows Local System account. This option prevents you from connecting to an external database by using Windows integrated authentication.
Specify a user service account	The vCenter Server service runs in an administrative user account with a user name and password that you provide. Important The user credentials that you provide must be of a user who is in the local administrator group and who has the Log on as a service privilege.

- 9 Select the type of database that you want to use and click **Next**.

Option	Description
Use an embedded database (PostgreSQL)	vCenter Server uses the embedded PostgreSQL database. This database is suitable for small scale deployments.
Use an external database	vCenter Server uses an existing external database. a Select your database from the list of available DSNs. b Type the user name and the password for the DSN. If your database uses Windows NT authentication, the user name and password text boxes are disabled.

- 10 For each component, accept the default port numbers, or if another service is using the defaults, enter alternative ports, and click **Next**.
- 11 (Optional) Change the default destination folders and click **Next**.

Important Do not use folders that end with an exclamation mark (!).

- 12 Review the summary of the installation settings and click **Install** to start the installation.

- 13 (Optional) After the installation finishes, click **Launch vSphere Web Client** to start the vSphere Web Client and log in to vCenter Server.
- 14 Click **Finish** to close the installer.

vCenter Server is installed in evaluation mode. You can activate vCenter Server by using the vSphere Web Client. For information about activating vCenter Server, see *vCenter Server and Host Management*.

Installing vCenter Server in an Environment with Multiple NICs

If you want to install vCenter Server with an external Platform Services Controller in an environment with multiple NICs, you must keep a record of the IP addresses or FQDNs that you use as system network names.

For example, if you want to install a Platform Services Controller on one virtual machine and vCenter Server on another virtual machine and each virtual machine has two NICs, you can use the following workflow:

- 1 Install a Platform Services Controller on one of the virtual machines and use one of its IP addresses or FQDNs as a system network name.
- 2 On the other virtual machine, start the installation of vCenter Server and use one of its IP addresses or FQDNs as a system network name.
- 3 When prompted to provide the system network name of the Platform Services Controller, enter the IP address or FQDN that you entered during the installation of the Platform Services Controller.
If you enter the other IP address or FQDN of the Platform Services Controller, you receive an error message.
- 4 After the installation completes, you can log in to the vSphere Web Client by using either of the NIC IP addresses or FQDNs of vCenter Server.

After You Install vCenter Server or Deploy the vCenter Server Appliance

5

After you install vCenter Server or deploy the vCenter Server Appliance, consider these postinstallation options before adding inventory for the vCenter Server to manage.

For information about configuring the vSphere Authentication Proxy service, see *vSphere Security*.

This section includes the following topics:

- [Log in to vCenter Server by Using the vSphere Web Client](#)
- [Install the VMware Enhanced Authentication Plug-in](#)
- [Collect vCenter Server Log Files](#)
- [Repoint vCenter Server to Another External Platform Services Controller](#)
- [Reconfigure a Standalone vCenter Server with an Embedded Platform Services Controller to a vCenter Server with an External Platform Services Controller](#)

Log in to vCenter Server by Using the vSphere Web Client

Log in to vCenter Server by using the vSphere Web Client to manage your vSphere inventory.

In vSphere 6.0 and later, the vSphere Web Client is installed as part of the vCenter Server on Windows or the vCenter Server Appliance deployment. This way, the vSphere Web Client always points to the same vCenter Single Sign-On instance.

Procedure

- 1 Open a Web browser and enter the URL for the vSphere Web Client:
`https://vcenter_server_ip_address_or_fqdn/vsphere-client.`
- 2 Enter the credentials of a user who has permissions on vCenter Server, and click **Login**.

- 3 If a warning message about an untrusted SSL certificate appears, select the appropriate action based on your security policy.

Option	Action
Ignore the security warning for this login session only.	Click Ignore .
Ignore the security warning for this login session, and install the default certificate so that the warning does not appear again.	Select Install this certificate and do not display any security warnings for this server and click Ignore . Select this option only if using the default certificate does not present a security problem in your environment.
Cancel and install a signed certificate before proceeding.	Click Cancel and ensure that a signed certificate is installed on the vCenter Server system before you attempt to connect again.

The vSphere Web Client connects to all the vCenter Server systems on which the specified user has permissions, allowing you to view and manage your inventory.

Install the VMware Enhanced Authentication Plug-in

The VMware Enhanced Authentication Plug-in provides Integrated Windows Authentication and Windows-based smart card functionality.

In this vSphere 6.5 release, the VMware Enhanced Authentication Plug-in replaces the Client Integration Plug-in from vSphere 6.0 releases and earlier. The Enhanced Authentication Plug-in provides Integrated Windows Authentication and Windows-based smart card functionality. These are the only two features carried over from the previous Client Integration Plug-in. The Enhanced Authentication Plug-in can function seamlessly if you already have the Client Integration Plug-in installed on your system from vSphere 6.0 or earlier. There are no conflicts if both plug-ins are installed.

Watch the video "vSphere Web Client after the Client Integration Plug-in Removal" for more information about the workflow changes to the vSphere Client:



vSphere Web Client after the Client Integration Plug-in Removal
(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_web_client_after_cip_removal)

Install the plug-in only once to enable all the functionality the plug-in delivers.

If you install the plug-in from an Internet Explorer browser, you must first disable Protected Mode and enable pop-up windows on your Web browser. Internet Explorer identifies the plug-in as being on the Internet instead of on the local intranet. In such cases, the plug-in is not installed correctly because Protected Mode is enabled for the Internet.

For information about supported browsers and operating systems, see the *vSphere Installation and Setup* documentation.

Prerequisites

If you use Microsoft Internet Explorer, disable Protected Mode.

Procedure

- 1 Open a Web browser and type the URL for the vSphere Web Client.
- 2 At the bottom of the vSphere Web Client login page, click **Download Enhanced Authentication Plug-in**.
- 3 If the browser blocks the installation either by issuing certificate errors or by running a pop-up blocker, follow the Help instructions for your browser to resolve the problem.
- 4 Save the plug-in to your computer, and run the executable.
- 5 Step through the installation wizard for both the VMware Enhanced Authentication Plug-in and the VMware Plug-in Service which are run in succession.
- 6 When the installations are complete, refresh your browser.
- 7 On the External Protocol Request dialog box, click **Launch Application** to run the Enhanced Authentication Plug-in.

The link to download the plug-in disappears from the login page.

Collect vCenter Server Log Files

After you install vCenter Server, you can collect the vCenter Server log files for diagnosing and troubleshooting purposes.

Note This procedure provides information about how to collect the log files for a Windows installation of vCenter Server. For information about exporting a support bundle and browsing the log files in the vCenter Server Appliance, see *vCenter Server Appliance Configuration*.

Procedure

- 1 Log in as an administrator on the Windows machine where vCenter Server is installed.
- 2 Navigate to **Start > Programs > VMware > Generate vCenter Server log bundle** to generate the log bundle.

You can generate vCenter Server log bundles even if you are unable to connect to the vCenter Server by using the vSphere Web Client

The log files for the vCenter Server system are generated and saved in a .tgz archive on your desktop.

Repoint vCenter Server to Another External Platform Services Controller

Joining external Platform Services Controller instances in the same vCenter Single Sign-On domain, ensures high availability of your system.

If an external Platform Services Controller stops responding or if you want to distribute the load of an external Platform Services Controller, you can repoint the vCenter Server instances to another Platform Services Controller in the same domain and site.

- You can repoint the vCenter Server instance to an existing functional Platform Services Controller instance with free load capacity in the same domain and site.
- You can install or deploy a new Platform Services Controller instance in the same domain and site to which to repoint the vCenter Server instance.

Prerequisites

- If the old Platform Services Controller instance has stopped responding, remove the node and clean up the stale vmdir data by running the `cmsso-util unregister` command. For information about decommissioning a Platform Services Controller instance, see <https://kb.vmware.com/kb/2106736>.
- Verify that the old and the new Platform Services Controller instances are in the same vCenter Single Sign-On domain and site by running the `vdcrcpadmin -f showservers` command. For information about using the command, see <https://kb.vmware.com/kb/2127057>.
- If you want to repoint a vCenter Server Appliance that is configured in a vCenter HA cluster, remove the vCenter HA configuration. For information about removing a vCenter HA configuration, see *vSphere Availability*.

Procedure

- 1 Log in to the vCenter Server instance.
 - For a vCenter Server Appliance, log in to the vCenter Server Appliance shell as root.
 - For a vCenter Server instance on Windows, log in as an administrator to the vCenter Server virtual machine or physical server.
- 2 If the vCenter Server instance runs on Windows, in the Windows command prompt, navigate to `C:\Program Files\VMware\vCenter Server\bin`.
- 3 Run the `cmsso-util repoint` command.

```
cmsso-util repoint --repoint-psc psc_fqdn_or_static_ip [--dc-port port_number]
```

where the square brackets [] enclose the command options.

Here, *pvc_fqdn_or_static_ip* is the system name used to identify the Platform Services Controller. This system name must be an FQDN or a static IP address.

Note The FQDN value is case-sensitive.

Use the `--dc-port port_number` option if the Platform Services Controller runs on a custom HTTPS port. The default value of the HTTPS port is 443.

- 4 Log in to the vCenter Server instance by using the vSphere Web Client to verify that the vCenter Server instance is running and can be managed.

The vCenter Server instance is registered with the new Platform Services Controller.

What to do next

If you repointed a vCenter Server Appliance that was configured in a vCenter HA cluster, you can reconfigure the vCenter HA cluster. For information about configuring vCenter HA, see *vSphere Availability*.

Reconfigure a Standalone vCenter Server with an Embedded Platform Services Controller to a vCenter Server with an External Platform Services Controller

If you have deployed or installed a standalone vCenter Server instance with an embedded Platform Services Controller and you want to extend your vCenter Single Sign-On domain with more vCenter Server instances, you can reconfigure and repoint the existing vCenter Server instance to an external Platform Services Controller.

Figure 5-1. Reconfiguration of a Standalone vCenter Server Instance with an Embedded Platform Services Controller and Repointing it to an External Platform Services Controller

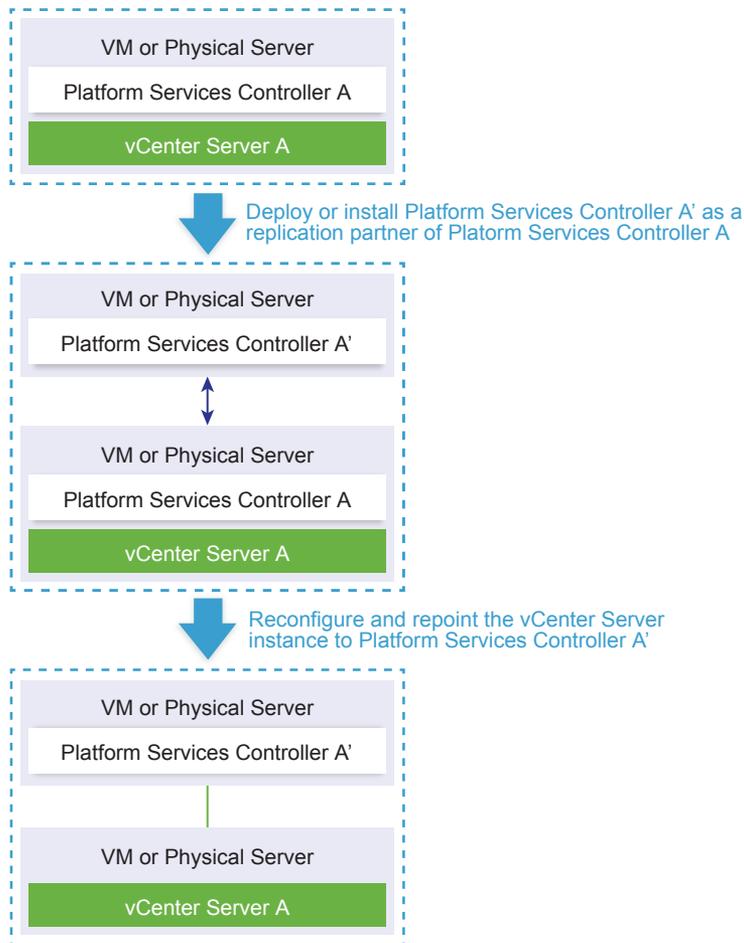


Table 5-1. Legend

Arrow or line	Description
	Replication agreement between two Platform Services Controller instances
	vCenter Server registration with an external Platform Services Controller
	Transition step

Note The reconfiguration of a vCenter Server instance with an embedded Platform Services Controller and repointing it to an external Platform Services Controller instance is a one-way process after which you cannot switch back to vCenter Server with an embedded Platform Services Controller.

Prerequisites

- Deploy or install the external Platform Services Controller instance as a replication partner of the existing embedded Platform Services Controller instance in the same vCenter Single Sign-On site.

Note You can determine the current vCenter Single Sign-On site by using the `vmfad-cli` command.

- For a vCenter Server Appliance with an embedded Platform Services Controller, log in to the appliance shell as root and run the command.

```
/usr/lib/vmware-vmafd/bin/vmafd-cli get-site-name --server-name localhost
```

- For a Windows installation of vCenter Server instance with an embedded Platform Services Controller, log in to the Windows machine as an administrator, open the Windows command prompt, and run the command.

```
C:\Program Files\VMware\vCenter Server\vmafdd\vmafd-cli get-site-name --server-name localhost
```

- Create snapshots of the vCenter Server with an embedded Platform Services Controller and the external Platform Services Controller instance, so that you can revert to the snapshots if the reconfiguration fails.
- If you want to reconfigure a vCenter Server Appliance with an embedded Platform Services Controller that is configured in a vCenter HA cluster, remove the vCenter HA configuration. For information about removing a vCenter HA configuration, see *vSphere Availability*.

Procedure

- 1 Log in to the vCenter Server instance with an embedded Platform Services Controller.

Option	Steps
For a vCenter Server Appliance with an embedded Platform Services Controller	Log in to the appliance shell as root. <ul style="list-style-type: none"> ■ If you have direct access to the appliance console, press Alt+F1. ■ If you want to connect remotely, use SSH or another remote console connection to start a session to the appliance.
For a Windows installation of vCenter Server with an embedded Platform Services Controller	Log in to the Windows machine as an administrator, open the Windows command prompt, and navigate to C:\Program Files\VMware\vCenter Server\bin.

- 2 Verify that all Platform Services Controller services are running.

Run the `service-control --status --all` command.

The Platform Services Controller services that must be running are VMware License Service, VMware Identity Management Service, VMware Security Token Service, VMware Certificate Service, and VMware Directory Service.

- 3 Run the `cmsso-util reconfigure` command.

```
cmsso-util reconfigure --repoint-psc psc_fqdn_or_static_ip --username username --domain-name domain_name --passwd password [--dc-port port_number]
```

where the square brackets [] enclose optional items.

Here, *psc_fqdn_or_static_ip* is the system name used to identify the external Platform Services Controller instance. This system name must be an FQDN or a static IP address.

Note The FQDN value is case-sensitive.

The options *username* and *password* are the administrator user name and password of the vCenter Single Sign-On *domain_name*.

Use the `--dc-port` option if the external Platform Services Controller runs on a custom HTTPS port. The default value of the HTTPS port is 443.

For example, if the external Platform Services Controller runs on a custom HTTPS port 449, you must run:

```
cmsso-util reconfigure --repoint-psc psc.acme.local --username administrator --domain-name vsphere.local --passwd Password1! --dc-port 449
```

- 4 Log in to the vCenter Server instance by using the vSphere Web Client to verify that the vCenter Server instance is running and can be managed.

The vCenter Server with an embedded Platform Services Controller is demoted, and the vCenter Server is redirected to the external Platform Services Controller.

