

vSphere Installation and Setup

vSphere 5.5

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see <http://www.vmware.com/support/pubs>.

EN-001266-04

vmware[®]

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

<http://www.vmware.com/support/>

The VMware Web site also provides the latest product updates.

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

docfeedback@vmware.com

Copyright © 2009–2013 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

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

Contents

About vSphere Installation and Setup	7
Updated Information	9
1 Introduction to vSphere Installation and Setup	11
Overview of the vSphere Installation and Setup Process	11
2 System Requirements	13
ESXi Hardware Requirements	13
Hardware Requirements for vCenter Server, the vSphere Web Client , vCenter Inventory Service, and vCenter Single Sign-On	17
vCenter Server Software Requirements	22
vSphere Web Client Software Requirements	22
Providing Sufficient Space for System Logging	23
Required Ports for vCenter Server	23
Required Ports for the vCenter Server Appliance	26
Conflict Between vCenter Server and IIS for Port 80	27
DNS Requirements for vSphere	27
Supported Remote Management Server Models and Minimum Firmware Versions	28
3 Before You Install vCenter Server	29
Preparing vCenter Server Databases	30
Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server	49
How vCenter Single Sign-On Affects vCenter Server Installation	51
vCenter Single Sign-On Deployment Modes	52
vCenter Single Sign-On and High Availability	54
vCenter Single Sign-On Components	56
Setting the vCenter Server Administrator User	56
Authenticating to the vCenter Server Environment	57
How vCenter Single Sign-On Affects Log In Behavior	57
Identity Sources for vCenter Server with vCenter Single Sign-On	58
Synchronizing Clocks on the vSphere Network	59
Using a User Account for Running vCenter Server	60
Installing vCenter Server on IPv6 Machines	60
JDBC URL Formats for the vCenter Server Database	61
Running the vCenter Server Installer from a Network Drive	62
Required Information for Installing or Upgrading vCenter Single Sign-On, Inventory Service, vCenter Server, and the vSphere Web Client	62
Microsoft SQL Database Set to Unsupported Compatibility Mode Causes vCenter Server Installation or Upgrade to Fail	68

- 4 Installing vCenter Server 69**
 - vCenter Server Installation and Sign-In Process 69
 - vCenter Server Components and Support Tools 72
 - Download the vCenter Server Installer 73
 - Install vCenter Single Sign-On, the vSphere Web Client , vCenter Inventory Service, and vCenter Server by Using Simple Install 73
 - Use Custom Install to Install vCenter Server and Required Components 76
 - Add a vCenter Single Sign-On Identity Source 85
 - Install or Upgrade vCenter Server Java Components Separately 90
 - Install or Upgrade vCenter Server tc Server Separately 91
 - vCenter Single Sign-On Installation Fails 91
 - Download and Deploy the VMware vCenter Server Appliance 92
- 5 After You Install vCenter Server 99**
 - Install vCenter Server Components 100
 - Creating vCenter Server Linked Mode Groups 109
 - Configuring VMware vCenter Server - tc Server Settings in vCenter Server 114
 - VMware vCenter Management Webservices Service Fails to Start 116
 - Back Up the Inventory Service Database on Windows 116
 - Restore an Inventory Service Database Backup on Windows 116
 - Back Up the Inventory Service Database on Linux 117
 - Restore an Inventory Service Database Backup on Linux 117
 - Reset the vCenter Inventory Service Database 118
 - Enable IPv6 Support for vCenter Inventory Service 119
- 6 Before You Install ESXi 121**
 - Options for Installing ESXi 121
 - Media Options for Booting the ESXi Installer 123
 - Using Remote Management Applications 134
 - Required Information for ESXi Installation 134
 - Download the ESXi Installer 135
- 7 Installing ESXi 137**
 - Installing ESXi Interactively 137
 - Installing, Upgrading, or Migrating Hosts Using a Script 140
 - Installing ESXi Using vSphere Auto Deploy 155
 - Using vSphere ESXi Image Builder CLI 223
- 8 Setting Up ESXi 245**
 - ESXi Autoconfiguration 246
 - About the Direct Console ESXi Interface 246
 - Set the Password for the Administrator Account 249
 - Configuring the BIOS Boot Settings 249
 - Host Fails to Boot After You Install ESXi in UEFI Mode 250
 - Network Access to Your ESXi Host 251
 - Configure the Network Settings on a Host That Is Not Attached to the Network 251
 - Managing ESXi Remotely 252
 - Configuring Network Settings 252

Storage Behavior	256
View System Logs	259
Configure Syslog on ESXi Hosts	260
Enable Lockdown Mode Using the Direct Console	261
Enable Lockdown Mode Using the vSphere Web Client	261
Enable ESXi Shell and SSH Access with the Direct Console User Interface	262
Set the Host Image Profile Acceptance Level	262
Reset the System Configuration	263
Remove All Custom Packages on ESXi	264
Disable Support for Non-ASCII Characters in Virtual Machine File and Directory Names	264
Disable ESXi	264
9 After You Install and Set Up ESXi	265
Managing the ESXi Host with the vSphere Web Client	265
Licensing ESXi Hosts	265
Index	267

About vSphere Installation and Setup

vSphere Installation and Setup describes how to install new configurations of VMware® vCenter Server and ESXi.

Intended Audience

vSphere Installation and Setup is intended for anyone who needs to install and set up vCenter Server and ESXi.

This information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacenter operations.

Updated Information

This *vSphere Installation and Setup* publication is updated with each release of the product or when necessary. This table provides the update history of *vSphere Installation and Setup*.

Revision	Description
EN-001266-04	<ul style="list-style-type: none">■ Removed the topic that discusses how to configure Windows NTP Client for Network Clock Synchronization. That information is available from the Microsoft Web site■ Updated “Synchronizing Clocks on the vSphere Network,” on page 59 to include link to Knowledge Base article that includes current time synchronization intervals.■ In the topic “vCenter Server Database Configuration Notes,” on page 30, added information about the required SQL Server Collation Model and ODBC System DSN minimum version for Microsoft SQL Server 2008 R2 Express and Microsoft SQL Server 2008.■ In the topics “Back Up the Inventory Service Database on Linux,” on page 117 and “Restore an Inventory Service Database Backup on Linux,” on page 117, corrected a typo in the command to stop VMware Inventory Service.
EN-001266-03	<ul style="list-style-type: none">■ Added information about additional vCenter Single Sign-On ports to Table 2-12.■ In the topic “Download and Deploy the VMware vCenter Server Appliance,” on page 92, added information in Step 1 and Step 2.■ In the topic “(Optional) Use a Script to Create the Oracle Database Schema,” on page 44, corrected the list of scripts in Step 3.■ In the topic “vCenter Server Software Requirements,” on page 22, added information about using the online VMware Compatibility Guide to find compatible host operating system versions for vCenter Server.■ In the topic “(Optional) Use a Script to Create Microsoft SQL Server Database Objects Manually,” on page 37, made the following corrections:<ul style="list-style-type: none">■ In Step 4, added script <code>insert_stats_proc_mssql.sql</code>, and corrected the name of script <code>VCDB_views_mssql.sql</code>.■ In Step 7, added <code>grant execute on insert_stats_proc to vCenter_db_user</code> and removed <code>grant execute on process temptable0_proc to vCenter_db_user</code> <code>grant execute on process temptable1_proc to vCenter_db_user</code> <code>grant execute on process temptable2_proc to vCenter_db_user</code>
EN-001266-02	<p>Corrected location of <code>dbschema</code> directory in the following topics:</p> <ul style="list-style-type: none">■ “Configure an Oracle Database User,” on page 43.■ “Use a Script to Create a Local or Remote Oracle Database,” on page 44.■ “(Optional) Use a Script to Create the Oracle Database Schema,” on page 44. <p>.</p>
EN-001266-01	Updated Table 2-8 .
EN-001266-00	Initial release.

Introduction to vSphere Installation and Setup

1

vSphere 5.x provides various options for installation and setup. To ensure a successful vSphere deployment, understand the installation and setup options, and the sequence of tasks required.


You have several options for installing and setting up ESXi, for managing vSphere with vCenter Server and the vSphere Web Client, and for the database setup that you use with vCenter Server.

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.

NOTE This overview does not apply to the vCenter Server Appliance. To deploy the vCenter Server Appliance, see [“Download and Deploy the VMware vCenter Server Appliance,”](#) on page 92.

The following illustration shows the order of installation for vCenter Server components.



Installing vSphere includes the following tasks:

- 1 Read the vSphere release notes.
- 2 Verify that your system meets vSphere hardware and software requirements. See [Chapter 2, “System Requirements,”](#) on page 13.
- 3 Install vCenter Server.
 - a (Optional) Set up the vCenter Server database. See [“Preparing vCenter Server Databases,”](#) on page 30.

NOTE For small installations (up to 5 hosts and 50 virtual machines), you can use the bundled Microsoft SQL Server 2008 Express database, which you can choose to have the vCenter Server installer create during the vCenter Server installation.

- b Review the topics in [Chapter 3, “Before You Install vCenter Server,”](#) on page 29.