What to do next

- You can deploy or install additional vCenter Server and Platform Services Controller instances in the vCenter Single Sign-On domain.
- If you reconfigured a vCenter Server Appliance with an embedded Platform Services Controller that was configured in a vCenter HA cluster, you can reconfigure the vCenter HA cluster. For information about configuring vCenter HA, see *vSphere Availability*.

File-Based Backup and Restore of vCenter Server Appliance

6

The vCenter Server Appliance supports a file-based backup and restore mechanism that helps you to recover your environment after failures.

In vSphere 6.5, you can use the vCenter Server Appliance Management Interface to create a file-based backup of the vCenter Server Appliance and Platform Services Controller appliance. After you create the backup, you can restore it by using the GUI installer of the appliance.

You use the vCenter Server Appliance Management Interface to perform a file-based backup of the vCenter Server core configuration, inventory, and historical data of your choice. The backed-up data is streamed over FTP, FTPS, HTTP, HTTPS, or SCP to a remote system. The backup is not stored on the vCenter Server Appliance.

You can perform a file-based restore only for a vCenter Server Appliance that you have previously backed up by using the vCenter Server Appliance Management Interface. You can perform such restore operation by using the GUI installer of the vCenter Server Appliance. The process consists of deploying a new vCenter Server Appliance and copying the data from the file-based backup to the new appliance.

You can also perform a restore operation by deploying a new vCenter Server Appliance and using the vCenter Server Appliance management interface to copy the data from the file-based backup to the new appliance.

Important If you back up a vCenter Server Appliance High Availability cluster, the backup operation only backs up the primary vCenter Server instance. Before restoring a vCenter Server Appliance High Availability cluster, you must power off the active, passive, and witness nodes. The restore operation restores the vCenter Server in non-vCenter Server High Availability mode. You must reconstruct the cluster after the restore operation completes successfully.

This section includes the following topics:

- [Considerations and Limitations for File-Based Backup and Restore](#)
- [Back up a vCenter Server Appliance by Using the vCenter Server Appliance Management Interface](#)
- [Restore a vCenter Server Appliance from a File-Based Backup](#)

Considerations and Limitations for File-Based Backup and Restore

When you backup or restore a vCenter Server environment, take into account these considerations and limitation.

Protocols

The following considerations apply to file-based backup and restore protocols:

- FTP and HTTP are not secure protocols
- Backup servers must support minimum of 10 simultaneous connections for each vCenter Server Appliance
- You must have write permissions for upload and read permissions for download
- Only explicit mode is supported for FTPS
- If you use HTTP or HTTPS, you must enable WebDAV on the backup Web server
- You can use only FTP, FTPS, HTTP, or HTTPS to transmit data through an HTTP proxy server
- You can use IPv4 and IPv6 URLs in file-based backup and restore of a vCenter Server Appliance. Mixed mode of IP versions between the backup server and the vCenter Server Appliance is unsupported.
- If you use SCP protocol, you must use the Linux backup server

Configuration

After a restore, the following configurations revert to the state when the backup was taken.

- Virtual machine resource settings
- Resource pool hierarchy and setting
- Cluster-host membership
- DRS configuration and rules

Storage DRS

If the configuration changes, the following might change after a restore.

- Datastore Cluster configuration
- Datastore Cluster membership
- Datastore I/O Resource Management (Storage I/O Control) settings
- Datastore-Datacenter membership
- Host-Datastore membership

Distributed Power Management

If you put a host into standby mode after a backup, the vCenter Server might force the host to exit standby mode when you restore to the backup.

Distributed Virtual Switch

If you use a distributed virtual switch, you are advised to export separately the distributed virtual switch configuration before you restore to a backup. You can import the configuration after the restore. If you omit this consideration, you may lose the changes made to a distributed virtual switch after the backup. For detailed steps, see the VMware knowledge base article at <http://kb.vmware.com/kb/2034602>.

Content Libraries

If you delete libraries or items after a backup, you cannot access or use these libraries or items after the restore. You can only delete such libraries or items. A warning message notifies you that there are missing files or folders in the storage backup.

If you create new items or item files after the backup, the Content Library Service has no record of the new items or files after the restore operation. A warning notifies you that extra folders or files were found on the storage backup.

If you create new libraries after the backup, the Content Library Service has no record of the new libraries after restore. The library content exists on the storage backing, but no warning is displayed. You must manually clean the new libraries.

Virtual Machine Life Cycle Operations

- Restoring vCenter Server from a backup that was taken during in-flight relocation operations in the vCenter Server instance.

After you restore vCenter Server, the vCenter Server view of the virtual machines might be out of sync with the ESXi view of the virtual machines. This is also true if you performed the backup during in-flight operations on vCenter Server. If virtual machines disappear after you restore vCenter Server, you can refer to the following cases.

- a The missing virtual machine is located on the destination ESXi host and is registered with the destination ESXi host, but it is either an orphan or not in the vCenter Server inventory. You must manually add the virtual machine to the vCenter Server inventory.
- b The missing virtual machine is located on the destination ESXi host, but it is not registered with the destination ESXi host and it is not in the vCenter Server inventory. You must manually register the virtual machine to the ESXi host and add the virtual machine back to the vCenter Server inventory.

- c The missing virtual machine is located on the destination ESXi host, but it is not registered with the destination ESXi host. In the vCenter Server instance, the missing virtual machine is marked as orphaned. You must remove the virtual machine from the vCenter Server inventory and add it again.
- Restoring vCenter Server from a backup that has an out-of-date linked clone virtual machine layout.
If you create a linked clone virtual machine after the backup and you restore vCenter Server from the old backup, then after the restore, the vCenter Server does not know about the new linked clone virtual machine until vCenter Server discovers the new linked clone virtual machine. If you remove all existing virtual machines before the new linked clone virtual machine is discovered, then the removal of existing virtual machines corrupts the new linked clone due to missing disks. In order to avoid this, you must wait until all linked clone virtual machines are discovered by the vCenter Server before you remove virtual machines.
- Restoring vCenter Server from a backup that was taken during virtual machine registration.
If you are registering a virtual machine during the backup and you restore vCenter Server from the old backup, then after the restore, the virtual machine is marked as orphaned in the vCenter Server instance. You must manually add the virtual machine to the vCenter Server inventory.

vSphere High Availability

Restoring vCenter Server from a backup might cause it to rollback to older version for the vSphere HA cluster state (HostList, ClusterConfiguration, VM protection state) while the hosts in the cluster have the latest version for the cluster state. You need to make sure the vSphere HA cluster state stays the same during restore and backup operations. Otherwise, the following problems might occur.

- If hosts are added or removed to or from the vSphere HA cluster after backup and before vCenter Server restore, virtual machines could potentially failover to hosts not being managed by the vCenter Server but are still part of the HA cluster.
- Protection state for new virtual machines is not updated on the vSphere HA agents on the hosts that are part of the vSphere HA cluster. As a result, virtual machines are not protected or unprotected.
- New cluster configuration state is not updated on the vSphere HA agents on the hosts that are part of the vSphere HA cluster.

vCenter High Availability

Restoring vCenter Server requires vCenter HA to be reconfigured.

Storage Policy Based Management

Restoring vCenter Server from a backup can lead to the following inconsistencies related to storage policies, storage providers, and virtual machines.

- Registered storage providers after backup are lost.
- Unregistered storage providers after backup re-appear and might show different provider status.

- Changes, such as create, delete, or update, performed on storage policies after backup are lost.
- Changes, such as create, delete, or update, performed on storage policy components after backup are lost.
- Default policy configuration changes for datastores performed after backup are lost.
- Changes in the storage policy association of the virtual machine and its disks, and in their policy compliance might occur.

Virtual Storage Area Network

Restoring vCenter Server from a backup might cause inconsistencies in the vSAN. For information on how to check vSAN health, see *Administering VMware vSAN*.

Patching

Restoring vCenter Server from a backup might result in missing security patches. You must apply them again after the restore is complete. For information on patching the vCenter Server Appliance, see *vSphere Upgrade*.

Back up a vCenter Server Appliance by Using the vCenter Server Appliance Management Interface

You can use the vCenter Server Appliance Management Interface to back up the vCenter Server instance. You can select whether to include historical data, such as stats, events, and tasks, in the backup file.

Note The backup operation for a vCenter High Availability cluster, backs up only the active node.

Prerequisites

- You must have an FTP, FTPS, HTTP, HTTPS, or SCP server up and running with sufficient disk space to store the backup.
- Dedicate a separate folder on your server for each file-based backup.

Procedure

- 1 In a Web browser, go to the vCenter Server Appliance Management Interface, <https://appliance-IP-address-or-FQDN:5480>.
- 2 Log in as root.
- 3 In the vCenter Server Appliance Management Interface, click **Summary**.
- 4 Click **Backup**.

The **Backup Appliance** wizard opens.

5 Enter the backup protocol and location details.

Option	Description
Backup protocol	Select the protocol to use to connect to your backup server. You can select FTP, FTPS, HTTP, HTTPS, or SCP. For FTP, FTPS, HTTP, or HTTPS the path is relative to the home directory configured for the service. For SCP, the path is absolute to the remote systems root directory.
Backup location	Enter the server address and backup folder in which to store the backup files.
Port	Enter the default or custom port of the backup server.
User name	Enter a user name of a user with write privileges on the backup server.
Password	Enter the password of the user with write privileges on the backup server.

6 (Optional) Select **Encrypt Backup Data** to encrypt your backup file and enter a password for the encryption.

If you select to encrypt the backup data, you must use the encryption password for the restore procedure.

7 Click **Next**.

8 On the Select parts to backup page, review the data that is backed up by default.

9 (Optional) Select **Stats, Events, and Tasks** to back up additional historical data from the database.

10 (Optional) In the **Description** text box, enter a description of the backup and click **Next**.

11 On the Ready to complete page, review the summary information for the backup and click **Finish**.

The **Backup Progress** window opens and indicates the progress of the backup operation.

12 After the backup process finishes, click **OK** to close the **Backup Progress** window.

You successfully created a backup file of the vCenter Server Appliance.

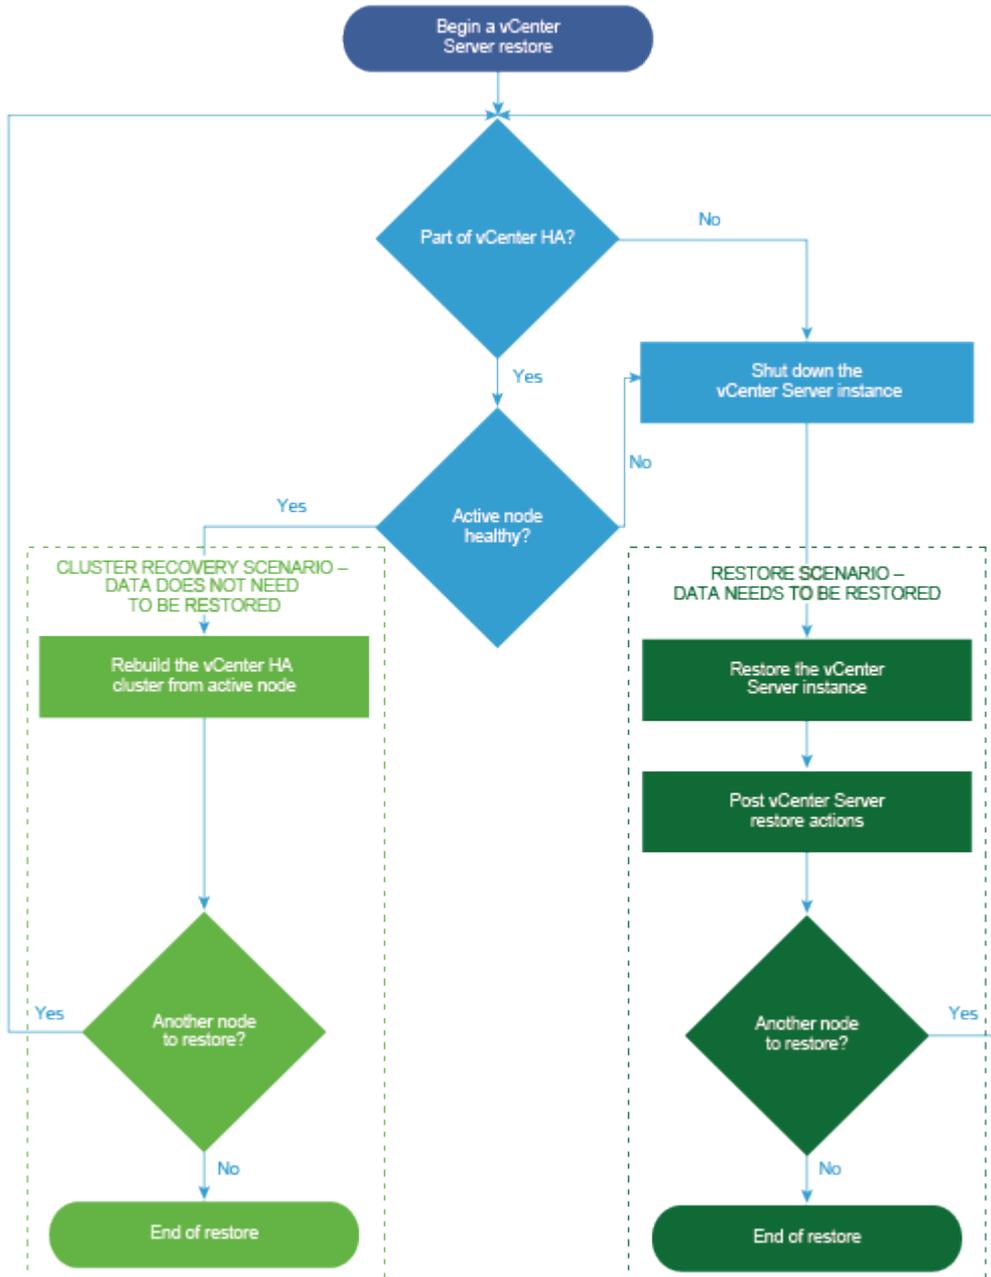
What to do next

If your file-based backup has failed, cancel the backup job.