- c Use the topic [“Required Information for Installing or Upgrading vCenter Single Sign-On, Inventory Service, vCenter Server, and the vSphere Web Client,”](#) on page 62 to create a worksheet with the information you need for installation.
 - d Install vCenter Single Sign-On, the vSphere Web Client, Inventory Service, vCenter Server, and other vCenter Server support tools as appropriate. Connect to vCenter Server from the vSphere Web Client. See [Chapter 4, “Installing vCenter Server,”](#) on page 69 and [Chapter 5, “After You Install vCenter Server,”](#) on page 99.
 - e (Optional) Create a Linked Mode Group or join vCenter Server to a Linked Mode Group. See [“Creating vCenter Server Linked Mode Groups,”](#) on page 109.
- 4 Install ESXi.
- a Review the topics in [Chapter 6, “Before You Install ESXi,”](#) on page 121.
 - b Determine the ESXi installation option to use. See [“Options for Installing ESXi,”](#) on page 121.
 - c Determine where you will locate and boot the ESXi installer. See [“Media Options for Booting the ESXi Installer,”](#) on page 123. If you are PXE-booting the installer, verify that your network PXE infrastructure is properly set up. See [“PXE Booting the ESXi Installer,”](#) on page 127.
 - d Use the topic [“Required Information for ESXi Installation,”](#) on page 134 to create a worksheet with the information you will need when you install ESXi.
 - e Install ESXi, using the installation option you chose:
 - [“Installing ESXi Interactively,”](#) on page 137
 - [“Installing, Upgrading, or Migrating Hosts Using a Script,”](#) on page 140
 - [“Installing ESXi Using vSphere Auto Deploy,”](#) on page 155
- 5 Configure ESXi boot and network settings, the direct console, and other settings. See [Chapter 8, “Setting Up ESXi,”](#) on page 245 and [Chapter 9, “After You Install and Set Up ESXi,”](#) on page 265.
- 6 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. Optionally, you can install the vSphere Syslog Collector to collect logs from all hosts. See [“Providing Sufficient Space for System Logging,”](#) on page 23, [“Configure Syslog on ESXi Hosts,”](#) on page 260, and [“Install or Upgrade vSphere Syslog Collector,”](#) on page 106.

System Requirements

Systems running vCenter Server and ESXi instances must meet specific hardware and operating system requirements.

If you are using Auto Deploy to provision ESXi hosts, see also “[Preparing for vSphere Auto Deploy](#),” on page 165.

This chapter includes the following topics:

- “[ESXi Hardware Requirements](#),” on page 13
- “[Hardware Requirements for vCenter Server, the vSphere Web Client, vCenter Inventory Service, and vCenter Single Sign-On](#),” on page 17
- “[vCenter Server Software Requirements](#),” on page 22
- “[vSphere Web Client Software Requirements](#),” on page 22
- “[Providing Sufficient Space for System Logging](#),” on page 23
- “[Required Ports for vCenter Server](#),” on page 23
- “[Required Ports for the vCenter Server Appliance](#),” on page 26
- “[Conflict Between vCenter Server and IIS for Port 80](#),” on page 27
- “[DNS Requirements for vSphere](#),” on page 27
- “[Supported Remote Management Server Models and Minimum Firmware Versions](#),” on page 28

ESXi Hardware Requirements

Make sure the host meets the minimum hardware configurations supported by ESXi 5.5.

Hardware and System Resources

To install and use ESXi 5.5, 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 5.5 will install and run only on servers with 64-bit x86 CPUs.
- ESXi 5.5 requires a host machine with at least two cores.
- ESXi 5.5 supports only LAHF and SAHF CPU instructions.
- ESXi 5.5 requires the NX/XD bit to be enabled for the CPU in the BIOS.
- ESXi supports a broad range of x64 multicore processors. For a complete list of supported processors, see the VMware compatibility guide at <http://www.vmware.com/resources/compatibility>.

- ESXi requires a minimum of 4GB of physical RAM. Provide at least 8GB of RAM to take full advantage of ESXi features and 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 10Gb Ethernet controllers. For a list of supported network adapter models, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.
- Any combination of one or more of the following controllers:
 - Basic SCSI controllers. Adaptec Ultra-160 or Ultra-320, LSI Logic Fusion-MPT, or most NCR/Symbios SCSI.
 - RAID controllers. Dell PERC (Adaptec RAID or LSI MegaRAID), HP Smart Array RAID, or IBM (Adaptec) ServeRAID controllers.
- 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 will be considered remote, not local. These disks will not be 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 5.5 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>. ESXi 5.5 supports installing on and booting from the following storage systems:

- SATA disk drives. SATA disk drives connected behind supported SAS controllers or supported on-board SATA controllers.

Supported SAS controllers include:

- LSI1068E (LSISAS3442E)
- LSI1068 (SAS 5)
- IBM ServeRAID 8K SAS controller
- Smart Array P400/256 controller
- Dell PERC 5.0.1 controller

Supported on-board SATA include:

- Intel ICH9
- NVIDIA MCP55
- ServerWorks HT1000

NOTE ESXi does not support using local, internal SATA drives on the host server to create VMFS datastores that are shared across multiple ESXi hosts.

- Serial Attached SCSI (SAS) disk drives. Supported for installing ESXi and for storing virtual machines on VMFS partitions.
- Dedicated SAN disk on Fibre Channel or iSCSI
- USB devices. Supported for installing ESXi.

- Software Fibre Channel over Ethernet (FCoE). See [“Installing and Booting ESXi with Software FCoE,”](#) on page 134.

ESXi Booting Requirements

vSphere 5.5 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. Network booting or provisioning with VMware Auto Deploy requires the legacy BIOS firmware and is not available with UEFI.

ESXi can boot from a disk larger than 2TB provided that 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 5.5 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 5.5.

Storage Requirements for ESXi 5.5 Installation

Installing ESXi 5.5 requires a boot device that is a minimum of 1GB in size. When booting from a local disk or SAN/iSCSI LUN, a 5.2GB disk is required to allow for the creation of the VMFS volume and a 4GB scratch partition on the boot device. If a smaller disk or LUN is used, the installer will attempt to allocate a scratch region on a separate local disk. If a local disk cannot be found the scratch partition, `/scratch`, will be located 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, VMware recommends that you do not leave `/scratch` on the ESXi host ramdisk.

To reconfigure `/scratch`, see [“Set the Scratch Partition from the vSphere Web Client,”](#) on page 258.

Due to the I/O sensitivity of USB and SD devices the installer does not create a scratch partition on these devices. As such, there is no tangible benefit to using large USB/SD devices as ESXi uses only the first 1GB. When installing 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. You should reconfigure `/scratch` to use a persistent datastore following the installation.

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, it is not necessary to 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.

Recommendation for Enhanced ESXi Performance

To enhance performance, install 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,”](#) on page 13. See also the technical papers on vSphere 5 performance at <http://www.vmware.com/resources/techresources/cat/91,203,96>.

Table 2-1. 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	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.
Disk location	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.
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	Faster processors improve ESXi performance. For certain workloads, larger caches improve ESXi performance.
Hardware compatibility	Use devices in your server that are supported by ESXi 5.5 drivers. See the <i>Hardware Compatibility Guide</i> at http://www.vmware.com/resources/compatibility .

Hardware Requirements for vCenter Server, the vSphere Web Client , vCenter Inventory Service, and vCenter Single Sign-On

vCenter Server host machines must meet hardware requirements.

vCenter Single Sign-On, the vSphere Web Client , vCenter Inventory Service, and vCenter Server Hardware Requirements

You can install vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server on the same host machine (as with vCenter Simple Install) or on different machines. See [Table 2-2](#).

The following tables list the hardware requirements for vCenter Single Sign-On and Inventory Service, running on separate host machines.

■ [Table 2-3](#)

■ [Table 2-4](#)

If you use Custom Install to install vCenter Single Sign-On, vCenter Inventory Service, and vCenter Server on the same host machine, the vCenter Single Sign-On, and Inventory Service memory and disk storage requirements are in addition to the requirements for vCenter Server. See [Table 2-5](#).

Table 2-2. Minimum Hardware Requirements for Simple Install Deployment of vCenter Single Sign-On, the vSphere Web Client , vCenter Inventory Service, and vCenter Server

Host Hardware for Simple Install Deployment	Minimum Requirement
Processor	Intel or AMD x64 processor with two or more logical cores, each with a speed of 2GHz.
Memory	12GB. Memory requirements are higher if the vCenter Server database runs on the same machine as vCenter Server. vCenter Server includes several Java services: VMware VirtualCenter Management Webservices (tc Server), Inventory Service, and Profile-Driven Storage Service. When you install vCenter Server, you select the size of your vCenter Server inventory to allocate memory for these services. The inventory size determines the maximum JVM heap settings for the services. You can adjust this setting after installation if the number of hosts in your environment changes. See the recommendations in Table 2-7 .
Disk storage	100GB recommended. 40-60GB of free disk space are required after installation, depending on the size of your inventory. You should provide more space to allow for future growth of your inventory. Disk storage requirements are higher if the vCenter Server database runs on the same machine as vCenter Server, depending on the size of the database. In vCenter Server 5.x, the default size for vCenter Server logs is 450MB larger than in vCenter Server 4.x. Make sure the disk space allotted to the log folder is sufficient for this increase.
Network speed	1Gbps

Table 2-3. Minimum Hardware Requirements for vCenter Single Sign-On, Running on a Separate Host Machine from vCenter Server

vCenter Single Sign-On Hardware	Requirement
Processor	Intel or AMD x64 processor with two or more logical cores, each with a speed of 2GHz.
Memory	3GB. If vCenter Single Sign-On runs on the same host machine as vCenter Server, see Table 2-2 or Table 2-5 .
Disk storage	2GB.
Network speed	1Gbps

Table 2-4. Minimum Hardware Requirements for vCenter Inventory Service, Running on a Separate Host Machine from vCenter Server

vCenter Inventory Service Hardware	Requirement
Processor	Intel or AMD x64 processor with two or more logical cores, each with a speed of 2GHz.
Memory	3GB. If vCenter Inventory Service runs on the same host machine as vCenter Server, see Table 2-2 or Table 2-5 .
Disk storage	<p>If vCenter Inventory Service runs on the same host machine as vCenter Server, these requirements are in addition to the disk space required for vCenter Server and any other applications running on the vCenter Server host machine. See Table 2-5.</p> <p>Disk storage requirements for Inventory Service depend on inventory size and the amount of activity in the virtual machines in the inventory. At typical activity rates, Inventory Service uses 6GB - 12GB of disk space for 15,000 virtual machines distributed among 1,000 hosts.</p> <p>A high rate of activity (more than 20 percent of your virtual machines changing per hour) results in write-ahead logs (WAL) being written to disk to handle updates, instead of in-line writes into existing disk usage. This high rate of activity is often associated with Virtual Desktop Infrastructure (VDI) use cases.</p> <p>In the following guidelines for required disk space, a small inventory is 1-100 hosts or 1-1000 virtual machines, and a large inventory is more than 400 hosts or 4000 virtual machines.</p> <ul style="list-style-type: none"> ■ Small inventory, low activity rate: 5GB. ■ Small inventory, high activity rate: 15GB. ■ Large inventory, low activity rate: 15GB. ■ Large inventory, high activity rate: 40GB-60GB.
Network speed	1Gbps

Table 2-5. Minimum Hardware Requirements for vCenter Server

vCenter Server Hardware	Requirement
CPU	Two 64-bit CPUs or one 64-bit dual-core processor.
Processor	2.0GHz or faster Intel 64 or AMD 64 processor. The Itanium (IA64) processor is not supported. Processor requirements might be higher if the database runs on the same machine.