Restore a vCenter Server Appliance from a File-Based Backup

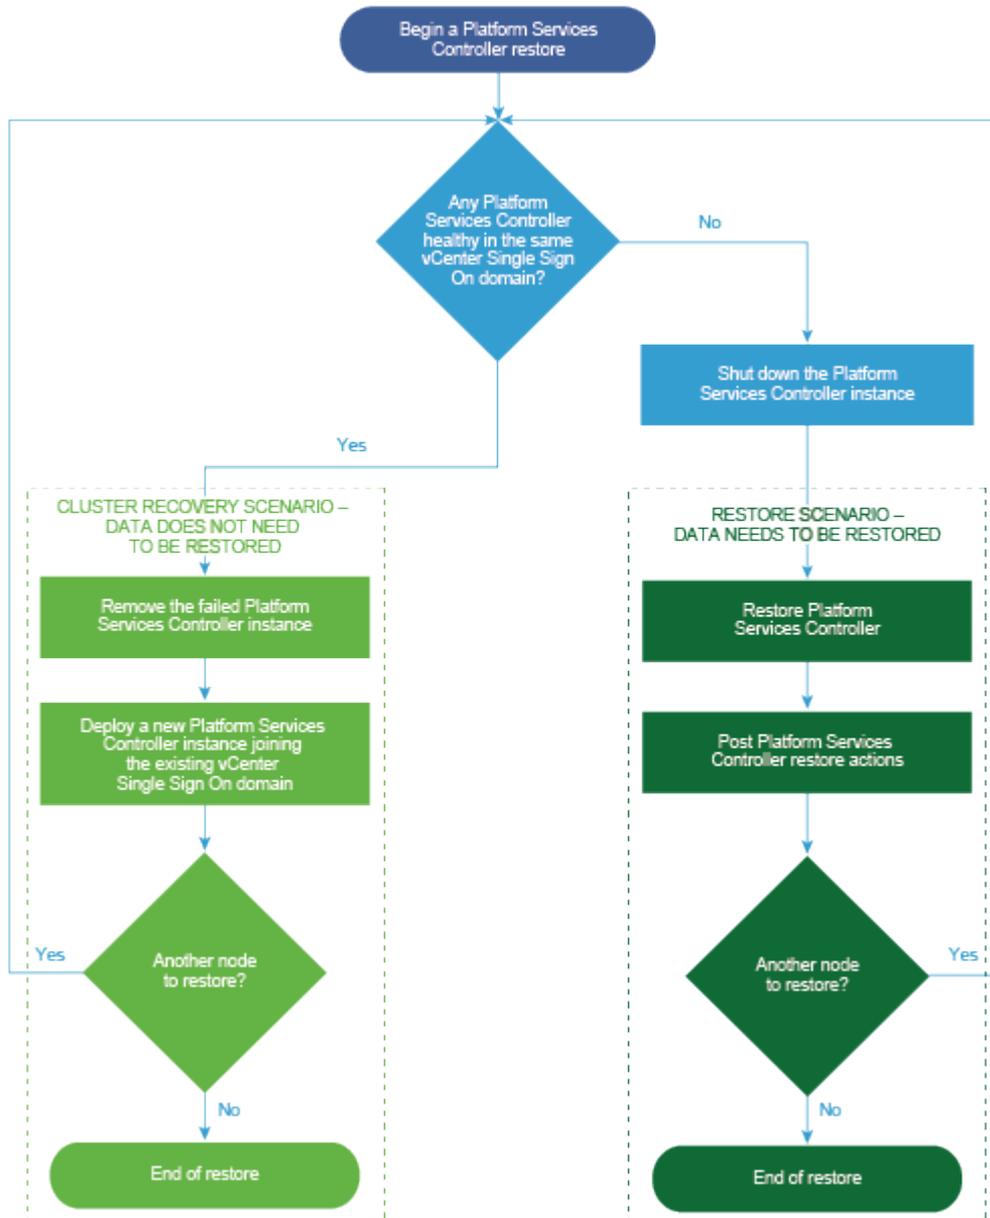
You can use the vCenter Server Appliance GUI installer to restore a vCenter Server Appliance to an ESXi host or a vCenter Server instance. The restore procedure has two stages. The first stage deploys a new vCenter Server Appliance. The second stage populates the newly deployed vCenter Server Appliance with the data stored in the file-based backup.

Figure 6-1. vCenter Server Appliance Restore Workflow



Perform a file-based restore of a Platform Services Controller only when the last Platform Services Controller in the domain fails. If there are other Platform Services Controller instances in the same vCenter Single Sign-On domain, deploy a new Platform Services Controller instance and join it to the existing Single Sign-On domain.

Figure 6-2. Platform Services Controller Appliance Restore Workflow



Prerequisites

- Verify that your system meets the minimum software and hardware requirements. See [System Requirements for the vCenter Server Appliance and Platform Services Controller Appliance](#).
- [Download and Mount the vCenter Server Appliance Installer](#).

- If the vCenter Server instance is part of a vCenter High Availability cluster, you must power off the active, passive, and witness nodes of the cluster before restoring the vCenter Server.

Procedure

1 Stage 1 - Deploy a New Appliance

In stage 1 of the restore process, you deploy the OVA file, which is included in the vCenter Server Appliance GUI installer.

2 Stage 2 - Transfer Data to the Newly Deployed Appliance

After the OVA deployment finishes, you are redirected to stage 2 of the restore process in which the data from the backup location is copied to the newly deployed vCenter Server Appliance.

Stage 1 - Deploy a New Appliance

In stage 1 of the restore process, you deploy the OVA file, which is included in the vCenter Server Appliance GUI installer.

As an alternative to performing the first stage of the restore with the GUI installer, you can deploy the OVA file of the new vCenter Server Appliance or Platform Services Controller appliance by using the vSphere Web Client or VMware Host Client. To deploy the OVA file on an ESXi host or vCenter Server instance 5.5 or 6.0, you can also use the vSphere Client. After the OVA deployment, you must log in to the appliance management interface of the newly deployed appliance to proceed with the second stage of the restore process.

Prerequisites

- Download and mount the vCenter Server Appliance installer. See [Download and Mount the vCenter Server Appliance Installer](#).
- If you plan to restore the vCenter Server Appliance on an ESXi host, verify that the target ESXi host is not in lockdown or maintenance mode.
- If you plan to restore the vCenter Server Appliance on a DRS cluster of a vCenter Server inventory, verify that the cluster contains at least one ESXi host that is not in lockdown or maintenance mode.
- If you plan to assign a static IP address to the appliance, verify that you have configured the forward and reverse DNS records for the IP address.
- If you are attempting to restore a vCenter Server instance that is still running, power off the backed up vCenter Server before you start the restore operation.

Procedure

- 1 In the vCenter Server Appliance installer, navigate to the `vcsa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.

- 2 On the Home page, click **Restore**.
- 3 Review the Introduction page to understand the restore process and click **Next**.
- 4 Read and accept the license agreement, and click **Next**.
- 5 On the Enter backup details page, enter the details of the backup file that you want to restore, and click **Next**.

Option	Description
Backup location type	Select the protocol to use to retrieve the backup from your backup server. You can select HTTPS, HTTP, SCP, FTPS, or FTP.
Backup location	Enter the server address and backup folder where the backup files are stored.
Port	Enter the default or custom port of the backup server.
User name	Enter the user name of a user with read privileges on the backup server.
Password	Enter the password of the user with read privileges on the backup server.
Encryption password	If the backup file was encrypted, enter the encryption password.

- 6 Review the backup information, and click **Next**.
- 7 Connect to the ESXi host or vCenter Server on which you want to deploy the vCenter Server Appliance to use for the restore operation.

Option	Steps
You can connect to an ESXi host on which to deploy the appliance to use for the restore operation.	<ol style="list-style-type: none"> 1 Enter the FQDN or IP address of the ESXi host. 2 Enter the HTTPS port of the ESXi host. 3 Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. 4 Click Next. 5 Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target ESXi host, and click Yes to accept the certificate thumbprint.
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the appliance to use for the restore operation.	<ol style="list-style-type: none"> 1 Enter the FQDN or IP address of the vCenter Server instance. 2 Enter the HTTPS port of the vCenter Server instance. 3 Enter the user name and password of user with vCenter Single Sign-On administrative privileges on the vCenter Server instance, for example, the administrator@your_domain_name user. 4 Click Next. 5 Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target vCenter Server instance, and click Yes to accept the certificate thumbprint. 6 Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next. <p>Note You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode.</p> <ol style="list-style-type: none"> 7 Select the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next.

- 8 Accept the certificate warning.
- 9 Enter a name for the vCenter Server Appliance, set up the password for the root user, and click **Next**.

- 10 Select the deployment size for the new vCenter Server Appliance depending on the size of your vSphere inventory.

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 CPUs and 10 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 CPUs and 16 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 CPUs and 24 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines
Large	Deploys an appliance with 16 CPUs and 32 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 CPUs and 48 GB of memory. Suitable for environments with up to 2,000 hosts or 35,000 virtual machines

- 11 Select the storage size for the new vCenter Server Appliance, and click **Next**.

Important You must consider the storage size of the appliance that you are restoring.

Storage Size Option	Description for Tiny Deployment Size	Description for Small Deployment Size	Description for Medium Deployment Size	Description for Large Deployment Size	Description for X-Large Deployment Size
Default	Deploys an appliance with 250 GB of storage.	Deploys an appliance with 290 GB of storage.	Deploys an appliance with 425 GB of storage.	Deploys an appliance with 640 GB of storage.	Deploys an appliance with 980 GB of storage.
Large	Deploys an appliance with 775 GB of storage.	Deploys an appliance with 820 GB of storage.	Deploys an appliance with 925 GB of storage.	Deploys an appliance with 990 GB of storage.	Deploys an appliance with 1030 GB of storage.
X-Large	Deploys an appliance with 1650 GB of storage.	Deploys an appliance with 1700 GB of storage.	Deploys an appliance with 1805 GB of storage.	Deploys an appliance with 1870 GB of storage.	Deploys an appliance with 1910 GB of storage.

- 12 From the list of available datastores, select the location where all the virtual machine configuration files and virtual disks will be stored and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**.
- 13 On the Configure network settings page review the settings populated from the backup file of the vCenter Server Appliance.
- 14 (Optional) Edit the network configuration to match the current network environment where the vCenter Server Appliance is restored.
- 15 On the Ready to complete stage 1 page, review the deployment settings for the restored vCenter Server Appliance and click **Finish** to start the OVA deployment process.

- 16 Wait for the OVA deployment to finish, and click **Continue** to proceed with stage 2 of the restore process to transfer the data to the newly deployed appliance.

Note If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Appliance Management Interface to transfer the data.

The newly deployed vCenter Server Appliance is running on the target server but the data is not copied from the backup location.

Stage 2 - Transfer Data to the Newly Deployed Appliance

After the OVA deployment finishes, you are redirected to stage 2 of the restore process in which the data from the backup location is copied to the newly deployed vCenter Server Appliance.

Procedure

- 1 Review the introduction to stage 2 of the restore process and click **Next**.
- 2 Review the backup details and click **Next**.
- 3 On the Ready to complete page, review the details, click **Finish**, and click **OK** to complete stage 2 of the restore process.

The restore process restarts the vCenter Server Appliance Management Service. You cannot access the vCenter Server Appliance Management API during the restart.

Important If a restore operation of a vCenter Server Appliance or a Platform Services Controller appliance VM results with a failure, you must power off and delete the partially restored VM. After that you can try to restore the VM again.

- 4 (Optional) After the restore process finishes, click the https://vcenter_server_appliance_fqdn/vsphere-client to go to the vSphere Web Client and log in to the vCenter Server instance in the vCenter Server Appliance, or click the https://vcenter_server_appliance_fqdn:443 to go to the vCenter Server Appliance Getting Started page.

- 5 Click **Close** to exit the wizard.

You are redirected to the vCenter Server Appliance Getting Started page.

6 Perform a post restore recovery to complete the restore process.

Restored Node Type	Action
vCenter Server Appliance with an external Platform Services Controller	<ul style="list-style-type: none"> a Log in to the restored vCenter Server Appliance Bash shell. b Run the script <code>/usr/bin/vcenter-restore</code>.
Platform Services Controller appliance	For all vCenter Server nodes in the domain <ul style="list-style-type: none"> a Log in to the restored vCenter Server Appliance Bash shell. b Run the script <code>/usr/bin/vcenter-restore</code>.
vCenter Server Appliance with an embedded Platform Services Controller	Post restore recovery is not required for this node type.

7 If the backed up vCenter node is part of a vCenter High Availability cluster, the last needs to be reconfigured after the restore operation completes successfully.

For information about how to perform backup and restore operations, see *vSphere Availability*.

Image-Based Backup and Restore of a vCenter Server Environment



You can use vSphere Data Protection or a third-party product that is integrated with VMware vSphere Storage APIs - Data Protection to back up and restore a virtual machine that contains vCenter Server, a vCenter Server Appliance, or a Platform Services Controller.

vSphere Data Protection is a disk-based backup and recovery solution that is powered by EMC. vSphere Data Protection is fully integrated with vCenter Server and lets you manage backup jobs while storing backups in deduplicated destination storage locations. After you deploy and configure vSphere Data Protection, you can access vSphere Data Protection by using the vSphere Web Client interface to select, schedule, configure, and manage backups and recoveries of virtual machines. During a backup, vSphere Data Protection creates a quiesced snapshot of the virtual machine. Deduplication is automatically performed with every backup operation.

You can use vSphere Data Protection to perform a full image backup of a virtual machine that contains vCenter Server, a vCenter Server Appliance, or a Platform Services Controller. The virtual machine must use a fully qualified domain name (FQDN) with correct DNS resolution, or the hostname must be configured to be an IP address. If the hostname is configured as an IP address, the IP address cannot be changed.

The following backups and recoveries are not supported:

- Incremental backups
- Differential backups
- Individual disk backups
- Virtual machines that have snapshots
- Virtual machines configured with Fault Tolerance

As an alternative to vSphere Data Protection, you can also use third-party products that are integrated with VMware vSphere Storage APIs - Data Protection.

VMware vSphere Storage APIs - Data Protection is a data protection framework that enables backup products to perform centralized, efficient, off-host LAN free backup of vSphere virtual machines. For information about VMware vSphere Storage APIs - Data Protection, see the VMware Web site. For information about the integration of backup products with VMware vSphere Storage APIs - Data Protection, contact your backup vendor.

This section includes the following topics:

- [Considerations and Limitations for Image-Based Backup and Restore](#)
- [Use vSphere Data Protection to Back Up a vCenter Server Environment](#)
- [Use vSphere Data Protection to Restore a vCenter Server Environment](#)

Considerations and Limitations for Image-Based Backup and Restore

When you restore a vCenter Server environment, take into account these considerations and limitations.

Note Restoring a vCenter Server or Platform Services Controller instance with DHCP network configuration results in changing its IP address. The changed IP address prevents some vCenter Server services from starting properly. To start all vCenter Server services successfully, after the restore, you must reconfigure the IP address of the restored vCenter Server or Platform Services Controller instance to the IP address that the instance was set to when you performed the backup.

Configuration

After a restore, the following configurations revert to the state when the backup was taken.

- Virtual machine resource settings
- Resource pool hierarchy and setting
- Cluster-host membership
- DRS configuration and rules

Storage DRS

If the configuration changes, the following might change after a restore.

- Datastore Cluster configuration
- Datastore Cluster membership
- Datastore I/O Resource Management (Storage I/O Control) settings
- Datastore-Datacenter membership
- Host-Datastore membership

Distributed Power Management

If you put a host into standby mode after a backup, the vCenter Server might force the host to exit standby mode when you restore to the backup.

Distributed Virtual Switch

If you use a distributed virtual switch, you are advised to export separately the distributed virtual switch configuration before you restore to a backup. You can import the configuration after the restore. If you omit this consideration, you may lose the changes made to a distributed virtual switch after the backup. For detailed steps, see the VMware knowledge base article at <http://kb.vmware.com/kb/2034602>.

Content Libraries

If you delete libraries or items after a backup, you cannot access or use these libraries or items after the restore. You can only delete such libraries or items. A warning message notifies you that there are missing files or folders in the storage backup.

If you create new items or item files after the backup, the Content Library Service has no record of the new items or files after the restore operation. A warning notifies you that extra folders or files were found on the storage backup.

If you create new libraries after the backup, the Content Library Service has no record of the new libraries after restore. The library content exists on the storage backing, but no warning is displayed. You must manually clean the new libraries.

Virtual Machine Life Cycle Operations

- Restoring vCenter Server from a backup that was taken while there are in-flight relocation operations within the vCenter Server instance.

After you restore vCenter Server, the vCenter Server view of the virtual machines may be out of sync with the ESXi view of the virtual machines. This is also true if you performed the backup while there were in-flight operations on vCenter Server. If virtual machines disappear after you restore vCenter Server, you can refer to the following cases.

- a The missing virtual machine is located on the destination ESXi host and is registered with the destination ESXi host, but it is not in the vCenter Server inventory. You must manually add the virtual machine to the vCenter Server inventory.
 - b The missing virtual machine is located on the destination ESXi host, but it is not registered with the destination ESXi host and it is not in the vCenter Server inventory. You must manually register the virtual machine to the ESXi and add the virtual machine back to the vCenter Server inventory.
 - c The missing virtual machine is located on the destination ESXi host, but it is not registered with the destination ESXi host. Within the vCenter Server instance, the missing virtual machine is marked as orphaned. You must remove the virtual machine from the vCenter Server inventory and add it again.
- Restoring vCenter Server from a backup that has an out of date linked clone virtual machine layout.

If you create a linked clone virtual machine after the backup and you restore vCenter Server from the old backup, then after the restore, vCenter Server does not know about the new linked clone virtual machine until vCenter Server discovers the new linked clone virtual machine. If you remove all existing virtual machines before the new linked clone virtual machine is discovered, then the removal of existing virtual machines corrupts the new linked clone due to missing disks. To avoid this corruption, you must wait until all linked clone virtual machines get discovered by the vCenter Server before you remove virtual machines.

vSphere High Availability

Restoring vCenter Server from a backup may cause it to roll back to older version for the vSphere HA cluster state (HostList, ClusterConfiguration, VM protection state) while the hosts in the cluster have the latest version for the cluster state. Ensure that the vSphere HA cluster state stays the same during restore and backup operations. Otherwise, the following potential problems are present.

- If hosts are added or removed to/from the vSphere HA cluster after backup and before vCenter Server restore, virtual machines could potentially fail over to hosts not managed by the vCenter Server but are still part of the HA cluster.
- Protection states for new virtual machines are not updated on the vSphere HA agents on the hosts which are part of the vSphere HA cluster. As a result, virtual machines are not protected/unprotected.
- New cluster configuration state is not updated on the vSphere HA agents on the hosts which are part of the vSphere HA cluster.

vCenter High Availability

Restoring vCenter Server requires vCenter HA to be reconfigured.

Storage Policy Based Management

Restoring vCenter Server from a backup can lead to the following inconsistencies related to storage policies, storage providers, and virtual machines.

- Registered storage providers after backup are lost.
- Unregistered storage providers after backup reappear and might show different provider status.
- Changes, such as create, delete, or update, performed on storage policies after backup are lost.
- Changes, such as create, delete, or update, performed on storage policy components after backup are lost.
- Default policy configuration changes for datastores performed after backup are lost.
- Changes in the storage policy association of the virtual machine and its disks, and in their policy compliance might occur.

Virtual Storage Area Network

Restoring vCenter Server from a backup may cause inconsistencies in the vSAN. For information how to check vSAN health, see *Administering VMware vSAN*.

Patching

Restoring vCenter Server from a backup might result in missing security patches. You must apply them again after the restore is complete. For information on patching the vCenter Server Appliance, see *vSphere Upgrade*.

Use vSphere Data Protection to Back Up a vCenter Server Environment

To perform an image-based backup of a virtual machine that contains vCenter Server, a vCenter Server Appliance, or a Platform Services Controller, you must first deploy and configure vSphere Data Protection and complete basic backup tasks.

The topology of your vCenter Server environment may vary and consist of many vCenter Server and Platform Services Controller instances. You must always perform the backup of all vCenter Server and Platform Services Controller instances simultaneously.

Important You can back up and restore only virtual machines that contain vCenter Server, vCenter Server Appliance, and Platform Services Controller. You cannot back up and restore physical machines that are running vCenter Server by using vSphere Data Protection.

Procedure

1 [Deploy the vSphere Data Protection OVF Template](#)

Deploy vSphere Data Protection to back up and restore virtual machines that contain vCenter Server, a vCenter Server Appliance, or a Platform Services Controller.

2 [Configure vSphere Data Protection](#)

During the initial vSphere Data Protection configuration you can configure the network settings and time zone information for your vSphere Data Protection Appliance. You use the **vSphere Data Protection configuration** wizard to register the vSphere Data Protection Appliance with vCenter Server.

3 [Create a Backup Job in vSphere Data Protection](#)

You can create backup jobs to associate the backup of a set of one or more VMs that contain vCenter Server, the vCenter Server Appliance, and Platform Services Controller with a backup schedule and specific retention policies.

4 [\(Optional\) Start a Backup Job Manually](#)

A backup operation starts automatically according to the scheduled date, time, and frequency configured in the backup job. If you want to run an existing backup job immediately, you can start the process manually.

Deploy the vSphere Data Protection OVF Template

Deploy vSphere Data Protection to back up and restore virtual machines that contain vCenter Server, a vCenter Server Appliance, or a Platform Services Controller.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

- Verify that your ESXi version is 5.1 or later.
- If a firewall is enabled in your environment, verify that port 902 is open for communication between the vSphere Data Protection Appliance and the ESXi host. See the *vSphere Data Protection* documentation.
- Use the vSphere Web Client to log in as an administrator to the vCenter Server instance that manages your environment.
- If your vCenter Server version is earlier than 6.5, verify that the VMware Client Integration Plug-in is installed for your browser. For more information, see the vSphere documentation for your vCenter Server version.

Procedure

- 1 Select **vCenter > Datacenters**.
- 2 On the **Objects** tab, click **Actions** and select **Deploy OVF Template**.
- 3 Navigate to the location of the vSphere Data Protection Appliance .ova file and click **Open**.
- 4 Verify the OVF template details and click **Next**.
- 5 Review the template details, click **Next**, and follow the prompts of the wizard to accept the license agreement.
- 6 On the Select name and folder page, enter an FQDN for the vSphere Data Protection Appliance, select the folder or data center where you want to deploy the vSphere Data Protection Appliance, and click **Next**.

The vSphere Data Protection configuration uses the name that you enter to find the vSphere Data Protection Appliance in the vCenter Server inventory. Do not change the vSphere Data Protection Appliance name after installation.

- 7 Select the host on which to deploy the vSphere Data Protection Appliance and click **Next**.
- 8 Select the virtual disk format and the storage location for the vSphere Data Protection Appliance and click **Next**.
- 9 Select the destination network for the vSphere Data Protection Appliance and click **Next**.

- 10 Enter the network settings such as the default gateway, DNS, network IP address, and netmask and click **Next**.

Confirm that the IP addresses are correct and match the entry in the DNS server. If you enter incorrect IP addresses, you must redeploy the vSphere Data Protection Appliance.

Note The vSphere Data Protection Appliance does not support DHCP. A static IP address is required.

- 11 On the Ready to complete page, confirm that all of the deployment options are correct, select **Power on after deployment**, and click **Finish**.

The vSphere Data Protection Appliance deployment process starts and the vSphere Data Protection Appliance boots in install mode.

Configure vSphere Data Protection

During the initial vSphere Data Protection configuration you can configure the network settings and time zone information for your vSphere Data Protection Appliance. You use the **vSphere Data Protection configuration** wizard to register the vSphere Data Protection Appliance with vCenter Server.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

- Read the *vSphere Data Protection Administration Guide* for the complete list of steps to configure vSphere Data Protection.
- Verify that enough disk space is available on the datastore. When an optional performance analysis test is run during the initial configuration of the appliance, 41 GB is required for each disk on each datastore. If the available space is not enough, the test reports a value of 0 for all of the read, write, and seek tests, and displays a final status of insufficient space.
- Use the vSphere Web Client to log in as an administrator to the vCenter Server instance that manages your environment.

Procedure

- 1 In the vSphere Web Client, select **vCenter Inventory Lists > Virtual Machines**.
- 2 Right-click the vSphere Data Protection Appliance and select **Open Console**.
After the installation files load, the Welcome screen for the vSphere Data Protection menu appears.
- 3 In a Web browser, navigate to vSphere Data Protection Configuration Utility URL.
https://ip_address_VDP_Appliance:8543/vdp-configure/

4 Log in as root.

The default password is changeme.

The **vSphere Data Protection configuration** wizard appears.

5 On the Network settings page of the wizard, enter or confirm the network and server information for the vSphere Data Protection Appliance, and click **Next**.

Ensure that the values are populated correctly, otherwise the initial configuration fails.

6 Select the appropriate time zone for your vSphere Data Protection Appliance and click **Next**.

7 On the VDP credentials page, select a new root password for the virtual appliance and click **Next**.

8 On the vCenter Registration page, register the appliance with vCenter Server:

- a In the **vCenter username** text box, enter a vCenter Server administrator user name. For example, **administrator@vsphere.local**.

If the user belongs to a domain account, enter the user name by using the *DOMAINUserName* format.

Important If you enter the vCenter Single Sign-On administrator user name in the user principal name (UPN) format, the tasks related to vSphere Data Protection operations do not appear in the Recent Tasks pane of the vSphere Web Client. If you want to use the vCenter Single Sign-On administrator user name, enter the vCenter Single Sign-On user name in UPN format.

- b In the **vCenter password** text box, enter the vCenter Server password.

- c Enter a vCenter FQDN or IP address.

- d Change the default vCenter Server HTTP port.

Enter a custom value for the HTTP port if you must connect to vCenter Server over the HTTP port, instead of the HTTPS port, which is used for all other communication.

- e Enter a vCenter HTTPS port (the default is 443).

- f Select the **Use vCenter for SSO authentication** check box.

- g (Optional) Click **Test Connection**.

A connection success message appears. If this message does not appear, troubleshoot your settings and repeat this step until a successful message appears.

9 Click **Next** and respond to the wizard prompts to complete the configuration.

Create a Backup Job in vSphere Data Protection

You can create backup jobs to associate the backup of a set of one or more VMs that contain vCenter Server, the vCenter Server Appliance, and Platform Services Controller with a backup schedule and specific retention policies.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that was used during the vSphere Data Protection configuration.

Procedure

- 1 On the vSphere Web Client Home page, click **vSphere Data Protection**.
- 2 From the **Backup Job Actions** menu, select **New** to run the **Create new backup job** wizard.
- 3 On the Job Type page, select **Guest Images** and click **Next**.
- 4 On the Data Type page, select **Full Image** and click **Next**.

You can see all the objects and virtual machines in the vCenter Server inventory.

- 5 On the Backup Targets page, select the VM that contains the vCenter Server or Platform Services Controller instance you want to back up, and click **Next**.
- 6 On the Schedule page, select the schedule for the backup job and click **Next**.
- 7 On the Retention Policy page, select a retention period and click **Next**.

Note When you enter a new maintenance period that follows the expiration of a backup, the vSphere Data Protection Appliance removes its reference to the backup data and you cannot restore the expired backup. The vSphere Data Protection Appliance determines whether the backup data is used by any other restore point, and if the system determines that the data is not used, the data is removed and the disk capacity becomes available.

- 8 On the Name page, enter a name for the backup job and click **Next**.
- 9 On the Ready to Complete page, review the summary information for the backup job and click **Finish**.

The newly created backup job is listed on the **Backup** tab. The backup job starts automatically according to the configured schedule.

(Optional) Start a Backup Job Manually

A backup operation starts automatically according to the scheduled date, time, and frequency configured in the backup job. If you want to run an existing backup job immediately, you can start the process manually.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that was used during the vSphere Data Protection configuration.

Procedure

- 1 On the vSphere Web Client Home page, click **vSphere Data Protection**.
- 2 On the **Backup** tab, select the backup job that you want to run.
- 3 Click **Backup now**, and select **Backup all sources**.

A dialog box confirms that the backup operation was successfully initiated.

Use vSphere Data Protection to Restore a vCenter Server Environment

You can use vSphere Data Protection or a third-party product that is integrated with VMware vSphere Storage APIs - Data Protection to restore a virtual machine that contains vCenter Server, vCenter Server Appliance, or Platform Services Controller.

You can use vSphere Data Protection to perform an image-based restore of a virtual machine that contains vCenter Server, a vCenter Server Appliance, or a Platform Services Controller. The virtual machine must use a fully qualified domain name (FQDN) with correct DNS resolution, or the host name of the machine must be configured to be an IP address. If the host name is configured as an IP address, the IP address cannot be changed.

You can restore a virtual machine to the original location by either overwriting the backed up virtual machine or by creating a new virtual machine that contains the restored vCenter Server, vCenter Server Appliance, or Platform Services Controller on the same ESXi host. You can also restore the virtual machine on a new ESXi host.

You can restore a virtual machine that contains vCenter Server or a Platform Services Controller instance directly on the ESXi host that is running the vSphere Data Protection Appliance when the vCenter Server service becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

Important Restoring virtual machines that have snapshots or that are configured with Fault Tolerance is unsupported.