Table 2-5. Minimum Hardware Requirements for vCenter Server (Continued)

vCenter Server Hardware	Requirement
Memory	<p>The amount of memory needed depends on your vCenter Server configuration.</p> <ul style="list-style-type: none"> ■ If vCenter Server is installed on a different host machine than vCenter Single Sign-On and vCenter Inventory Service, 4GB of RAM are required. ■ If vCenter Server, vCenter Single Sign-On and vCenter Inventory Service are installed on the same host machine (as with vCenter Simple Install), 10GB of RAM are required. <p>Memory requirements are higher if the vCenter Server database runs on the same machine as vCenter Server. vCenter Server includes several Java services: VMware VirtualCenter Management Webservices (tc Server), Inventory Service, and Profile-Driven Storage Service. When you install vCenter Server, you select the size of your vCenter Server inventory to allocate memory for these services. The inventory size determines the maximum JVM heap settings for the services. You can adjust this setting after installation if the number of hosts in your environment changes. See the recommendations in Table 2-7.</p>
Disk storage	<p>The amount of disk storage needed for the vCenter Server installation depends on your vCenter Server configuration.</p> <ul style="list-style-type: none"> ■ If vCenter Server is installed on a different host machine than vCenter Single Sign-On and vCenter Inventory Service, 4GB are required. ■ If vCenter Server, vCenter Single Sign-On and vCenter Inventory Service are installed on the same host machine (as with vCenter Simple Install), at least 40-60GB of free disk space are required after installation, depending on the size of your inventory. You should provide more space to allow for future growth of your inventory. For guidelines about the disk space required for vCenter Single Sign-On and Inventory Service, see Table 2-3 and Table 2-4. <p>Disk storage requirements are higher if the vCenter Server database runs on the same machine as vCenter Server, depending on the size of those databases.</p> <p>In vCenter Server 5.x, the default size for vCenter Server logs is 450MB larger than in vCenter Server 4.x. Make sure the disk space allotted to the log folder is sufficient for this increase.</p>
Microsoft SQL Server 2008 R2 Express disk	<p>Up to 2GB free disk space to decompress the installation archive. Approximately 1.5GB of these files are deleted after the installation is complete.</p>
Network speed	1Gbps

NOTE Installing vCenter Server on a network drive or USB flash drive is not supported.

For the hardware requirements of your database, see your database documentation. The database requirements are in addition to the vCenter Server requirements if the database and vCenter Server run on the same machine.

vSphere Web Client Hardware Requirements

The vSphere Web Client has two components: A Java server and an Adobe Flex client application running in a browser.

Table 2-6. Hardware Requirements for the vSphere Web Client Server Component

vSphere Web Client Server Hardware	Requirement
Memory	At least 2GB: 1GB for the Java heap, and 1GB for <ul style="list-style-type: none"> ■ The resident code ■ The stack for Java threads ■ Global/bss segments for the Java process
CPU	2GHz processor with two or more cores
Disk Storage	At least 2GB free disk space
Networking	Gigabit connection recommended

JVM heap settings for vCenter Server

The JVM heap settings for vCenter Server depend on your inventory size. See [“Configuring VMware vCenter Server - tc Server Settings in vCenter Server,”](#) on page 114.

Table 2-7. JVM Heap Settings for vCenter Server

vCenter Server Inventory	VMware VirtualCenter Management Webservices (tc Server)	Inventory Service	Profile-Driven Storage Service
Small inventory (1-100 hosts or 1-1000 virtual machines)	1GB	3GB	1GB
Medium inventory (100-400 hosts or 1000-4000 virtual machines)	2GB	6GB	2GB
Large inventory (More than 400 hosts or 4000 virtual machines)	3GB	12GB	4GB

VMware vCenter Server Appliance Hardware Requirements and Recommendations

Table 2-8. Hardware Requirements for VMware vCenter Server Appliance

VMware vCenter Server Appliance Hardware	Requirement
Disk storage on the host machine	For most deployments, the vCenter Server Appliance requires at least 70GB of disk space, and is limited to a maximum size of 125GB. The required disk space depends on the size of your vCenter Server inventory. The vCenter Server Appliance can be deployed with thin-provisioned virtual disks that can grow to the maximum size of 125GB. If the host machine does not have enough free disk space to accommodate the growth of the vCenter Server Appliance virtual disks, vCenter Server might cease operation, and you will not be able to manage your vSphere environment.
Memory in the VMware vCenter Server Appliance	<p>Using the embedded PostgreSQL database, the vCenter Server Appliance supports up to 100 hosts or 3000 virtual machines, and has the following memory requirements</p> <ul style="list-style-type: none"> ■ Very small inventory (10 or fewer hosts, 100 or fewer virtual machines): at least 8GB. ■ Small inventory (10-50 hosts or 100-1500 virtual machines): at least 16GB. ■ Medium inventory (the maximum inventory supported with the embedded database; 50-100 hosts or 1500-3000 virtual machines): at least 24GB. <p>Using an external Oracle database, the vCenter Server Appliance supports up to 1000 hosts or 10000 registered virtual machines, and 10000 powered-on virtual machines, and has the following memory requirements:</p> <ul style="list-style-type: none"> ■ Very small inventory (10 or fewer hosts, 100 or fewer virtual machines): at least 4GB. ■ Small inventory (10-100 hosts or 100-1000 virtual machines): at least 8GB. ■ Medium inventory (100-400 hosts or 1000-4000 virtual machines): at least 16GB. ■ Large inventory (More than 400 hosts or 4000 virtual machines): at least 32GB.

For inventory and other configuration limits in the vCenter Server Appliance, see *Configuration Maximums*.

Table 2-9. JVM Heap Settings for VMware vCenter Server Appliance

vCenter Server Appliance Inventory	VMware VirtualCenter Management Webservices (tc Server)	Inventory Service	Profile-Driven Storage Service
Small inventory (1-100 hosts or 1-1000 virtual machines)	512MB	3GB	1GB
Medium inventory (100-400 hosts or 1000-4000 virtual machines)	512MB	6GB	2GB
Large inventory (More than 400 hosts or 4000 virtual machines)	1GB	12GB	4GB

See [“Configuring VMware vCenter Server - tc Server Settings in vCenter Server,”](#) on page 114.

vCenter Server Software Requirements

Make sure 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 its database.

For a list of supported operating systems, see the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility>. In the drop-down **What are you looking for** menu on that page, select **Host OS** and use the selection boxes to find compatible operating systems for your version of vCenter Server.

vCenter Server requires the Microsoft .NET 3.5 SP1 Framework. If it is not installed on your system, the vCenter Server installer installs it. The .NET 3.5 SP1 installation might require Internet connectivity to download more files.

NOTE If your vCenter Server host machine uses a non-English operating system, install both the Microsoft .NET Framework 3.5 SP1 and Microsoft .NET Framework 3.5 Language Pack through Windows Update. Windows Update automatically selects the correct localized version for your operating system. The .NET Framework installed through the vCenter Server installer includes only the English version.

vCenter Server 5.5 removes support for Windows Server 2003 as a host operating system. See the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php>.

vCenter Server 5.5 removes support for Windows Server 2008 SP1 as a host operating system. Upgrade Windows Server 2008 SP1 hosts to SP2 before upgrading vCenter Server to version 5.5. See the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php> and the Microsoft Software Lifecycle Policy at <http://support.microsoft.com/lifecycle/#ServicePackSupport>.

If you plan to use the Microsoft SQL Server 2008 R2 Express database that is bundled with vCenter Server, Microsoft Windows Installer version 4.5 (MSI 4.5) is required on your system. You can download MSI 4.5 from the Microsoft Web site. You can also install MSI 4.5 directly from the vCenter Server `autorun.exe` installer.

The VMware vCenter Server Appliance can be deployed only on hosts that are running ESX version 4.x or ESXi version 4.x or later.

vSphere Web Client Software Requirements

Make sure that your browser supports the vSphere Web Client.

VMware has tested and supports the following guest operating systems and browser versions for the vSphere Web Client.

Table 2-10. Supported guest operating systems and browser versions for the vSphere Web Client .

Operating system	Browser
Windows 32-bit and 64-bit	<p>Microsoft Internet Explorer 8, 9 (64-bit only), and 10.</p> <p>Mozilla Firefox: the latest browser version, and the one previous version at the time the vSphere 5.5 is produced.</p> <p>Google Chrome: the latest browser version, and the one previous version at the time the vSphere 5.5 is produced.</p>
Mac OS	<p>Mozilla Firefox: the latest browser version, and the one previous version at the time the vSphere 5.5 is produced.</p> <p>Google Chrome: the latest browser version, and the one previous version at the time the vSphere 5.5 is produced.</p>

Later versions of these browsers are likely to work, but have not been tested.

The vSphere Web Client requires the Adobe Flash Player version 11.5.0 or later to be installed with the appropriate plug-in for your browser.

Providing Sufficient Space for System Logging

ESXi 5.x uses a new log infrastructure. If your host is deployed with Auto Deploy, 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 for system logging exists.

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.

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

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 5.x autoconfigures logs to best suit your installation, and provides enough space to accommodate log messages.

Table 2-11. 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)	10240KB	10	100MB
VirtualCenter Agent (vpxa)	5120KB	10	50MB
vSphere HA agent (Fault Domain Manager, fdm)	5120KB	10	50MB

For information about setting up a remote log server, see [“Configure Syslog on ESXi Hosts,”](#) on page 260 and [“Install or Upgrade vSphere Syslog Collector,”](#) on page 106.

Required Ports for vCenter Server

The VMware vCenter Server system must be able to send data to every managed host and receive data from every vSphere Web Client. To enable migration and provisioning activities between managed hosts, the source and destination hosts must be able to receive data from each other.

For information about ports required for the vCenter Server Appliance, see [“Required Ports for the vCenter Server Appliance,”](#) on page 26.

VMware uses designated ports for communication. Additionally, the managed hosts monitor designated ports for data from the vCenter Server system. If a firewall exists between any of these elements and Windows firewall service is in use, the installer opens the ports during the installation. 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, a firewall is enabled by default.

Table 2-12. Ports Required for Communication Between Components

Port	Description
22	SSH Server (vSphere Client)
53	DNS Client
80	<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 <code>http://server</code> instead of <code>https://server</code>. WS-Management (also requires port 443 to be open)</p> <p>If you use a custom Microsoft SQL database (not the bundled SQL Server 2008 database) that is stored on the same host machine as the vCenter Server, port 80 is used by the SQL Reporting Service. When you install vCenter Server, the installer will prompt you to change the HTTP port for vCenter Server. Change the vCenter Server HTTP port to a custom value to ensure a successful installation. Microsoft Internet Information Services (IIS) also use port 80. See “Conflict Between vCenter Server and IIS for Port 80,” on page 27.</p>
88	Control interface RPC for Kerberos, used by vCenter Single Sign-On
111	RPC service that is used for the NIS register by the vCenter Server Appliance
123	NTP Client
135	Used to join vCenter Virtual Appliance to an Active Directory domain.
161	SNMP Server
389	<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. The vCenter Server system needs to bind to port 389, even if you are not joining this vCenter Server instance to a Linked Mode 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>
427	The CIM client uses the Service Location Protocol, version 2 (SLPv2) to find CIM servers.
443	<p>The default port that the vCenter Server system uses to listen for connections from the vSphere Client. To enable the vCenter Server system to receive data from the vSphere Client, open port 443 in the firewall.</p> <p>The vCenter Server system also uses port 443 to monitor data transfer from SDK clients. 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) ■ vSphere Client access to vSphere Update Manager ■ Third-party network management client connections to vCenter Server ■ Third-party network management clients access to hosts
513	vCenter Virtual Appliance used for logging activity
636	For vCenter Server Linked Mode, this is the SSL port of the local instance. 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 SSL service on any port from 1025 through 65535.
902	<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 vSphere Client and the hosts. The vSphere Client uses this port to display virtual machine consoles</p>
903	<p>Access a virtual machine console from the vSphere Client when the vSphere Client is connected directly to the ESXi host (no vCenter Server).</p> <p>MKS transactions (xinetd/vmware-authd-mks)</p>
1234, 1235	vSphere Replication
2012	Control interface RPC for vCenter Single Sign-On vmdir.
2013	Control interface RPC for Kerberos, used by vCenter Single Sign-On
2014	RPC port for all VMCA (VMware Certificate Authority) APIs

Table 2-12. Ports Required for Communication Between Components (Continued)

Port	Description
2049	Transactions from NFS storage devices This port is used on the VMkernel interface.
3260	Transactions to iSCSI storage devices
5900-5964	RFB protocol, which is used by management tools such as VNC
5988	CIM transactions over HTTP
5989	CIM XML transactions over HTTPS
6501	Auto Deploy service
6502	Auto Deploy management
7005	vCenter Single Sign-On
7009	vCenter Single Sign-On
7080	vCenter Single Sign-On
7331	vSphere Web Client - HTML5 Remote Console
7444	vCenter Single Sign-On HTTPS
8000	Requests from vMotion
8009	AJP connector port for vCenter Server Appliance communication with Tomcat
8080	Web Services HTTP. Used for the VMware VirtualCenter Management Web Services.
8100	Traffic between hosts for vSphere Fault Tolerance (FT)
8182	Traffic between hosts for vSphere High Availability (HA)
8200	Traffic between hosts for vSphere Fault Tolerance (FT)
8443	Web Services HTTPS. Used for the VMware VirtualCenter Management Web Services.
9009	Used to allow a vCenter Server Appliance to communicate with the vSphere Web Client.
9090	vSphere Web Client HTTP
9443	vSphere Web Client HTTPS
9875 - 9877	vSphere Web Client Java Management Extension (JMX). Dynamically acquired upon the vSphere Web Client service starting.
10080	vCenter Inventory Service HTTP
10109	vCenter Inventory Service Management
10111	vCenter Inventory Service Linked Mode Communication
10443	vCenter Inventory Service HTTPS
11711	vCenter Single Sign-On LDAP
11712	vCenter Single Sign-On LDAPS
12721	VMware Identity Management service
49000 - 65000	vCenter Single Sign-On - VMware Identity Management Service. Dynamically acquired when the VMware Identity Management Service starts.
60099	Web Service change service notification port

To have the vCenter Server system use a different port to receive vSphere Web Client data, see the *vCenter Server and Host Management* documentation.

For a discussion of firewall configuration, see the *vSphere Security* documentation.

Required Ports for the vCenter Server Appliance

The VMware vCenter Server system must be able to send data to every managed host and receive data from every vSphere Web Client. For migration and provisioning activities between managed hosts, the source and destination hosts must be able to receive data from each other.

For information about ports required for vCenter Server on Windows, see [“Required Ports for vCenter Server,”](#) on page 23.

VMware uses designated ports for communication. Additionally, the managed hosts monitor designated ports for data from the vCenter Server system. The vCenter Server Appliance is preconfigured to use the ports listed in [Table 2-13](#). 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.

Table 2-13. Ports Required for the vCenter Server Appliance

Port	Description
80	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 <code>http://server</code> instead of <code>https://server</code> .
443	The vCenter Server system uses port 443 to monitor data transfer from SDK clients.
902	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.
8080	Web Services HTTP. Used for the VMware VirtualCenter Management Web Services.
8443	Web Services HTTPS. Used for the VMware VirtualCenter Management Web Services.
10080	vCenter Inventory Service HTTP
10443	vCenter Inventory Service HTTPS
10109	vCenter Inventory Service database
514	vSphere Syslog Collector server
1514	vSphere Syslog Collector server (SSL)
6500	Network coredump server (UDP)
6501	Auto Deploy service
6502	Auto Deploy management
9090	vSphere Web Client HTTP
9443	vSphere Web Client HTTPS
5480	vCenter Server Appliance Web user interface HTTPS
5489	vCenter Server Appliance Web user interface CIM service
22	System port for SSHD

To have the vCenter Server system use a different port to receive vSphere Web Client data, see the *vCenter Server and Host Management* documentation.

For a discussion of firewall configuration, see the *vSphere Security* documentation.

Conflict Between vCenter Server and IIS for Port 80

vCenter Server and Microsoft Internet Information Service (IIS) both use port 80 as the default port for direct HTTP connections. This conflict can cause vCenter Server to fail to restart after the installation of vSphere Authentication Proxy.

Problem

vCenter Server fails to restart after the installation of vSphere Authentication Proxy is complete.

Cause

If you do not have IIS installed when you install vSphere Authentication Proxy, the installer prompts you to install IIS. Because IIS uses port 80, which is the default port for vCenter Server direct HTTP connections, vCenter Server fails to restart after the installation of vSphere Authentication Proxy is complete. See [“Required Ports for vCenter Server,”](#) on page 23.

Solution

- ◆ To resolve a conflict between IIS and vCenter Server for port 80, take one of the following actions.

Option	Description
If you installed IIS before installing vCenter Server	Change the port for vCenter Server direct HTTP connections from 80 to another value.
If you installed vCenter Server before installing IIS	Before restarting vCenter Server, change the binding port of the IIS default Web site from 80 to another value.

DNS Requirements for vSphere

You install vCenter Server, like any other network server, on a 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.

Ensure that the ESXi host management interface has a valid DNS resolution from the vCenter Server and all vSphere Web Clients. Ensure that the vCenter Server has a valid DNS resolution from all ESXi hosts and all vSphere Web Clients.

Ensure that the vCenter Server is installed on a machine that has a resolvable fully qualified domain name (FQDN). To check that the FQDN is resolvable, type **nslookup your_vCenter_Server_fqdn** at a command line prompt. If the FQDN is resolvable, the **nslookup** command returns the IP and name of the domain controller machine.

Ensure that DNS reverse lookup returns a fully qualified domain name when queried with the IP address of the vCenter Server. When you install vCenter Server, 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 of the vCenter Server from its IP address. Reverse lookup is implemented using PTR records. To create a PTR record, see the documentation for your vCenter Server host operating system.

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). Ping the computer name to test the connection. For example, if the computer name is `host-1.company.com`, run the following command in the Windows command prompt:

```
ping host-1.company.com
```

If you can ping the computer name, the name is updated in DNS.

Supported Remote Management Server Models and Minimum Firmware Versions

You can use remote management applications to install ESXi or for remote management of hosts.

Table 2-14. Supported Remote Management Server Models and Firmware Versions

Remote Controller Make and Model	Firmware Version	Java
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
IBM RSA 2	1.03, 1.2	1.6.0_22

Before You Install vCenter Server

You can install vCenter Server on a physical system or on a virtual machine running on an ESXi host. You can also download the VMware vCenter Server Appliance, a preconfigured Linux-based virtual machine optimized for running vCenter Server.