Figure 7-1. vCenter Server Restore Workflow

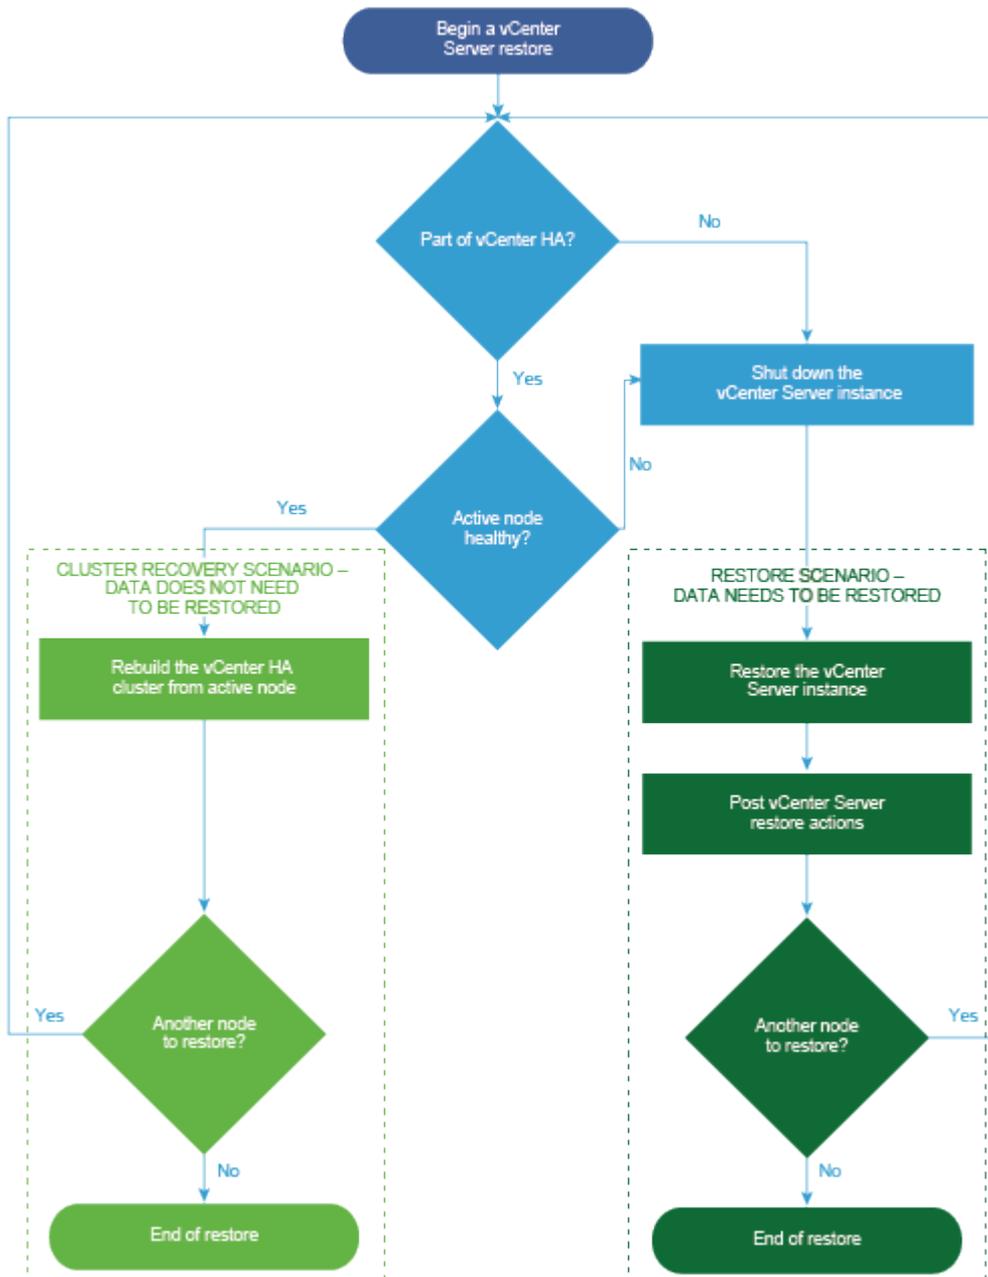
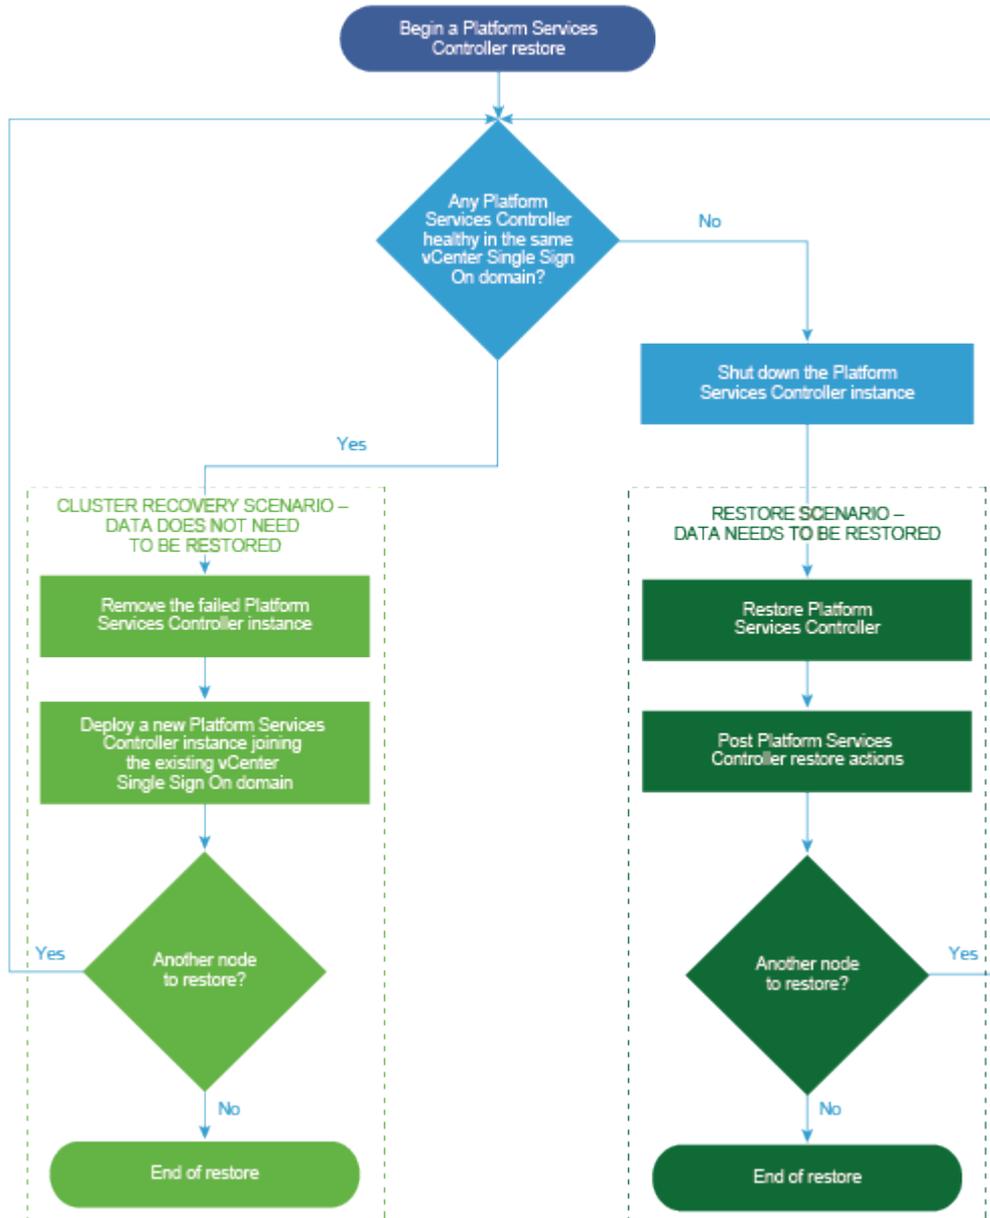


Figure 7-2. Platform Services Controller Restore Workflow



Restore a vCenter Server Instance with an Embedded Platform Services Controller

Your environment might consist of vCenter Server or a vCenter Server Appliance with an embedded Platform Services Controller. You can use vSphere Data Protection to restore a vCenter Server environment with an embedded Platform Services Controller.

Important You can back up and restore only virtual machines that contain vCenter Server, vCenter Server Appliance, and Platform Services Controller. You cannot back up and restore physical machines that are running vCenter Server by using vSphere Data Protection.

- [Restore the Failed vCenter Server Virtual Machine to the Original Location](#)

You can restore to the original location the full image backup of a virtual machine that contains vCenter Server with an embedded Platform Services Controller manually by using the **Restore backup** wizard.

- [Restore the Failed vCenter Server Virtual Machine to a New Location](#)

You can manually restore the full image backup of a virtual machine that contains a vCenter Server with an embedded Platform Services Controller by using the **Restore backup** wizard.

- [Restore the Failed vCenter Server Virtual Machine With the Direct-to-Host Emergency Restore Operation](#)

The direct-to-host emergency restore operation lets you restore the virtual machine that contains vCenter Server with an embedded Platform Services Controller when vCenter Server becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

What to do next

Power on the restored virtual machine that contains the restored vCenter Server instance.

Restore the Failed vCenter Server Virtual Machine to the Original Location

You can restore to the original location the full image backup of a virtual machine that contains vCenter Server with an embedded Platform Services Controller manually by using the **Restore backup** wizard.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

- Deploy and configure the vSphere Data Protection Appliance.
- Back up a virtual machine with running vCenter Server. See [Use vSphere Data Protection to Back Up a vCenter Server Environment](#).
- Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that was used during the vSphere Data Protection configuration.
- Verify that the virtual machine you want to restore is powered off.

Procedure

- 1 On the vSphere Web Client Home page, click **vSphere Data Protection**.
- 2 Click the **Restore** tab.
- 3 (Optional) Filter the backups to narrow your search.

- 4 Select a virtual machine listed in the Name column, and select one or more backup items that you want to restore.

When you select a virtual machine, you can see the list of the performed backups for that virtual machine.

- 5 Click **Restore** to start the **Restore backup** wizard.
- 6 On the Select Backup page, verify that the list of backups is correct, remove the backups that you want to exclude from the restore operation, and click **Next**.
- 7 On the Set Restore Options page, leave the **Restore to Original Location** check box selected.

Important If the virtual disk of the original virtual machine has been removed or deleted, you cannot restore the virtual machine to its original location. The VMDK must be restored to a new location.

- 8 (Optional) Under **Advanced options**, select a new datastore to power on the virtual machine after it is restored and to reconnect the NIC.
- 9 Click **Next**.
- 10 On the Ready to complete page, review the summary of your restore requests, and click **Finish** to start the restore operation.

Note If you selected to reconnect the NIC during the restore process, verify that the network configuration for the newly created virtual machine is correct. The NIC of the new virtual machine might use the same IP address as the original virtual machine, which causes conflicts.

An information dialog box confirms that the restore operation was successfully initiated. You can monitor the restore progress in the Recent Tasks pane.

Restore the Failed vCenter Server Virtual Machine to a New Location

You can manually restore the full image backup of a virtual machine that contains a vCenter Server with an embedded Platform Services Controller by using the **Restore backup** wizard.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

- Deploy and configure the vSphere Data Protection Appliance.
- Back up a virtual machine with running vCenter Server. See [Use vSphere Data Protection to Back Up a vCenter Server Environment](#).
- Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that was used during the vSphere Data Protection configuration.
- Verify that the virtual machine you want to restore is powered off.

Procedure

- 1 On the vSphere Web Client Home page, click **vSphere Data Protection**.
- 2 Click the **Restore** tab.
- 3 (Optional) Filter the backups to narrow your search.
- 4 Select a virtual machine listed in the Name column, and select one or more backup items that you want to restore.

When you select a virtual machine, you can see the list of the performed backups for that virtual machine.

- 5 Click **Restore** to start the **Restore backup** wizard.
- 6 On the Select Backup page, verify that the list of backups is correct, remove the backups that you want to exclude from the restore operation, and click **Next**.
- 7 On the Set Restore Options page, deselect the **Restore to Original Location** check box to set the restore options for each backup that you are restoring to a new location.
- 8 Enter the name of the new virtual machine and click **Choose** to select a new host for the virtual machine.
- 9 Select the datastore in which to restore the virtual machine, and click **Next**.
- 10 (Optional) Under **Advanced options**, select a new datastore to power on the virtual machine after it is restored and to reconnect the NIC.
- 11 Click **Next**.
- 12 On the Ready to complete page, review the summary of your restore requests, and click **Finish** to start the restore operation.

Note If you selected to reconnect the NIC during the restore process, verify that the network configuration for the newly created virtual machine is correct. The NIC of the new virtual machine might use the same IP address as the original virtual machine, which causes conflicts.

An information dialog box confirms that the restore operation was successfully initiated. You can monitor the restore progress in the Recent Tasks pane.

Restore the Failed vCenter Server Virtual Machine With the Direct-to-Host Emergency Restore Operation

The direct-to-host emergency restore operation lets you restore the virtual machine that contains vCenter Server with an embedded Platform Services Controller when vCenter Server becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

vSphere Data Protection depends on vCenter Server for many of the vSphere Data Protection core operations. When vCenter Server becomes unavailable, an emergency restore operation can restore the virtual machine that contains the vCenter Server, vCenter Server Appliance, or Platform Services Controller directly on the ESXi host that is running the vSphere Data Protection Appliance. The **Emergency Restore** tab displays a list of VMs that have been backed up by the vSphere Data Protection Appliance. These VMs that contain vCenter Server or Platform Services Controller instances can be restored as new VMs on the ESXi host where the vSphere Data Protection Appliance is running. For best practices, recommendations, and limitations of the emergency restore operation, see the *vSphere Data Protection* documentation.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

Back up the vCenter Server virtual machine or the vCenter Server Appliance by using vSphere Data Protection.

Procedure

- 1 In a Web browser navigate to `http://host-name/ui` or `http://host-IP-address/ui`.

Here, *host-name* is the name of the ESXi host and *host-IP-address* is the IP of the ESXi host on which the vSphere Data Protection Appliance resides. Log in as an administrator to the VMware Host Client.

- a Right-click **Host** in the VMware Host Client inventory and select **Disconnect from vCenter Server** from the menu.
- b Click **Disconnect from vCenter Server** when prompted to disassociate the host from vCenter Server.

Note If the ESXi host is version 5.1, log in to the vSphere Client instead the VMware Host Client and, on the **Summary** tab, click **Disassociate Host from vCenter Server**.

- 2 In a Web browser, navigate to the vSphere Data Protection Configure Utility.

`https://ip_address_VDP_Appliance:8543/vdp-configure/`.

- 3 On the **Emergency Restore** tab, select the virtual machine that will serve as the restore point, and click **Restore**.

- 4 In the Host Credentials dialog box, enter valid host credentials and click **OK**.
- 5 In the Restore a Backup dialog box, enter a new name.
- 6 Select a datastore as the destination target for the backup, and click **Restore**.

Caution The datastore capacity size is listed. Make sure you select a datastore with enough disk space to accommodate the restore. Insufficient space causes the restore to fail.

The restored virtual machine is listed in the inventory at the vSphere host level. Restoring to a more specific inventory path is unsupported.

Restoring a vCenter Server Environment with a Single External Platform Services Controller

Your environment might consist of many vCenter Server instances that are registered with a single Platform Services Controller. You can use vSphere Data Protection to restore a virtual machine that contains a Platform Services Controller. You can also use vSphere Data Protection to restore either virtual machines that contain vCenter Server instances or vCenter Server Appliance instances that are registered with a single external Platform Services Controller.

Note If vCenter Server and Platform Services Controller instances fail at the same time, you must first restore the Platform Services Controller and then the vCenter Server instances.

- **Restore the Failed Platform Services Controller**

You can install or deploy a Platform Services Controller and register several vCenter Server instances with the same Platform Services Controller. You can use vSphere Data Protection to restore your environment if the external Platform Services Controller fails.

- **Restore Failed vCenter Server Instances**

You can install or deploy a Platform Services Controller and register several vCenter Server instances with the same Platform Services Controller. You can use vSphere Data Protection to restore the whole environment, so that if any of the vCenter Server instances fails, you can restore the failed vCenter Server instance.

Restore the Failed Platform Services Controller

You can install or deploy a Platform Services Controller and register several vCenter Server instances with the same Platform Services Controller. You can use vSphere Data Protection to restore your environment if the external Platform Services Controller fails.

Important You can back up and restore only virtual machines that contain vCenter Server, vCenter Server Appliance, and Platform Services Controller. You cannot back up and restore physical machines that are running vCenter Server by using vSphere Data Protection.

Prerequisites

Back up the virtual machine on which the Platform Services Controller resides.

Procedure

1 Restore the Failed Platform Services Controller Virtual Machine With the Direct-to-Host Emergency Restore Operation

The direct-to-host emergency restore operation lets you restore the virtual machine that contains Platform Services Controller when vCenter Server becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

2 Run the vcenter-restore Script

After you complete the restore process of the Platform Services Controller, you must run the vcenter-restore script on the vCenter Server instances registered with the restored Platform Services Controller.

Restore the Failed Platform Services Controller Virtual Machine With the Direct-to-Host Emergency Restore Operation

The direct-to-host emergency restore operation lets you restore the virtual machine that contains Platform Services Controller when vCenter Server becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

vSphere Data Protection depends on vCenter Server for many of the vSphere Data Protection core operations. When vCenter Server becomes unavailable, an emergency restore operation can restore the virtual machine that contains the vCenter Server, vCenter Server Appliance, or Platform Services Controller directly on the ESXi host that is running the vSphere Data Protection Appliance. The **Emergency Restore** tab displays a list of VMs that have been backed up by the vSphere Data Protection Appliance. These VMs that contain vCenter Server or Platform Services Controller instances can be restored as new VMs on the ESXi host where the vSphere Data Protection Appliance is running. For best practices, recommendations, and limitations of the emergency restore operation, see the *vSphere Data Protection* documentation.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Procedure

1 In a Web browser navigate to `http://host-name/ui` or `http://host-IP-address/ui`.

Here, *host-name* is the name of the ESXi host and *host-IP-address* is the IP of the ESXi host on which the vSphere Data Protection Appliance resides. Log in as an administrator to the VMware Host Client.

- a Right-click **Host** in the VMware Host Client inventory and select **Disconnect from vCenter Server** from the menu.
- b Click **Disconnect from vCenter Server** when prompted to disassociate the host from vCenter Server.

Note If the ESXi host is version 5.1, log in to the vSphere Client instead the VMware Host Client and, on the **Summary** tab, click **Disassociate Host from vCenter Server**.

- 2 In a Web browser, navigate to the vSphere Data Protection Configure Utility.
https://ip_address_VDP_Appliance:8543/vdp-configure/.
- 3 On the **Emergency Restore** tab, select the virtual machine that will serve as the restore point, and click **Restore**.
- 4 In the Host Credentials dialog box, enter valid host credentials and click **OK**.
- 5 In the Restore a Backup dialog box, enter a new name.
- 6 Select a datastore as the destination target for the backup, and click **Restore**.

Caution The datastore capacity size is listed. Make sure you select a datastore with enough disk space to accommodate the restore. Insufficient space causes the restore to fail.

The restored virtual machine is listed in the inventory at the vSphere host level. Restoring to a more specific inventory path is unsupported.

Run the vcenter-restore Script

After you complete the restore process of the Platform Services Controller, you must run the vcenter-restore script on the vCenter Server instances registered with the restored Platform Services Controller.