This chapter includes the following topics:

- [“Preparing vCenter Server Databases,”](#) on page 30
- [“Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server,”](#) on page 49
- [“How vCenter Single Sign-On Affects vCenter Server Installation,”](#) on page 51
- [“vCenter Single Sign-On Deployment Modes,”](#) on page 52
- [“vCenter Single Sign-On and High Availability,”](#) on page 54
- [“vCenter Single Sign-On Components,”](#) on page 56
- [“Setting the vCenter Server Administrator User,”](#) on page 56
- [“Authenticating to the vCenter Server Environment,”](#) on page 57
- [“How vCenter Single Sign-On Affects Log In Behavior,”](#) on page 57
- [“Identity Sources for vCenter Server with vCenter Single Sign-On,”](#) on page 58
- [“Synchronizing Clocks on the vSphere Network,”](#) on page 59
- [“Using a User Account for Running vCenter Server,”](#) on page 60
- [“Installing vCenter Server on IPv6 Machines,”](#) on page 60
- [“JDBC URL Formats for the vCenter Server Database,”](#) on page 61
- [“Running the vCenter Server Installer from a Network Drive,”](#) on page 62
- [“Required Information for Installing or Upgrading vCenter Single Sign-On, Inventory Service, vCenter Server, and the vSphere Web Client,”](#) on page 62
- [“Microsoft SQL Database Set to Unsupported Compatibility Mode Causes vCenter Server Installation or Upgrade to Fail,”](#) on page 68

Preparing vCenter Server Databases

vCenter Server and vSphere Update Manager require databases to store and organize server data.

Each vCenter Server instance must have its own database. For small installations (up to 5 hosts and 50 virtual machines), you can use the bundled Microsoft SQL Server 2008 Express database, which you can choose to have the vCenter Server installer create during the vCenter Server installation. Larger installations require a supported database, prepared as described in the subtopics of this section.

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 do not need to install a new database server for the vCenter Server installation to work. During vCenter Server installation, you can point the vCenter Server system to any existing supported database. vCenter Server supports Oracle and Microsoft SQL Server databases. Update Manager 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.



CAUTION If you have a VirtualCenter database that you want to preserve, do not perform a fresh installation of vCenter Server. See the *vSphere Upgrade* documentation.

VMware recommends using separate databases for vCenter Server and Update Manager. For a small deployments, a separate database for Update Manager might not be necessary.

vCenter Server Database Configuration Notes

After you choose a supported database type, make sure you understand any special configuration requirements.

Table 3-1 is not a complete list of databases supported with vCenter Server. For information about specific database versions and service pack configurations supported with vCenter Server, see the [VMware Product Interoperability Matrixes](#). This topic is intended only to provide special database configuration notes not listed in the Product Interoperability Matrixes.

vCenter Server databases require a UTF code set.

Contact your DBA for the appropriate database credentials.

Table 3-1. Configuration Notes for Databases Supported with vCenter Server

Database Type	Configuration Notes
Microsoft SQL Server 2008 R2 Express	<p>Bundled database that you can use for small deployments of up to 5 hosts and 50 virtual machines.</p> <p>SQL Server Collation Model: SQL_Latin1_General_CP1_CI_AS.</p> <p>ODBC System DSN minimum version: SQL Native Client 10.0 (version 2009.100.4000.00), which you can obtain as a free download from the microsoft.com Download Center.</p> <p>NOTE This database is not supported for the vCenter Server Appliance.</p>
Microsoft SQL Server 2005	<p>Ensure that the machine has a valid ODBC DSN entry.</p> <p>If Microsoft SQL Server 2005 is not already installed and the machine has MSXML Core Services 6.0 installed, remove MSXML Core Services 6.0 before installing Microsoft SQL Server 2005. If you cannot remove it using the Add or Remove Programs utility, use the Windows Installer CleanUp utility. See http://support.microsoft.com/kb/968749.</p> <p>NOTE This database is not supported for the vCenter Server Appliance.</p>

Table 3-1. Configuration Notes for Databases Supported with vCenter Server (Continued)

Database Type	Configuration Notes
Microsoft SQL Server 2008	<p>Ensure that the machine has a valid ODBC DSN entry.</p> <p>SQL Server Collation Model: SQL_Latin1_General_CP1_CI_AS.</p> <p>ODBC System DSN minimum version: SQL Native Client 10.0 (version 2009.100.4000.00), which you can obtain as a free download from the microsoft.com Download Center.</p> <p>NOTE This database is not supported for the vCenter Server Appliance.</p>
Oracle	<p>Ensure that the machine has a valid ODBC DSN entry.</p> <p>After you complete the vCenter Server installation, take the following steps:</p> <ul style="list-style-type: none"> ■ Apply the latest patch to the Oracle client and server. ■ Copy the Oracle JDBC driver (ojdbc14.jar or ojdbc5.jar) to the vCenter Server installation directory, in the <code>tomcat\lib</code> subdirectory: <i>vCenter install location\Infrastructure\tomcat\lib</i>. ■ In the Services section of the Windows Administrative Tools control panel, restart the VMware VirtualCenter Management Webservices service. <p>The vCenter Server installer attempts to copy the Oracle JDBC driver from the Oracle client location to the vCenter Server installation directory. If the Oracle JDBC driver is not found in the Oracle client location, the vCenter Server installer prompts you to copy the file manually. You can download the file from the oracle.com Web site.</p>

Create a 64-Bit DSN

The vCenter Server system must have a 64-bit DSN. This requirement applies to all supported databases.

Procedure

- 1 Select **Control Panel > Administrative Tools > Data Sources (ODBC)**.
- 2 Use the application to create a system DSN.
If you have a Microsoft SQL database, create the system DSN for the SQL Native Client driver.
- 3 Test the connectivity.

The system now has a DSN that is compatible with vCenter Server. When the vCenter Server installer prompts you for a DSN, select the 64-bit DSN.

Confirm That vCenter Server Can Communicate with the Local Database

If your database is located on the same machine on which vCenter Server will be installed, and you have changed the name of this machine, make sure 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.

The name change has no effect on communication with remote databases. You can skip this procedure if your database is remote.

Check with your database administrator or the database vendor to make sure all components of the database are working after you rename the server.

Prerequisites

- Make sure 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 Ping the computer name to test this connection.

For example, if the computer name is `host-1.company.com`, run the following command in the Windows command prompt:

```
ping host-1.company.com
```

If you can ping the computer name, the name is updated in DNS.

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.

Configure Microsoft SQL Server Databases

To use a Microsoft SQL database for your vCenter Server repository, configure your database to work with vCenter Server.

Procedure

- 1 [Create a SQL Server Database and User for vCenter Server](#) on page 33
You must create a database and user for vCenter Server. To simplify the process, you can use a script.
- 2 [Set Database Permissions By Manually Creating Database Roles and the VMW Schema](#) on page 34
By using this recommended method, available with vCenter Server 5.x, the vCenter Server database administrator can set permissions for vCenter Server users and administrators to be granted through Microsoft SQL Server database roles.
- 3 [Set Database Permissions by Using the dbo Schema and the db_owner Database Role](#) on page 35
If you use Microsoft SQL Server database, the simplest way to assign permissions for a vCenter Server database user is through the database role **db_owner**. Assign the **db_owner** role to the vCenter Server database user on both the vCenter and MSDB databases.
- 4 [Use a Script to Create a vCenter Server User by Using the dbo Schema and db_owner Database Role](#) on page 35
If you set database permissions by using the dbo schema and db_owner database role, you can use a script to create a vCenter Server user with the db_owner database role.
- 5 [Use a Script to Create a Microsoft SQL Server Database Schema and Roles](#) on page 36
In this recommended method of configuring the SQL database, you create the custom schema VMW, instead of using the existing dbo schema.
- 6 [\(Optional\) Use a Script to Create Microsoft SQL Server Database Objects Manually](#) on page 37
You can create database objects manually with this method of configuring the SQL database.

7 [Configure a SQL Server ODBC Connection](#) on page 40

After you create a vCenter Server user, establish a connection with a SQL Server database. This connection is required to install a vCenter Server system.

8 [Configure Microsoft SQL Server TCP/IP for JDBC](#) on page 41

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.

9 [\(Optional\) Configure a Microsoft SQL Server Database User to Enable Database Monitoring](#) on page 42

vCenter Server Database Monitoring captures metrics that enable the administrator to assess the status and health of the database server. Enabling Database Monitoring helps the administrator prevent vCenter downtime because of a lack of resources for the database server.

Create a SQL Server Database and User for vCenter Server

You must create a database and user for vCenter Server. To simplify the process, you can use a script.

In the script, you can customize the location of the data and log files.

The user that is created by this script is not subject to any security policy. Change the passwords as appropriate.

Procedure

- 1 Log in to a Microsoft SQL Server Management Studio session as the sysadmin (SA) or a user account with **sysadmin** privileges.
- 2 Run the following script.

The script is located in the vCenter Server installation package at `/<installation directory>/vCenter-Server/dbschema/DB_and_schema_creation_scripts_MSSQL.txt`.

```
use [master]
go
CREATE DATABASE [VCDB] ON PRIMARY
(NAME = N'vcdb', FILENAME = N'C:\VCDB.mdf', SIZE = 3000KB, FILEGROWTH = 10% )
LOG ON
(NAME = N'vcdb_log', FILENAME = N'C:\VCDB.ldf', SIZE = 1000KB, FILEGROWTH = 10%)
COLLATE SQL_Latin1_General_CP1_CI_AS
go
use VCDB
go
sp_addlogin @loginame=[vpxuser], @passwd=N'vpxuser!0', @defdb='VCDB',
@deflanguage='us_english'
go
ALTER LOGIN [vpxuser] WITH 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.

What to do next

See [“Set Database Permissions By Manually Creating Database Roles and the VMW Schema,”](#) on page 34.

Set Database Permissions By Manually Creating Database Roles and the VMW Schema

By using this recommended method, available with vCenter Server 5.x, the vCenter Server database administrator can set permissions for vCenter Server users and administrators to be granted through Microsoft SQL Server database roles.

VMware recommends this method because it removes the requirement to set up the database dbo schema and **db_owner** role for vCenter Server users who install and upgrade vCenter Server.

Alternatively, you can assign vCenter Server database permissions by creating and assigning the **db_owner** role and letting the vCenter Server installer create the default schema that assigns database user permissions to that role. See [“Set Database Permissions by Using the dbo Schema and the db_owner Database Role,”](#) on page 35.

Prerequisites

Create the vCenter Server database. See [“Create a SQL Server Database and User for vCenter Server,”](#) on page 33

Procedure

- 1 Create the database VCDB and the database schema VMW in VCDB.
- 2 Assign the default schema VMW to the user [vpxuser].
- 3 In the vCenter Server database, create the user role **VC_ADMIN_ROLE**.
- 4 In the vCenter Server database, grant privileges to the **VC_ADMIN_ROLE**.
 - a Grant the schema permissions **ALTER**, **REFERENCES**, and **INSERT**.
 - b Grant the permissions **CREATE TABLE**, **VIEW**, and **CREATE PROCEDURES**.
- 5 In the vCenter Server database, create the **VC_USER_ROLE**.
- 6 In the vCenter Server database, grant the schema permissions **SELECT**, **INSERT**, **DELETE**, **UPDATE**, and **EXECUTE** to the **VC_USER_ROLE**.
- 7 Grant the **VC_USER_ROLE** to the user [vpxuser].
- 8 Grant the **VC_ADMIN_ROLE** to the user [vpxuser].
- 9 In the MSDB database, create the user [vpxuser].
- 10 In the MSDB database, create the user role **VC_ADMIN_ROLE**.
- 11 Grant privileges to the **VC_ADMIN_ROLE** in MSDB.
 - a On the MSDB tables syscategories, sysjobsteps, and sysjobs, grant the **SELECT** permission to the user [vpxuser].
 - b On the MSDB stored procedures sp_add_job, sp_delete_job, sp_add_jobstep, sp_update_job, sp_add_jobserver, sp_add_jobschedule, and sp_add_category, grant the **EXECUTE** permission to the role **VC_ADMIN_ROLE**.
- 12 In the MSDB database, grant the **VC_ADMIN_ROLE** to the user [vpxuser].
- 13 Connect to the vCenter Server database as user [vpxuser] and create the ODBC DSN.
- 14 Install vCenter Server.

- 15 Revoke the **VC_ADMIN_ROLE** from the user **[vpxuser]** in the vCenter Server database.

After you revoke the role, you can leave the role as inactive for use in future upgrades, or remove the role for increased security. If you remove the role, you must recreate the role and assign it to the user **[vpxuser]** before any future upgrade of vCenter Server.

The hardcoded **dbo** role is removed from `Vcdb_mssql.sql`.

What to do next

[“Use a Script to Create a Microsoft SQL Server Database Schema and Roles,”](#) on page 36

Set Database Permissions by Using the dbo Schema and the db_owner Database Role

If you use Microsoft SQL Server database, the simplest way to assign permissions for a vCenter Server database user is through the database role **db_owner**. Assign the **db_owner** role to the vCenter Server database user on both the vCenter and MSDB databases.

Alternatively, experienced database administrators can set permissions by creating database roles and the VMW schema manually. See [“Set Database Permissions By Manually Creating Database Roles and the VMW Schema,”](#) on page 34 and [“Use a Script to Create a Microsoft SQL Server Database Schema and Roles,”](#) on page 36. That method, available beginning with vSphere 5.0, is recommended, because it gives the database administrator greater control over database permissions. The recommended method also removes the requirement to set up the database **dbo** schema and **db_owner** role for vCenter Server users who install and upgrade vCenter Server.

Prerequisites

Create the vCenter Server database. See [“Create a SQL Server Database and User for vCenter Server,”](#) on page 33

Procedure

- 1 Assign the role **dbo** to the vCenter Server and Microsoft SQL databases.
- 2 For any user who will install or upgrade vCenter Server, assign the user the default schema **dbo**.

When you install vCenter Server, the installer uses the default **dbo** schema to assign permissions to the **db_owner** role.

Use a Script to Create a vCenter Server User by Using the dbo Schema and db_owner Database Role

If you set database permissions by using the **dbo** schema and **db_owner** database role, you can use a script to create a vCenter Server user with the **db_owner** database role.

Alternatively, experienced database administrators can set permissions by creating database roles and the VMW and SQL Server database schemas. See [“Set Database Permissions By Manually Creating Database Roles and the VMW Schema,”](#) on page 34 [“Use a Script to Create a Microsoft SQL Server Database Schema and Roles,”](#) on page 36. That method, available beginning with vSphere 5.0, is recommended, because it gives the database administrator greater control over database permissions. That method removes the requirement to set up the database role **dbo** and **db_owner** schema for vCenter Server users who install and upgrade vCenter Server.

Prerequisites

Create the vCenter Server database. See [“Create a SQL Server Database and User for vCenter Server,”](#) on page 33

Procedure

- 1 Log in to a Microsoft SQL Server Management Studio session as the sysadmin (SA) or a user account with **sysadmin** privileges.
- 2 Run the following script.

The script is located in the vCenter Server installation package */installation directory/vCenter-Server/dbschema/DB_and_schema_creation_scripts_MSSQL.txt* file.

```
use Vcdb
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
use MSDB
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
```

What to do next

[“Configure a SQL Server ODBC Connection,”](#) on page 40

Use a Script to Create a Microsoft SQL Server Database Schema and Roles

In this recommended method of configuring the SQL database, you create the custom schema VMW, instead of using the existing dbo schema.

This method requires that you create new database roles and grant them to the database *user*. See [“Set Database Permissions By Manually Creating Database Roles and the VMW Schema,”](#) on page 34 and [“Use a Script to Create a Microsoft SQL Server Database Schema and Roles,”](#) on page 36.

Prerequisites

Create the SQL Server database and user for vCenter Server. You can create the database manually or by using a script. See [“Create a SQL Server Database and User for vCenter Server,”](#) on page 33

Procedure

- 1 Log in to a Microsoft SQL Server Management Studio session as the sysadmin (SA) or a user account with sysadmin privileges.
- 2 Run the following script.

The script is located in the vCenter Server installation package at */installation directory/vCenter-Server/dbschema/DB_and_schema_creation_scripts_MSSQL.txt*

```
CREATE SCHEMA [VMW]
go
ALTER USER [vpxuser] WITH DEFAULT_SCHEMA = [VMW]
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
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 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

```

(Optional) Use a Script to Create Microsoft SQL Server Database Objects Manually

You can create database objects manually with this method of configuring the SQL database.

Alternatively, you can configure a SQL Server ODBC connection and run the Install package. The vCenter Server installer will create database objects. See [“Configure a SQL Server ODBC Connection,”](#) on page 40.

Using a script to create database objects manually requires that you take one of the following actions.

- Grant the **db_owner** role to the database *user* in VCDB and in MSDB. See [“Set Database Permissions by Using the dbo Schema and the db_owner Database Role,”](#) on page 35 and [“Use a Script to Create a vCenter Server User by Using the dbo Schema and db_owner Database Role,”](#) on page 35.

- Grant the **VC_ADMIN_ROLE** to the database *user* in VCDB and in MSDB, and grant the **VC_USER_ROLE** to the database *user* in VCDB. See [“Set Database Permissions By Manually Creating Database Roles and the VMW Schema,”](#) on page 34.

Prerequisites

Create the SQL Server database. You can create the SQL Server database manually or by using a script. See [“Create a SQL Server Database and User for vCenter Server,”](#) on page 33

Procedure

- 1 Log in to a Microsoft SQL Server Management Studio session as *user_name* for a user account that you created on the vCenter Server and MSDB databases.
- 2 Locate the dbschema scripts in the vCenter Server installation package `/installation_directory/vCenter-Server/dbschema` directory.
- 3 Open the `sql` file through Microsoft SQL Server Management Studio and replace all occurrences of `$schema` with the schema name in the file `VCDB_mssql.SQL`.
- 4 Run the scripts in 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 order shown here.

```
VCDB_mssql.SQL
insert_stats_proc_mssql.sql
load_stats_proc_mssql.sql
purge_stat2_proc_mssql.sql
purge_stat3_proc_mssql.sql
purge_usage_stats_proc_mssql.sql
stats_rollup1_proc_mssql.sql
stats_rollup2_proc_mssql.sql
stats_rollup3_proc_mssql.sql
cleanup_events_mssql.sql
delete_stats_proc_mssql.sql
upsert_last_event_proc_mssql.sql
load_usage_stats_proc_mssql.sql
TopN_DB_mssql.sql
calc_topn1_proc_mssql.sql
calc_topn2_proc_mssql.sql
calc_topn3_proc_mssql.sql
calc_topn4_proc_mssql.sql
clear_topn1_proc_mssql.sql
clear_topn2_proc_mssql.sql
clear_topn3_proc_mssql.sql
clear_topn4_proc_mssql.sql
rule_topn1_proc_mssql.sql
rule_topn2_proc_mssql.sql
rule_topn3_proc_mssql.sql
rule_topn4_proc_mssql.sql
process_license_snapshot_mssql.sql
l_stats_rollup3_proc_mssql.sql
l_purge_stat2_proc_mssql.sql
l_purge_stat3_proc_mssql.sql
l_stats_rollup1_proc_mssql.sql
l_stats_rollup2_proc_mssql.sql
VCDB_views_mssql.sql
```

- 5 (Optional) You can also run the following scripts to enable database health monitoring.