Procedure

- 1 Log in to the vCenter Server virtual machine.
 - For a vCenter Server Appliance, log in to the appliance shell as root.
 - For a vCenter Server installed on Windows, log in to the virtual machine OS as an administrator.
- 2 Verify that no vCenter Server services are running.
 - For a vCenter Server Appliance, run the `service-control --status --all` command in the appliance shell.
 - For a vCenter Server instance installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

- 3 Run the `vcenter-restore` script to complete the restore operation and start all the vCenter Server services.

Option	Action
For a vCenter Server Appliance	Run the <code>vcenter-restore</code> script in the appliance shell. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre>
For vCenter Server installed on Windows	<ol style="list-style-type: none"> From the Windows command prompt, navigate to the <code>vcenter-restore</code> script. By default, the script is located in <code>C:\Program Files\VMware\vCenter Server\</code>. Run the <code>vcenter-restore</code> script. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre> <p>Note If you do not provide arguments three subsequent times, the script closes after notifying you that the required arguments were not provided.</p>

Here, `psc_administrator_username` is the vCenter Single Sign-On administrator user name which must be in UPN format.

- 4 Verify that all vCenter Server services are running.
- ◆ For a vCenter Server Appliance deployed as an appliance, run the `service-control --status --all` command in the appliance shell.
 - ◆ For a vCenter Server installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

Restore Failed vCenter Server Instances

You can install or deploy a Platform Services Controller and register several vCenter Server instances with the same Platform Services Controller. You can use vSphere Data Protection to restore the whole environment, so that if any of the vCenter Server instances fails, you can restore the failed vCenter Server instance.

Important You can back up and restore only virtual machines that contain vCenter Server, vCenter Server Appliance, and Platform Services Controller. You cannot back up and restore physical machines that are running vCenter Server by using vSphere Data Protection.

You must restore each failed vCenter Server.

Prerequisites

Back up the virtual machines on which the vCenter Server instances reside.

- [Restore the Failed vCenter Server Virtual Machine to the Original Location](#)

You can restore to the original location the full image backup of a virtual machine that contains a vCenter Server instance manually by using the **Restore backup** wizard.

- [Restore the Failed vCenter Server Virtual Machine to a New Location](#)

You can manually restore the full image backup of a virtual machine that contains a vCenter Server instance by using the **Restore backup** wizard.

- [Restore the Failed vCenter Server Virtual Machine With the Direct-to-Host Emergency Restore Operation](#)

The direct-to-host emergency restore operation lets you restore the virtual machine that contains vCenter Server when vCenter Server becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

What to do next

Power on the restored virtual machine that contains the restored vCenter Server instance.

Restore the Failed vCenter Server Virtual Machine to the Original Location

You can restore to the original location the full image backup of a virtual machine that contains a vCenter Server instance manually by using the **Restore backup** wizard.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

- Deploy and configure the vSphere Data Protection Appliance.
- Back up a virtual machine with running vCenter Server. See [Use vSphere Data Protection to Back Up a vCenter Server Environment](#).
- Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that was used during the vSphere Data Protection configuration.
- Verify that the virtual machine you want to restore is powered off.

Procedure

- 1 On the vSphere Web Client Home page, click **vSphere Data Protection**.
- 2 Click the **Restore** tab.
- 3 (Optional) Filter the backups to narrow your search.
- 4 Select a virtual machine listed in the Name column, and select one or more backup items that you want to restore.

When you select a virtual machine, you can see the list of the performed backups for that virtual machine.

- 5 Click **Restore** to start the **Restore backup** wizard.
- 6 On the Select Backup page, verify that the list of backups is correct, remove the backups that you want to exclude from the restore operation, and click **Next**.

- 7 On the Set Restore Options page, leave the **Restore to Original Location** check box selected.

Important If the virtual disk of the original virtual machine has been removed or deleted, you cannot restore the virtual machine to its original location. The VMDK must be restored to a new location.

- 8 (Optional) Under **Advanced options**, select a new datastore to power on the virtual machine after it is restored and to reconnect the NIC.
- 9 Click **Next**.
- 10 On the Ready to complete page, review the summary of your restore requests, and click **Finish** to start the restore operation.

Note If you selected to reconnect the NIC during the restore process, verify that the network configuration for the newly created virtual machine is correct. The NIC of the new virtual machine might use the same IP address as the original virtual machine, which causes conflicts.

- 11 Verify that no vCenter Server services are running.
- For a vCenter Server Appliance, run the `service-control --status --all` command in the appliance shell.
 - For a vCenter Server instance installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.
- 12 Run the `vcenter-restore` script to complete the restore operation and start all the vCenter Server services.

Option	Action
For a vCenter Server Appliance	Run the <code>vcenter-restore</code> script in the appliance shell. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre>
For vCenter Server installed on Windows	<ol style="list-style-type: none"> 1 From the Windows command prompt, navigate to the <code>vcenter-restore</code> script. By default, the script is located in <code>C:\Program Files\VMware\vCenter Server\</code>. 2 Run the <code>vcenter-restore</code> script. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre>

Note If you do not provide arguments three subsequent times, the script closes after notifying you that the required arguments were not provided.

Here, `psc_administrator_username` is the vCenter Single Sign-On administrator user name which must be in UPN format.

- 13 Verify that all vCenter Server services are running.
- ◆ For a vCenter Server Appliance deployed as an appliance, run the `service-control --status --all` command in the appliance shell.
 - ◆ For a vCenter Server installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

An information dialog box confirms that the restore operation was successfully initiated. You can monitor the restore progress in the Recent Tasks pane.

Restore the Failed vCenter Server Virtual Machine to a New Location

You can manually restore the full image backup of a virtual machine that contains a vCenter Server instance by using the **Restore backup** wizard.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

- Deploy and configure the vSphere Data Protection Appliance.
- Back up a virtual machine with running vCenter Server. See [Use vSphere Data Protection to Back Up a vCenter Server Environment](#).
- Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that was used during the vSphere Data Protection configuration.
- Verify that the virtual machine you want to restore is powered off.

Procedure

- 1 On the vSphere Web Client Home page, click **vSphere Data Protection**.
- 2 Click the **Restore** tab.
- 3 (Optional) Filter the backups to narrow your search.
- 4 Select a virtual machine listed in the Name column, and select one or more backup items that you want to restore.

When you select a virtual machine, you can see the list of the performed backups for that virtual machine.

- 5 Click **Restore** to start the **Restore backup** wizard.
- 6 On the Select Backup page, verify that the list of backups is correct, remove the backups that you want to exclude from the restore operation, and click **Next**.
- 7 On the Set Restore Options page, deselect the **Restore to Original Location** check box to set the restore options for each backup that you are restoring to a new location.
- 8 Enter the name of the new virtual machine and click **Choose** to select a new host for the virtual machine.
- 9 Select the datastore in which to restore the virtual machine, and click **Next**.
- 10 (Optional) Under **Advanced options**, select a new datastore to power on the virtual machine after it is restored and to reconnect the NIC.
- 11 Click **Next**.

- 12 On the Ready to complete page, review the summary of your restore requests, and click **Finish** to start the restore operation.

Note If you selected to reconnect the NIC during the restore process, verify that the network configuration for the newly created virtual machine is correct. The NIC of the new virtual machine might use the same IP address as the original virtual machine, which causes conflicts.

- 13 Verify that no vCenter Server services are running.

- For a vCenter Server Appliance, run the `service-control --status --all` command in the appliance shell.
- For a vCenter Server instance installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

- 14 Run the `vcenter-restore` script to complete the restore operation and start all the vCenter Server services.

Option	Action
For a vCenter Server Appliance	Run the <code>vcenter-restore</code> script in the appliance shell. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre>
For vCenter Server installed on Windows	<ol style="list-style-type: none"> 1 From the Windows command prompt, navigate to the <code>vcenter-restore</code> script. By default, the script is located in <code>C:\Program Files\VMware\vCenter Server\</code>. 2 Run the <code>vcenter-restore</code> script. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre>

Note If you do not provide arguments three subsequent times, the script closes after notifying you that the required arguments were not provided.

Here, `psc_administrator_username` is the vCenter Single Sign-On administrator user name which must be in UPN format.

- 15 Verify that all vCenter Server services are running.

- ◆ For a vCenter Server Appliance deployed as an appliance, run the `service-control --status --all` command in the appliance shell.
- ◆ For a vCenter Server installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

An information dialog box confirms that the restore operation was successfully initiated. You can monitor the restore progress in the Recent Tasks pane.

Restore the Failed vCenter Server Virtual Machine With the Direct-to-Host Emergency Restore Operation

The direct-to-host emergency restore operation lets you restore the virtual machine that contains vCenter Server when vCenter Server becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

vSphere Data Protection depends on vCenter Server for many of the vSphere Data Protection core operations. When vCenter Server becomes unavailable, an emergency restore operation can restore the virtual machine that contains the vCenter Server, vCenter Server Appliance, or Platform Services Controller directly on the ESXi host that is running the vSphere Data Protection Appliance. The **Emergency Restore** tab displays a list of VMs that have been backed up by the vSphere Data Protection Appliance. These VMs that contain vCenter Server or Platform Services Controller instances can be restored as new VMs on the ESXi host where the vSphere Data Protection Appliance is running. For best practices, recommendations, and limitations of the emergency restore operation, see the *vSphere Data Protection* documentation.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

Back up the vCenter Server virtual machine or the vCenter Server Appliance by using vSphere Data Protection.

Procedure

- 1 In a Web browser navigate to `http://host-name/ui` or `http://host-IP-address/ui`.
Here, *host-name* is the name of the ESXi host and *host-IP-address* is the IP of the ESXi host on which the vSphere Data Protection Appliance resides. Log in as an administrator to the VMware Host Client.
 - a Right-click **Host** in the VMware Host Client inventory and select **Disconnect from vCenter Server** from the menu.
 - b Click **Disconnect from vCenter Server** when prompted to disassociate the host from vCenter Server.

Note If the ESXi host is version 5.1, log in to the vSphere Client instead the VMware Host Client and, on the **Summary** tab, click **Disassociate Host from vCenter Server**.

- 2 In a Web browser, navigate to the vSphere Data Protection Configure Utility.
`https://ip_address_VDP_Appliance:8543/vdp-configure/`.
- 3 On the **Emergency Restore** tab, select the virtual machine that will serve as the restore point, and click **Restore**.
- 4 In the Host Credentials dialog box, enter valid host credentials and click **OK**.
- 5 In the Restore a Backup dialog box, enter a new name.

- 6 Select a datastore as the destination target for the backup, and click **Restore**.

Caution The datastore capacity size is listed. Make sure you select a datastore with enough disk space to accommodate the restore. Insufficient space causes the restore to fail.

The restored virtual machine is listed in the inventory at the vSphere host level. Restoring to a more specific inventory path is unsupported.

- 7 Verify that no vCenter Server services are running.
- For a vCenter Server Appliance, run the `service-control --status --all` command in the appliance shell.
 - For a vCenter Server instance installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.
- 8 Run the `vcenter-restore` script to complete the restore operation and start all the vCenter Server services.

Option	Action
For a vCenter Server Appliance	Run the <code>vcenter-restore</code> script in the appliance shell. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre>
For vCenter Server installed on Windows	<ol style="list-style-type: none"> 1 From the Windows command prompt, navigate to the <code>vcenter-restore</code> script. By default, the script is located in <code>C:\Program Files\VMware\vCenter Server\</code>. 2 Run the <code>vcenter-restore</code> script. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre> <p>Note If you do not provide arguments three subsequent times, the script closes after notifying you that the required arguments were not provided.</p>

Here, `psc_administrator_username` is the vCenter Single Sign-On administrator user name which must be in UPN format.

- 9 Verify that all vCenter Server services are running.
- ◆ For a vCenter Server Appliance deployed as an appliance, run the `service-control --status --all` command in the appliance shell.
 - ◆ For a vCenter Server installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

Restoring a vCenter Server Environment with Multiple Platform Services Controller Instances

You can use vSphere Data Protection to restore an environment in which the vCenter Server instances are registered with different Platform Services Controller instances, and the infrastructure data is replicated between the Platform Services Controller instances.

Important You can back up and restore only virtual machines that contain vCenter Server, vCenter Server Appliance, and Platform Services Controller. You cannot back up and restore physical machines that are running vCenter Server by using vSphere Data Protection.

- [Restore a Single Failed Platform Services Controller](#)

Your environment might contain multiple vCenter Server instances registered with different external Platform Services Controller instances that replicate their data. You can use vSphere Data Protection to back up and restore the whole environment so that if a Platform Services Controller fails, you can restore the failed Platform Services Controller.

- [Restore All Failed Platform Services Controller Instances](#)

Your environment might contain multiple vCenter Server instances registered with different external Platform Services Controller instances that replicate their data. You can use vSphere Data Protection for backing up and restoring the whole environment. If all Platform Services Controller instances fail, you can restore the environment.

- [Restore a Failed vCenter Server Instance](#)

Your environment might contain multiple vCenter Server instances registered with different external Platform Services Controller instances, while the infrastructure data is replicated between the Platform Services Controller instances. You can use vSphere Data Protection to restore any failed vCenter Server instance.

Restore a Single Failed Platform Services Controller

Your environment might contain multiple vCenter Server instances registered with different external Platform Services Controller instances that replicate their data. You can use vSphere Data Protection to back up and restore the whole environment so that if a Platform Services Controller fails, you can restore the failed Platform Services Controller.

Important You can back up and restore only virtual machines that contain vCenter Server, vCenter Server Appliance, and Platform Services Controller. You cannot back up and restore physical machines that are running vCenter Server by using vSphere Data Protection.

Procedure

- 1 [Repoint vCenter Server to Another External Platform Services Controller](#)

Joining external Platform Services Controller instances in the same vCenter Single Sign-On domain, ensures high availability of your system.

2 Deploy a New Platform Services Controller

If a Platform Services Controller instance fails in an environment that contains multiple Platform Services Controller instances, you must deploy a new Platform Services Controller instance and join it to an active node in the same vCenter Single Sign-On domain, and site.

3 Repoint Back the vCenter Server Instances to the Restored Platform Services Controller

After you restore a failed Platform Services Controller in an environment that contains multiple vCenter Server instances registered with different external Platform Services Controller instances, you must repoint back the vCenter Server instances to the restored Platform Services Controller.

Repoint vCenter Server to Another External Platform Services Controller

Joining external Platform Services Controller instances in the same vCenter Single Sign-On domain, ensures high availability of your system.

If an external Platform Services Controller stops responding or if you want to distribute the load of an external Platform Services Controller, you can repoint the vCenter Server instances to another Platform Services Controller in the same domain and site.

- You can repoint the vCenter Server instance to an existing functional Platform Services Controller instance with free load capacity in the same domain and site.
- You can install or deploy a new Platform Services Controller instance in the same domain and site to which to repoint the vCenter Server instance.

Prerequisites

- If the old Platform Services Controller instance has stopped responding, remove the node and clean up the stale vmdir data by running the `cmsso-util unregister` command. For information about decommissioning a Platform Services Controller instance, see <https://kb.vmware.com/kb/2106736>.
- Verify that the old and the new Platform Services Controller instances are in the same vCenter Single Sign-On domain and site by running the `vdcrepadmin -f showservers` command. For information about using the command, see <https://kb.vmware.com/kb/2127057>.
- If you want to repoint a vCenter Server Appliance that is configured in a vCenter HA cluster, remove the vCenter HA configuration. For information about removing a vCenter HA configuration, see *vSphere Availability*.

Procedure

- 1 Log in to the vCenter Server instance.
 - For a vCenter Server Appliance, log in to the vCenter Server Appliance shell as root.
 - For a vCenter Server instance on Windows, log in as an administrator to the vCenter Server virtual machine or physical server.
- 2 If the vCenter Server instance runs on Windows, in the Windows command prompt, navigate to `C:\Program Files\VMware\vCenter Server\bin`.