```
job_dbm_performance_data_mssql.sql
process_performance_data_mssql.sql
```

- 6 For all supported editions of Microsoft SQL Server (except Microsoft SQL Server 2005 Express and Microsoft SQL Server 2008 R2 Express), run these scripts to set up scheduled jobs on the database.

These scripts ensure that the SQL Server Agent service is running.

```
job_schedule1_mssql.sql
job_schedule2_mssql.sql
job_schedule3_mssql.sql
job_cleanup_events_mssql.sql
job_topn_past_day_mssql.sql
job_topn_past_week_mssql.sql
job_topn_past_month_mssql.sql
job_topn_past_year_mssql.sql
job_property_bulletin_mssql.sql
```

- 7 For all the procedures you created in [Step 4](#), grant the execute privilege to the vCenter Server database.

```
grant execute on insert_stats_proc to vCenter_db_user
grant execute on purge_stat2_proc to vCenter_db_user
grant execute on purge_stat3_proc to vCenter_db_user
grant execute on purge_usage_stat_proc to vCenter_db_user
grant execute on stats_rollup1_proc to vCenter_db_user
grant execute on stats_rollup2_proc to vCenter_db_user
grant execute on stats_rollup3_proc to vCenter_db_user
grant execute on cleanup_events_tasks_proc to vCenter_db_user
grant execute on delete_stats_proc to vCenter_db_user
grant execute on upsert_last_event_proc to vCenter_db_user
grant execute on load_usage_stats_proc to vCenter_db_user
grant execute on load_stats_proc to vCenter_db_user
grant execute on calc_topn1_proc to vCenter_db_user
grant execute on calc_topn2_proc to vCenter_db_user
grant execute on calc_topn3_proc to vCenter_db_user
grant execute on calc_topn4_proc to vCenter_db_user
grant execute on clear_topn1_proc to vCenter_db_user
grant execute on clear_topn2_proc to vCenter_db_user
grant execute on clear_topn3_proc to vCenter_db_user
grant execute on clear_topn4_proc to vCenter_db_user
grant execute on rule_topn1_proc to vCenter_db_user
grant execute on rule_topn2_proc to vCenter_db_user
grant execute on rule_topn3_proc to vCenter_db_user
grant execute on rule_topn4_proc to vCenter_db_user
grant execute on process_license_snapshot_proc to vCenter_db_user
grant execute on l_stats_rollup3_proc to vCenter_db_user
grant execute on l_purge_stat2_proc to vCenter_db_user
grant execute on l_purge_stat3_proc to vCenter_db_user
grant execute on l_stats_rollup1_proc to vCenter_db_user
grant execute on l_stats_rollup2_proc to vCenter_db_user
```

If you ran the script `process_performance_data_mssql.sql` in [Step 4](#), grant the following execute privilege to the vCenter Server database.

```
grant execute on process_performance_data_proc to vCenter_db_user
```

- 8 On the machine on which you intend to install vCenter Server, create a DSN that points to the database server with the schema.

- 9 Run the vCenter Server installer.
- 10 If a database reinitialization warning message appears in the vCenter Server installer, select **Do not overwrite, leave my existing database in place** and continue the installation.

This message appears if you are using a database that has vCenter Server tables that were created by a previous installation. The message does not appear if the database is clean.

If you leave your existing database in place, you cannot join a Linked Mode group during the installation. You can join after the installation is complete. See [“Join a Linked Mode Group After Installation,”](#) on page 112.

- 11 When prompted, provide the database user login.

Configure a SQL Server ODBC Connection

After you create a vCenter Server user, establish a connection with a SQL Server database. This connection is required to install a vCenter Server system.

If you use SQL Server for vCenter Server, do not use the master database.

See your Microsoft SQL ODBC documentation for specific instructions regarding 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

- Review the required database patches specified in [“vCenter Server Database Configuration Notes,”](#) on page 30.
- Create a database using SQL Server Management Studio on the SQL Server. See [“Create a SQL Server Database and User for vCenter Server,”](#) on page 33
- Set database permissions using one of the following options:
 - Option 1 (recommended): Follow the procedures in [“Set Database Permissions By Manually Creating Database Roles and the VMW Schema,”](#) on page 34 and [“Use a Script to Create a Microsoft SQL Server Database Schema and Roles,”](#) on page 36
 - Option 2 (alternative): Follow the procedures in [“Set Database Permissions by Using the dbo Schema and the db_owner Database Role,”](#) on page 35 and [“Use a Script to Create a vCenter Server User by Using the dbo Schema and db_owner Database Role,”](#) on page 35.

Procedure

- 1 On your vCenter Server system, select **Settings > Control Panel > Administrative Tools > Data Sources (ODBC)**.
- 2 Click the **System DSN** tab and do one of the following.
 - To modify an existing SQL Server ODBC connection, select the connection from the System Data Source list and click **Configure**.
 - To create a new SQL Server ODBC connection, click **Add**, select **SQL Native Client**, and click **Finish**.
- 3 Type an ODBC datastore name (DSN) in the **Name** text box.
For example, VMware vCenter Server.
- 4 (Optional) Type an ODBC DSN description in the **Description** text box.

- 5 Select the server name from the **Server** drop-down menu.
Type the SQL Server host name in the text box if it is not in the drop-down menu.
- 6 Select one of the authentication methods.
 - **Integrate Windows authentication.** Optionally, enter the Service Principal Name (SPN).
 - **SQL Server authentication.** Type 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 For SQL Server 2005 and SQL Server 2008 editions, 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 local.

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**.

What to do next

Optionally, you can enable Database Monitoring for Microsoft SQL database users. Otherwise, install vCenter Server.

(Optional) Configure a Microsoft SQL Server Database User to Enable Database Monitoring

vCenter Server Database Monitoring captures metrics that enable the administrator to assess the status and health of the database server. Enabling Database Monitoring helps the administrator prevent vCenter downtime because of a lack of resources for the database server.

Database Monitoring for vCenter Server enables administrators to monitor the database server CPU, memory, I/O, data storage, and other environment factors for stress conditions. Statistics are stored in the vCenter Server Profile Logs.

You can enable Database Monitoring for a user before or after you install vCenter Server. You can also perform this procedure while vCenter Server is running.

Procedure

- 1 Log in to a SQL Server Management Studio session as the sysadmin (SA) or to a user account with sysadmin privileges.
- 2 Run the following SQL commands to grant additional permissions to vCenter Server database login:

```
use master
go
grant VIEW SERVER STATE to user
go
```

vCenter Database Monitoring is enabled.

Configure Oracle Databases

To use an Oracle database for your vCenter Server repository, configure your database to work with vCenter Server.

Procedure

- 1 [Configure an Oracle Database User](#) on page 43
To use an Oracle database when you install vCenter Server, you must configure the database user.
- 2 [Use a Script to Create a Local or Remote Oracle Database](#) on page 44
When you use an Oracle database with vCenter Server, the database must have certain table spaces and privileges. To simplify the process of creating the database, you can run a script. You also can create the database manually.
- 3 [\(Optional\) Use a Script to Create the Oracle Database Schema](#) on page 44
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.
- 4 [Configure an Oracle Connection for Local Access](#) on page 46
Configure a connection for local access if you install vCenter Server on the same system as the Oracle database.
- 5 [Configure an Oracle Database Connection for Remote Access](#) on page 46
Before a vCenter Server system can access the Oracle database remotely, you must configure an Oracle connection.
- 6 [Connect to an Oracle Database Locally](#) on page 47
Before a vCenter Server system can connect to an Oracle database locally, you must set up the connection.

- 7 (Optional) [Configure an Oracle Database User to Enable Database Monitoring](#) on page 48
- vCenter Server Database Monitoring captures metrics that enable the administrator to assess the status and health of the database server. Enabling Database Monitoring helps the administrator prevent vCenter downtime because of a lack of resources for the database server.

Configure an Oracle Database User

To use an Oracle database when you install vCenter Server, you must configure the database user.

You can configure an Oracle database for vCenter Server either locally on the same Microsoft Windows machine as vCenter Server or remotely on a network-connected Linux, UNIX or Microsoft Windows host.

Prerequisites

Review the software requirements for vCenter Server with Oracle.

Procedure

- 1 Log in to a SQL*Plus session with the system account.
- 2 Run the following SQL command to create a vCenter Server database user with the correct permissions.

The script is located in the vCenter Server installation package */installation directory/vCenter-Server/dbschema/DB_and_schema_creation_scripts_oracle.txt* file.

In this example, the user name is VPXADMIN.

```
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_tablespaces to VPXADMIN;
grant select on dba_temp_files to VPXADMIN;
grant select on dba_data_files 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>'
```

- 3 (Optional) After you have successfully installed vCenter Server with the Oracle database, you can revoke the following privileges.

```
revoke select on dba_tablespaces from VPXADMIN;
revoke select on dba_temp_files from VPXADMIN;
revoke select on dba_data_files from VPXADMIN;
```

You now have an Oracle database user that you can reference in the vCenter Server installer.

What to do next

Create the Oracle database, including all necessary table spaces and privileges.

Use a Script to Create a Local or Remote Oracle Database

When you use an Oracle database with vCenter Server, the database must have certain table spaces and privileges. To simplify the process of creating the database, you can run a script. You also can create the database manually.

When using the script, you can customize the location of the data and log files. The user created by this script does not follow any security policy. The passwords are provided only for convenience. Change the passwords as appropriate.

Procedure

- 1 Log in to a SQL*Plus session with the system account.
- 2 Run the following script.

The script is located in the vCenter Server installation package */installation directory/vCenter-Server/dbschema/DB_and_schema_creation_scripts_oracle.txt* file.