- 3 Run the `cmsso-util repoint` command.

```
cmsso-util repoint --repoint-psc psc_fqdn_or_static_ip [--dc-port port_number]
```

where the square brackets [] enclose the command options.

Here, *psc_fqdn_or_static_ip* is the system name used to identify the Platform Services Controller. This system name must be an FQDN or a static IP address.

Note The FQDN value is case-sensitive.

Use the `--dc-port port_number` option if the Platform Services Controller runs on a custom HTTPS port. The default value of the HTTPS port is 443.

- 4 Log in to the vCenter Server instance by using the vSphere Web Client to verify that the vCenter Server instance is running and can be managed.

The vCenter Server instance is registered with the new Platform Services Controller.

What to do next

If you repointed a vCenter Server Appliance that was configured in a vCenter HA cluster, you can reconfigure the vCenter HA cluster. For information about configuring vCenter HA, see *vSphere Availability*.

Deploy a New Platform Services Controller

If a Platform Services Controller instance fails in an environment that contains multiple Platform Services Controller instances, you must deploy a new Platform Services Controller instance and join it to an active node in the same vCenter Single Sign-On domain, and site.

You can deploy the new Platform Services Controller instance by using one of the deployment methods.

- [Deploy a Platform Services Controller Appliance by Using the GUI](#)
- [Deploy a vCenter Server Appliance or Platform Services Controller Appliance by Using the CLI](#)
- [Installing vCenter Server and Platform Services Controller on Windows](#)

After you deploy the new Platform Services Controller instance, you can repoint the vCenter Server instances back to it.

Repoint Back the vCenter Server Instances to the Restored Platform Services Controller

After you restore a failed Platform Services Controller in an environment that contains multiple vCenter Server instances registered with different external Platform Services Controller instances, you must repoint back the vCenter Server instances to the restored Platform Services Controller.

Prerequisites

Verify that the external Platform Services Controller instances are within a single site and replicate the infrastructure data within a single domain.

Procedure

- 1 Log in to the vCenter Server instance.
 - For a vCenter Server Appliance, log in to the vCenter Server Appliance shell as root.
 - For a vCenter Server instance on Windows, log in as an administrator to the vCenter Server virtual machine or physical server.
- 2 If the vCenter Server instance runs on Windows, in the Windows command prompt, navigate to `C:\Program Files\VMware\vCenter Server\bin`.
- 3 Run the `cmsso-util` script.

```
cmsso-util repoint --repoint-psc psc_fqdn_or_static_ip [--dc-port port_number]
```

where the square brackets [] enclose the command options.

Here, *psc_fqdn_or_static_ip* is the system name used to identify the Platform Services Controller. This system name must be an FQDN or a static IP address.

Use the `--dc-port port_number` option if the Platform Services Controller runs on a custom HTTPS port. The default value of the HTTPS port is 443.

- 4 Log in to the vCenter Server instance by using the vSphere Web Client to verify that the vCenter Server instance is running and can be managed.

The vCenter Server instance is registered with the restored Platform Services Controller.

Restore All Failed Platform Services Controller Instances

Your environment might contain multiple vCenter Server instances registered with different external Platform Services Controller instances that replicate their data. You can use vSphere Data Protection for backing up and restoring the whole environment. If all Platform Services Controller instances fail, you can restore the environment.

Important You can back up and restore only virtual machines that contain vCenter Server, vCenter Server Appliance, and Platform Services Controller. You cannot back up and restore physical machines that are running vCenter Server by using vSphere Data Protection.

Restore the virtual machine and the services of the most recently backed up Platform Services Controller instance by using the direct-to-host emergency restore. After the restore is complete, deploy new Platform Services Controller instances and join them to the restored Platform Services Controller instance. After the deployment, you can repoint the vCenter Server instances to the newly deployed Platform Services Controller instances.

Procedure

1 Restore the Most Recently Backed Up Platform Services Controller Virtual Machine With the Direct-to-Host Emergency Restore Operation

The direct-to-host emergency restore operation lets you restore the VM that contains Platform Services Controller when vCenter Server becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

2 Run the vcenter-restore Script

After you complete the restore process of the Platform Services Controller, you must run the vcenter-restore script on the vCenter Server instances registered with the restored Platform Services Controller.

3 Deploy Platform Services Controllers

If more than one Platform Services Controller instances fail in an environment that contains multiple Platform Services Controller instances, you must deploy new Platform Services Controller instances and join them to the active nodes in the same vCenter Single Sign-On domain, and site.

4 Repoint Back the Connections Between vCenter Server and Platform Services Controller Instances

After you restore the failed Platform Services Controller instances in an environment that contains multiple vCenter Server instances registered with different external Platform Services Controller instances, you must repoint back the vCenter Server instances to the restored Platform Services Controller nodes.

Restore the Most Recently Backed Up Platform Services Controller Virtual Machine With the Direct-to-Host Emergency Restore Operation

The direct-to-host emergency restore operation lets you restore the VM that contains Platform Services Controller when vCenter Server becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

vSphere Data Protection depends on vCenter Server for many of the vSphere Data Protection core operations. When vCenter Server becomes unavailable, an emergency restore operation can restore the virtual machine that contains the vCenter Server, vCenter Server Appliance, or Platform Services Controller directly on the ESXi host that is running the vSphere Data Protection Appliance. The **Emergency Restore** tab displays a list of VMs that have been backed up by the vSphere Data Protection Appliance. These VMs that contain vCenter Server or Platform Services Controller instances can be restored as new VMs on the ESXi host where the vSphere Data Protection Appliance is running. For best practices, recommendations, and limitations of the emergency restore operation, see the *vSphere Data Protection* documentation.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Procedure

- 1 In a Web browser navigate to `http://host-name/ui` or `http://host-IP-address/ui`.

Here, *host-name* is the name of the ESXi host and *host-IP-address* is the IP of the ESXi host on which the vSphere Data Protection Appliance resides. Log in as an administrator to the VMware Host Client.

- a Right-click **Host** in the VMware Host Client inventory and select **Disconnect from vCenter Server** from the menu.
- b Click **Disconnect from vCenter Server** when prompted to disassociate the host from vCenter Server.

Note If the ESXi host is version 5.1, log in to the vSphere Client instead the VMware Host Client and, on the **Summary** tab, click **Disassociate Host from vCenter Server**.

- 2 In a Web browser, navigate to the vSphere Data Protection Configure Utility.
`https://ip_address_VDP_Appliance:8543/vdp-configure/`.
- 3 On the **Emergency Restore** tab, select the virtual machine that will serve as the restore point, and click **Restore**.
- 4 In the Host Credentials dialog box, enter valid host credentials and click **OK**.
- 5 In the Restore a Backup dialog box, enter a new name.
- 6 Select a datastore as the destination target for the backup, and click **Restore**.

Caution The datastore capacity size is listed. Make sure you select a datastore with enough disk space to accommodate the restore. Insufficient space causes the restore to fail.

The restored virtual machine is listed in the inventory at the vSphere host level. Restoring to a more specific inventory path is unsupported.

Run the vcenter-restore Script

After you complete the restore process of the Platform Services Controller, you must run the vcenter-restore script on the vCenter Server instances registered with the restored Platform Services Controller.

Procedure

- 1 Log in to the vCenter Server virtual machine.
 - For a vCenter Server Appliance, log in to the appliance shell as root.
 - For a vCenter Server installed on Windows, log in to the virtual machine OS as an administrator.

- 2 Verify that no vCenter Server services are running.
 - For a vCenter Server Appliance, run the `service-control --status --all` command in the appliance shell.
 - For a vCenter Server instance installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.
- 3 Run the `vcenter-restore` script to complete the restore operation and start all the vCenter Server services.

Option	Action
For a vCenter Server Appliance	Run the <code>vcenter-restore</code> script in the appliance shell. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre>
For vCenter Server installed on Windows	<ol style="list-style-type: none"> 1 From the Windows command prompt, navigate to the <code>vcenter-restore</code> script. By default, the script is located in <code>C:\Program Files\VMware\vCenter Server\</code>. 2 Run the <code>vcenter-restore</code> script. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre> <p>Note If you do not provide arguments three subsequent times, the script closes after notifying you that the required arguments were not provided.</p>

Here, `psc_administrator_username` is the vCenter Single Sign-On administrator user name which must be in UPN format.

- 4 Verify that all vCenter Server services are running.
 - ◆ For a vCenter Server Appliance deployed as an appliance, run the `service-control --status --all` command in the appliance shell.
 - ◆ For a vCenter Server installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

Deploy Platform Services Controllers

If more than one Platform Services Controller instances fail in an environment that contains multiple Platform Services Controller instances, you must deploy new Platform Services Controller instances and join them to the active nodes in the same vCenter Single Sign-On domain, and site.

You can deploy the new Platform Services Controller instances by using one of the deployment methods.

- [Deploy a Platform Services Controller Appliance by Using the GUI](#)
- [Deploy a vCenter Server Appliance or Platform Services Controller Appliance by Using the CLI](#)
- [Installing vCenter Server and Platform Services Controller on Windows](#)

After you deploy the new Platform Services Controller instances, you can repoint the vCenter Server instances back to them.

Repoint Back the Connections Between vCenter Server and Platform Services Controller Instances

After you restore the failed Platform Services Controller instances in an environment that contains multiple vCenter Server instances registered with different external Platform Services Controller instances, you must repoint back the vCenter Server instances to the restored Platform Services Controller nodes.

Prerequisites

Verify that the external Platform Services Controller instances are within a single site and replicate the infrastructure data within a single domain.

Procedure

- 1 Log in to the vCenter Server instance.
 - For a vCenter Server Appliance, log in to the vCenter Server Appliance shell as root.
 - For a vCenter Server instance on Windows, log in as an administrator to the vCenter Server virtual machine or physical server.
- 2 If the vCenter Server instance runs on Windows, in the Windows command prompt, navigate to `C:\Program Files\VMware\vCenter Server\bin`.
- 3 Run the `cmsso-util` script.

```
cmsso-util repoint --repoint-psc psc_fqdn_or_static_ip [--dc-port port_number]
```

where the square brackets [] enclose the command options.

Here, *psc_fqdn_or_static_ip* is the system name used to identify the Platform Services Controller. This system name must be an FQDN or a static IP address.

Use the `--dc-port port_number` option if the Platform Services Controller runs on a custom HTTPS port. The default value of the HTTPS port is 443.

- 4 Log in to the vCenter Server instance by using the vSphere Web Client to verify that the vCenter Server instance is running and can be managed.

The vCenter Server instance is registered with the restored Platform Services Controller.

Restore a Failed vCenter Server Instance

Your environment might contain multiple vCenter Server instances registered with different external Platform Services Controller instances, while the infrastructure data is replicated between the Platform Services Controller instances. You can use vSphere Data Protection to restore any failed vCenter Server instance.

Important You can back up and restore only virtual machines that contain vCenter Server, vCenter Server Appliance, and Platform Services Controller. You cannot back up and restore physical machines that are running vCenter Server by using vSphere Data Protection.

You must restore each failed vCenter Server.

Prerequisites

Back up the virtual machines on which the vCenter Server instances reside.

- [Restore the Failed vCenter Server Virtual Machine to the Original Location](#)
You can restore to the original location the full image backup of a virtual machine that contains a vCenter Server instance manually by using the **Restore backup** wizard.
- [Restore the Failed vCenter Server Virtual Machine to a New Location](#)
You can manually restore the full image backup of a virtual machine that contains a vCenter Server instance by using the **Restore backup** wizard.
- [Restore the Failed vCenter Server Virtual Machine With the Direct-to-Host Emergency Restore Operation](#)
The direct-to-host emergency restore operation lets you restore the virtual machine that contains vCenter Server when vCenter Server becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

Restore the Failed vCenter Server Virtual Machine to the Original Location

You can restore to the original location the full image backup of a virtual machine that contains a vCenter Server instance manually by using the **Restore backup** wizard.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

- Deploy and configure the vSphere Data Protection Appliance.
- Back up a virtual machine with running vCenter Server. See [Use vSphere Data Protection to Back Up a vCenter Server Environment](#).
- Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that was used during the vSphere Data Protection configuration.

- Verify that the virtual machine you want to restore is powered off.

Procedure

- 1 On the vSphere Web Client Home page, click **vSphere Data Protection**.
- 2 Click the **Restore** tab.
- 3 (Optional) Filter the backups to narrow your search.
- 4 Select a virtual machine listed in the Name column, and select one or more backup items that you want to restore.

When you select a virtual machine, you can see the list of the performed backups for that virtual machine.

- 5 Click **Restore** to start the **Restore backup** wizard.
- 6 On the Select Backup page, verify that the list of backups is correct, remove the backups that you want to exclude from the restore operation, and click **Next**.
- 7 On the Set Restore Options page, leave the **Restore to Original Location** check box selected.

Important If the virtual disk of the original virtual machine has been removed or deleted, you cannot restore the virtual machine to its original location. The VMDK must be restored to a new location.

- 8 (Optional) Under **Advanced options**, select a new datastore to power on the virtual machine after it is restored and to reconnect the NIC.
- 9 Click **Next**.
- 10 On the Ready to complete page, review the summary of your restore requests, and click **Finish** to start the restore operation.

Note If you selected to reconnect the NIC during the restore process, verify that the network configuration for the newly created virtual machine is correct. The NIC of the new virtual machine might use the same IP address as the original virtual machine, which causes conflicts.

- 11 Verify that no vCenter Server services are running.
 - For a vCenter Server Appliance, run the `service-control --status --all` command in the appliance shell.
 - For a vCenter Server instance installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

- 12 Run the `vcenter-restore` script to complete the restore operation and start all the vCenter Server services.

Option	Action
For a vCenter Server Appliance	Run the <code>vcenter-restore</code> script in the appliance shell. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre>
For vCenter Server installed on Windows	<ol style="list-style-type: none"> From the Windows command prompt, navigate to the <code>vcenter-restore</code> script. By default, the script is located in <code>C:\Program Files\VMware\vCenter Server\</code>. Run the <code>vcenter-restore</code> script. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre> <p>Note If you do not provide arguments three subsequent times, the script closes after notifying you that the required arguments were not provided.</p>

Here, `psc_administrator_username` is the vCenter Single Sign-On administrator user name which must be in UPN format.

- 13 Verify that all vCenter Server services are running.

- ◆ For a vCenter Server Appliance deployed as an appliance, run the `service-control --status --all` command in the appliance shell.
- ◆ For a vCenter Server installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

An information dialog box confirms that the restore operation was successfully initiated. You can monitor the restore progress in the Recent Tasks pane.

Restore the Failed vCenter Server Virtual Machine to a New Location

You can manually restore the full image backup of a virtual machine that contains a vCenter Server instance by using the **Restore backup** wizard.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

- Deploy and configure the vSphere Data Protection Appliance.
- Back up a virtual machine with running vCenter Server. See [Use vSphere Data Protection to Back Up a vCenter Server Environment](#).
- Use the vSphere Web Client to log in to the vCenter Server instance that manages your environment. Log in as the user with administrator privileges that was used during the vSphere Data Protection configuration.
- Verify that the virtual machine you want to restore is powered off.

Procedure

- 1 On the vSphere Web Client Home page, click **vSphere Data Protection**.
- 2 Click the **Restore** tab.
- 3 (Optional) Filter the backups to narrow your search.
- 4 Select a virtual machine listed in the Name column, and select one or more backup items that you want to restore.

When you select a virtual machine, you can see the list of the performed backups for that virtual machine.

- 5 Click **Restore** to start the **Restore backup** wizard.
- 6 On the Select Backup page, verify that the list of backups is correct, remove the backups that you want to exclude from the restore operation, and click **Next**.
- 7 On the Set Restore Options page, deselect the **Restore to Original Location** check box to set the restore options for each backup that you are restoring to a new location.
- 8 Enter the name of the new virtual machine and click **Choose** to select a new host for the virtual machine.
- 9 Select the datastore in which to restore the virtual machine, and click **Next**.
- 10 (Optional) Under **Advanced options**, select a new datastore to power on the virtual machine after it is restored and to reconnect the NIC.
- 11 Click **Next**.
- 12 On the Ready to complete page, review the summary of your restore requests, and click **Finish** to start the restore operation.

Note If you selected to reconnect the NIC during the restore process, verify that the network configuration for the newly created virtual machine is correct. The NIC of the new virtual machine might use the same IP address as the original virtual machine, which causes conflicts.

- 13 Verify that no vCenter Server services are running.
 - For a vCenter Server Appliance, run the `service-control --status --all` command in the appliance shell.
 - For a vCenter Server instance installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

- 14 Run the `vcenter-restore` script to complete the restore operation and start all the vCenter Server services.

Option	Action
For a vCenter Server Appliance	Run the <code>vcenter-restore</code> script in the appliance shell. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre>
For vCenter Server installed on Windows	<ol style="list-style-type: none"> From the Windows command prompt, navigate to the <code>vcenter-restore</code> script. By default, the script is located in <code>C:\Program Files\VMware\vCenter Server\</code>. Run the <code>vcenter-restore</code> script. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre> <p>Note If you do not provide arguments three subsequent times, the script closes after notifying you that the required arguments were not provided.</p>

Here, `psc_administrator_username` is the vCenter Single Sign-On administrator user name which must be in UPN format.

- 15 Verify that all vCenter Server services are running.

- ◆ For a vCenter Server Appliance deployed as an appliance, run the `service-control --status --all` command in the appliance shell.
- ◆ For a vCenter Server installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

An information dialog box confirms that the restore operation was successfully initiated. You can monitor the restore progress in the Recent Tasks pane.

Restore the Failed vCenter Server Virtual Machine With the Direct-to-Host Emergency Restore Operation

The direct-to-host emergency restore operation lets you restore the virtual machine that contains vCenter Server when vCenter Server becomes unavailable or when you cannot access the vSphere Data Protection user interface by using the vSphere Web Client.

vSphere Data Protection depends on vCenter Server for many of the vSphere Data Protection core operations. When vCenter Server becomes unavailable, an emergency restore operation can restore the virtual machine that contains the vCenter Server, vCenter Server Appliance, or Platform Services Controller directly on the ESXi host that is running the vSphere Data Protection Appliance. The **Emergency Restore** tab displays a list of VMs that have been backed up by the vSphere Data Protection Appliance. These VMs that contain vCenter Server or Platform Services Controller instances can be restored as new VMs on the ESXi host where the vSphere Data Protection Appliance is running. For best practices, recommendations, and limitations of the emergency restore operation, see the *vSphere Data Protection* documentation.

Note This procedure describes the steps by using vSphere Data Protection 6.1.3. The steps might vary if you use a different version of vSphere Data Protection.

Prerequisites

Back up the vCenter Server virtual machine or the vCenter Server Appliance by using vSphere Data Protection.

Procedure

- 1 In a Web browser navigate to `http://host-name/ui` or `http://host-IP-address/ui`.

Here, *host-name* is the name of the ESXi host and *host-IP-address* is the IP of the ESXi host on which the vSphere Data Protection Appliance resides. Log in as an administrator to the VMware Host Client.

- a Right-click **Host** in the VMware Host Client inventory and select **Disconnect from vCenter Server** from the menu.
- b Click **Disconnect from vCenter Server** when prompted to disassociate the host from vCenter Server.

Note If the ESXi host is version 5.1, log in to the vSphere Client instead of the VMware Host Client and, on the **Summary** tab, click **Disassociate Host from vCenter Server**.

- 2 In a Web browser, navigate to the vSphere Data Protection Configure Utility.
`https://ip_address_VDP_Appliance:8543/vdp-configure/`.
- 3 On the **Emergency Restore** tab, select the virtual machine that will serve as the restore point, and click **Restore**.
- 4 In the Host Credentials dialog box, enter valid host credentials and click **OK**.
- 5 In the Restore a Backup dialog box, enter a new name.
- 6 Select a datastore as the destination target for the backup, and click **Restore**.

Caution The datastore capacity size is listed. Make sure you select a datastore with enough disk space to accommodate the restore. Insufficient space causes the restore to fail.

The restored virtual machine is listed in the inventory at the vSphere host level. Restoring to a more specific inventory path is unsupported.

- 7 Verify that no vCenter Server services are running.
 - For a vCenter Server Appliance, run the `service-control --status --all` command in the appliance shell.
 - For a vCenter Server instance installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

- 8 Run the `vcenter-restore` script to complete the restore operation and start all the vCenter Server services.

Option	Action
For a vCenter Server Appliance	Run the <code>vcenter-restore</code> script in the appliance shell. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre>
For vCenter Server installed on Windows	<ol style="list-style-type: none"> From the Windows command prompt, navigate to the <code>vcenter-restore</code> script. By default, the script is located in <code>C:\Program Files\VMware\vCenter Server\</code>. Run the <code>vcenter-restore</code> script. <pre>vcenter-restore -u <i>psc_administrator_username</i> -p <i>psc_administrator_password</i></pre> <p>Note If you do not provide arguments three subsequent times, the script closes after notifying you that the required arguments were not provided.</p>

Here, `psc_administrator_username` is the vCenter Single Sign-On administrator user name which must be in UPN format.

- 9 Verify that all vCenter Server services are running.
- ◆ For a vCenter Server Appliance deployed as an appliance, run the `service-control --status --all` command in the appliance shell.
 - ◆ For a vCenter Server installed on Windows, from the Windows **Start** menu, select **Control Panel > Administrative Tools > Services**.

Troubleshooting ESXi Booting

The ESXi booting troubleshooting topics provide solutions to problems that you might encounter during the ESXi booting.

This section includes the following topics:

- [Host Stops Unexpectedly at Bootup When Sharing a Boot Disk with Another Host](#)
- [Host Fails to Boot After You Install ESXi in UEFI Mode](#)

Host Stops Unexpectedly at Bootup When Sharing a Boot Disk with Another Host

When more than one host, either physical or virtual, boots from the same shared physical disk or LUN, they cannot use the same scratch partition.

Problem

The host stops at bootup when sharing a boot disk with another host.

Cause

More than one ESXi host can share the same physical disk or LUN. When two such hosts also have the same scratch partition configured, either of the hosts can fail at bootup.

Solution

- 1 Set the hosts to boot sequentially, and boot the hosts.

This setting lets you start the hosts so that you can change the scratch partition for one of them.

- 2 From the vSphere Web Client, connect to the vCenter Server.
- 3 Select the host in the inventory.
- 4 Click the **Manage** tab.
- 5 Click **Settings**.
- 6 Under System, select **Advanced System Settings**.
- 7 Select **ScratchConfig**.

The **ScratchConfig.CurrentScratchLocation** text box shows the current location of the scratch partition.

- 8 In the **ScratchConfig.ConfiguredScratchLocation** text box, enter a directory path that is unique for this host.

For example, */vmfs/volumes/DatastoreUUID/DatastoreFolder*.

- 9 Reboot the host for the changes to take effect.

Host Fails to Boot After You Install ESXi in UEFI Mode

When you install ESXi on a host machine in UEFI mode, the machine might fail to boot.

Problem

When you reboot after installing ESXi on a host machine in UEFI mode, the reboot might fail. This problem is accompanied by an error message similar to Unexpected network error. No boot device available.

Cause

The host system fails to recognize the disk that ESXi is installed on as the boot disk.

Solution

- 1 While the error message is displayed on screen, press F11 to display boot options.
- 2 Select an option similar to **Add boot option**.
The wording of the option might vary, depending on your system.
- 3 Select the file `\EFI\BOOT\BOOTx64.EFI` on the disk that you installed ESXi on.
- 4 Change the boot order so that the host boots from the option that you added.

Troubleshooting vCenter Server Installation or Deployment

9

The vCenter Server installation or deployment troubleshooting topics provide solutions to problems that you might encounter during the vCenter Server installation or vCenter Server Appliance deployment process.

This section includes the following topics:

- [Collecting Logs for Troubleshooting a vCenter Server Installation or Upgrade](#)
- [Attempt to Install a Platform Services Controller After a Prior Installation Failure](#)
- [Microsoft SQL Database Set to Unsupported Compatibility Mode Causes vCenter Server Installation or Upgrade to Fail](#)

Collecting Logs for Troubleshooting a vCenter Server Installation or Upgrade

You can collect installation or upgrade log files for vCenter Server. If an installation or upgrade fails, checking the log files can help you identify the source of the failure.

You can choose the Installation Wizard method or the manual method for saving and recovering log files for a vCenter Server for Windows installation failure.

You can also collect deployment log files for vCenter Server Appliance.

- [Collect Installation Logs by Using the Installation Wizard](#)

You can use the Setup Interrupted page of the installation wizard to browse to the generated .zip file of the vCenter Server for Windows installation log files.
- [Retrieve Installation Logs Manually](#)

You can retrieve the installation log files manually for examination.
- [Collect Deployment Log Files for the vCenter Server Appliance](#)

If the vCenter Server Appliance deployment fails, you can retrieve the log files and examine them for the reason of the failure.
- [Export a vCenter Server Support Bundle for Troubleshooting](#)

You can export the support bundle of the vCenter Server instance in the vCenter Server Appliance for troubleshooting using the URL displayed on the DCUI home screen.

Collect Installation Logs by Using the Installation Wizard

You can use the Setup Interrupted page of the installation wizard to browse to the generated .zip file of the vCenter Server for Windows installation log files.

If the installation fails, the Setup Interrupted page appears with the log collection check boxes selected by default.

Procedure

- 1 Leave the check boxes selected and click **Finish**.

The installation files are collected in a .zip file on your desktop, for example, `VMware-VCS-logs-time-of-installation-attempt.zip`, where *time-of-installation-attempt* displays the year, month, date, hour, minutes, and seconds of the installation attempt.

- 2 Retrieve the log files from the .zip file on your desktop.

What to do next

Examine the log files to determine the cause of failure.

Retrieve Installation Logs Manually

You can retrieve the installation log files manually for examination.

Procedure

- 1 Navigate to the installation log file locations.

- `%PROGRAMDATA%\VMware\vCenterServer\logs` directory, usually `C:\ProgramData\VMware\vCenterServer\logs`
- `%TEMP%` directory, usually `C:\Users\username\AppData\Local\Temp`

The files in the `%TEMP%` directory include `vc-install.txt`, `vminst.log`, `pkgmgr.log`, `pkgmgr-comp-msi.log`, and `vim-vcs-msi.log`.

- 2 Open the installation log files in a text editor for examination.

Collect Deployment Log Files for the vCenter Server Appliance

If the vCenter Server Appliance deployment fails, you can retrieve the log files and examine them for the reason of the failure.

The full path to the log files is displayed in the vCenter Server Appliance deployment wizard.

In case of firstboot failure, you can download the support bundle on a Windows host machine and examine the log files to determine which firstboot script failed. See [Export a vCenter Server Support Bundle for Troubleshooting](#).

Procedure

- 1 On the Windows machine that you use for deploying the vCenter Server Appliance, navigate to the log files folder.

If you are logged in as an administrator, by default this is the
 C:\Users\Administrator\AppData\Local\VMware\CIP\vcSaInstaller folder.

- 2 Open the installation log files in a text editor for examination.

Export a vCenter Server Support Bundle for Troubleshooting

You can export the support bundle of the vCenter Server instance in the vCenter Server Appliance for troubleshooting using the URL displayed on the DCUI home screen.

You can also collect the support bundle from the vCenter Server Appliance Bash shell by running the `vc-support.sh` script.

The support bundle is exported in `.tgz` format.

Procedure

- 1 Log in to the Windows host machine on which you want to download the bundle.
- 2 Open a Web browser and enter the URL to the support bundle displayed in the DCUI.

`https://appliance-fully-qualified-domain-name:443/appliance/support-bundle`

- 3 Enter the user name and password of the root user.
- 4 Click **Enter**.

The support bundle is downloaded as `.tgz` file on your Windows machine.

- 5 (Optional) To determine which firstboot script failed, examine the `firstbootStatus.json` file.

If you ran the `vc-support.sh` script in the vCenter Server Appliance Bash shell, to examine the `firstbootStatus.json` file, run

```
cat /var/log/firstboot/firstbootStatus.json
```

Attempt to Install a Platform Services Controller After a Prior Installation Failure

When you want to replicate Platform Services Controller data, you might not be able to join a vCenter Single Sign-On domain in an existing Platform Services Controller.

Problem

When you try to install a Platform Services Controller, either embedded or external, and join the Platform Services Controller to a vCenter Single Sign-On domain or site, the installation might fail and the failure might leave incomplete data in the Platform Services Controller federation.

Cause

The Platform Services Controller data is not cleaned up when an installation of a Platform Services Controller fails. Consider the following scenario:

- 1 Install Platform Services Controller A.
- 2 When you try to install Platform Services Controller B and join it to the same domain as Platform Services Controller A, the installation fails.
- 3 Second attempt to install Platform Services Controller B and join it to the same domain as Platform Services Controller A fails, because Platform Services Controller A contains incomplete data.

Solution

- 1 Log in as an administrator to the machine on which you install Platform Services Controller A.
- 2 At the command prompt navigate to the `vdcleavefed` command.

The `vdcleavefed` command is located at `C:\Program Files\VMware\vCenter Server\vmmdir\` on Windows and `/usr/lib/vmware-vmmdir/bin/` on Linux.

- 3 Run the `vdcleavefed` command to delete the data.

```
vdcleavefed -h Platform-Services-Controller-B-System-Name -u Administrator
```

- 4 Install Platform Services Controller B.

Microsoft SQL Database Set to Unsupported Compatibility Mode Causes vCenter Server Installation or Upgrade to Fail

vCenter Server installation with a Microsoft SQL database fails when the database is set to compatibility mode with an unsupported version.

Problem

The following error message appears: The DB User entered does not have the required permissions needed to install and configure vCenter Server with the selected DB. Please correct the following error(s): %s

Cause

The database version must be supported for vCenter Server. For SQL, even if the database is a supported version, if it is set to run in compatibility mode with an unsupported version, this error occurs. For example, if SQL 2008 is set to run in SQL 2000 compatibility mode, this error occurs.

Solution

- ◆ Make sure the vCenter Server database is a supported version and is not set to compatibility mode with an unsupported version. See the VMware Product Interoperability Matrixes at http://partnerweb.vmware.com/comp_guide2/sim/interop_matrix.php?

Decommissioning ESXi and vCenter Server

10

The decommissioning topics provide information on how to remove ESXi and vCenter Server from your host machines.

This section includes the following topics:

- [Decommission an ESXi Host](#)
- [Uninstall vCenter Server](#)

Decommission an ESXi Host

If you do not want your server to be an ESXi host, you can decommission the ESXi host machine.

Procedure

- 1 Remove VMFS datastores on the internal disks so that the internal disks are no longer set up to store virtual machines.
- 2 Change the boot setting in the BIOS so that the host no longer boots into ESXi.
- 3 Install another operating system in its place.

Uninstall vCenter Server

You must have administrator privileges to uninstall VMware vCenter Server.

Important If you are using the embedded PostgreSQL database, uninstalling vCenter Server causes the embedded database to be uninstalled, and all data is lost.

Prerequisites

If you are uninstalling the vCenter Server system, remove the hosts from the Hosts and Clusters inventory.

Procedure

- 1 As an administrator user on the Windows system, click **Start > Control Panel > Programs and Features**.
- 2 Select **VMware vCenter Server** from the list and click **Remove**.

- 3 Click **Remove** to confirm that you want to remove the program.
- 4 Click **Finish**.
- 5 Reboot the system.