```
CREATE SMALLFILE TABLESPACE "VPX" DATAFILE '/u01/app/oracle/oradata/vcdb/vpx01.dbf'
SIZE 1G AUTOEXTEND ON NEXT 10M MAXSIZE UNLIMITED LOGGING EXTENT MANAGEMENT LOCAL SEGMENT
SPACE MANAGEMENT AUTO;
```

For a Windows installation, change the directory path to the vpx01.dbf file.

You now have an Oracle database that you can use with vCenter Server.

What to do next

You can run a script to create the database schema.

(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.

To have the vCenter Server installer create your schema for you, see [“Configure an Oracle Connection for Local Access,”](#) on page 46 or [“Configure an Oracle Database Connection for Remote Access,”](#) on page 46, depending on your environment.

Prerequisites

Create the Oracle database and user. You can create the Oracle database and user manually or by using scripts.

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 sequence on the database.

path is the directory path to the */installation directory/vCenter-Server/dbschema* folder.

```
@path/VCDB_oracle.SQL
@path/VCDB_views_oracle.SQL
@path/insert_stats_proc_oracle.sql
@path/load_stats_proc_oracle.sql
@path/purge_stat2_proc_oracle.sql
@path/purge_stat3_proc_oracle.sql
@path/purge_usage_stats_proc_oracle.sql
@path/stats_rollup1_proc_oracle.sql
@path/stats_rollup2_proc_oracle.sql
@path/stats_rollup3_proc_oracle.sql
@path/cleanup_events_oracle.sql
@path/delete_stats_proc_oracle.sql
@path/load_usage_stats_proc_oracle.sql
@path/TopN_DB_oracle.sql
@path/calc_topn1_proc_oracle.sql
@path/calc_topn2_proc_oracle.sql
@path/calc_topn3_proc_oracle.sql
@path/calc_topn4_proc_oracle.sql
@path/clear_topn1_proc_oracle.sql
@path/clear_topn2_proc_oracle.sql
@path/clear_topn3_proc_oracle.sql
@path/clear_topn4_proc_oracle.sql
@path/rule_topn1_proc_oracle.sql
@path/rule_topn2_proc_oracle.sql
@path/rule_topn3_proc_oracle.sql
@path/rule_topn4_proc_oracle.sql
@path/process_license_snapshot_oracle.sql
@path/l_purge_stat2_proc_oracle.sql
@path/l_purge_stat3_proc_oracle.sql
@path/l_stats_rollup1_proc_oracle.sql
@path/l_stats_rollup2_proc_oracle.sql
@path/l_stats_rollup3_proc_oracle.sql
```

- 4 (Optional) You can also run the following scripts to enable database health monitoring.

```
job_dbm_performance_data_oracle.sql
process_performance_data_oracle.sql
```

- 5 For all supported editions of Oracle Server, run these scripts to set up scheduled jobs on the database.

```
@path/job_schedule1_oracle.sql
@path/job_schedule2_oracle.sql
@path/job_schedule3_oracle.sql
@path/job_cleanup_events_oracle.sql
@path/job_topn_past_day_oracle.sql
@path/job_topn_past_week_oracle.sql
@path/job_topn_past_month_oracle.sql
@path/job_topn_past_year_oracle.sql
@path/job_property_bulletin_oracle.sql
```

You now have a database schema that is compatible with vCenter Server.

- 6 On the machine that you are installing vCenter Server on, create a DSN that points to the database server that has the schema.
- 7 Run the vCenter Server installer.

- 8 If a database reinitialization warning message appears in the vCenter Server installer, select **Do not overwrite, leave my existing database in place** and continue the installation.

This message appears if you are using a database that has vCenter Server tables that were created by a previous installation. The message does not appear if the database is clean.

If you leave your existing database in place, you cannot join a Linked Mode group during the installation. You can join after the installation is complete. See [“Join a Linked Mode Group After Installation,”](#) on page 112.

- 9 When prompted, provide the database user login.

The Oracle database schema is created.

Configure an Oracle Connection for Local Access

Configure a connection for local access if you install vCenter Server on the same system as the Oracle database.

Prerequisites

Review the required database patches specified in [“vCenter Server Database Configuration Notes,”](#) on page 30. If you do not prepare your database correctly, the vCenter Server installer displays error and warning messages.

Procedure

- 1 Download Oracle 10g or Oracle 11g from the Oracle Web site.
- 2 Install Oracle 10g or Oracle 11g, and create a database.
- 3 Configure the TNS Service Name option in the ODBC DSN.

The TNS Service Name is the net service name for the database to which you want to connect. You can find the net service name in the `tnsnames.ora` file located in the `NETWORK\ADMIN` folder in the Oracle database installation location.

The database is configured for local access.

Configure an Oracle Database Connection for Remote Access

Before a vCenter Server system can access the Oracle database remotely, you must configure an Oracle connection.

Prerequisites

Review the required database patches specified in [“vCenter Server Database Configuration Notes,”](#) on page 30. If you do not prepare your database correctly, the vCenter Server installer displays error and warning messages.

Procedure

- 1 Install the Oracle client on the vCenter Server system machine.
- 2 Download and install the ODBC driver.
- 3 Create a new tablespace for a vCenter Server system using a SQL statement such as the following statement.

```
CREATE TABLESPACE "VPX" DATAFILE 'C:\Oracle\ORADATA\VPX\VPX.dat' SIZE 1000M AUTOEXTEND ON
NEXT 500K;
```

- 4 Create a user, such as `vpxAdmin`, for accessing the tablespace through ODBC.

```
CREATE USER vpxAdmin IDENTIFIED BY vpxadmin DEFAULT TABLESPACE vpx;
```

- 5 Grant permissions to the user, in one of the following ways.

- Grant **dba** permission to the user.
- Grant the following permissions to the user.

```
grant connect to user
grant resource to user
grant create view to user
grant create sequence to user
grant create table to user
grant create materialized view to user
grant execute on dbms_lock to user
grant execute on dbms_job to user
grant unlimited tablespace to user # To ensure space is sufficient
```

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.

- 6 Use a text editor or the Net8 Configuration Assistant to edit the `tnsnames.ora` file located in the directory `C:\Oracle\Oraxx\NETWORK\ADMIN`, where `xx` is either `10g` or `11g`.

Add the following entry, where `HOST` is the managed host to which the client must connect.

```
VPX =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS=(PROTOCOL=TCP)(HOST=vpzd-Oracle)(PORT=1521))
)
(CONNECT_DATA =
(SERVICE_NAME = VPX)
)
)
```

- 7 Configure the TNS Service Name option in the ODBC DSN.

The TNS Service Name is the net service name for the database to which you want to connect, in this case, `VPX`. You can find the net service name in the `tnsnames.ora` file.

Connect to an Oracle Database Locally

Before a vCenter Server system can connect to an Oracle database locally, you must set up the connection.

Procedure

- 1 Create a new tablespace for a vCenter Server system using a SQL statement such as the following statement.

```
CREATE TABLESPACE "VPX" DATAFILE 'C:\Oracle\ORADATA\VPX\VPX.dat' SIZE 1000M AUTOEXTEND ON
NEXT 500K;
```

- 2 Create a user, such as `vpAdmin`, for accessing the tablespace through ODBC.

```
CREATE USER vpAdmin IDENTIFIED BY vpadmin DEFAULT TABLESPACE vpx;
```

- 3 Grant permissions to the user, in one of the following ways.

- Grant **dba** permission to the user.
- Grant the following permissions to the user.

```
grant connect to user
grant resource to user
grant create view to user
grant create sequence to user
grant create table to user
grant create materialized view to user
grant execute on dbms_lock to user
grant execute on dbms_job to user
grant unlimited tablespace to user # To ensure space is sufficient
```

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.

- 4 Create an ODBC connection to the database.

The following code shows example settings.

```
Data Source Name: VMware vCenter Server TNS Service Name: VPX User Id: vpxAdmin
```

You now have a database that you can connect to locally.

What to do next

Optionally, you can enable Database Monitoring for Oracle database users. Otherwise, install vCenter Server.

(Optional) Configure an Oracle Database User to Enable Database Monitoring

vCenter Server Database Monitoring captures metrics that enable the administrator to assess the status and health of the database server. Enabling Database Monitoring helps the administrator prevent vCenter downtime because of a lack of resources for the database server.

Database Monitoring for vCenter Server enables administrators to monitor the database server CPU, memory, I/O, data storage, and other environment factors for stress conditions. Statistics are stored in the vCenter Server Profile Logs.

Enable Database Monitoring for a user before or after you install vCenter Server. You can perform this procedure while vCenter Server is running.

Procedure

- 1 Log in to a SQL*Plus session with the system account.
- 2 Run the following SQL commands to grant additional permissions to the vCenter Server database user:

```
grant select on v_$system_event to user;
grant select on v_$sysmetric_history to user;
grant select on v_$sysstat to user;
grant select on dba_data_files to user;
grant select on v_$loghist to user;
```

vCenter Database Monitoring is enabled.

Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server

Before installing vCenter Single Sign-On, Inventory Service, and vCenter Server, review the prerequisites.

Prerequisites for Understanding and Preparing for the Installation Process

- vCenter Server versions 5.1 and later require vCenter Single Sign-On and Inventory Service. You must install these components in this order: vCenter Single Sign-On, the vSphere Web Client, Inventory Service, and vCenter Server. Review the topics in the section [“How vCenter Single Sign-On Affects vCenter Server Installation,”](#) on page 51
- Review the release notes for known issues or special installation notes.
- Gather the information that the vCenter Single Sign-On, Inventory Service, and vCenter Server installation wizards require. See [“Required Information for Installing or Upgrading vCenter Single Sign-On, Inventory Service, vCenter Server, and the vSphere Web Client,”](#) on page 62.
- Decide whether the vCenter Server instance will be a standalone instance or in a Linked Mode group. See [“Creating vCenter Server Linked Mode Groups,”](#) on page 109.
- Download the vCenter Server installer from the VMware Web site.

System Prerequisites

- Verify that your system meets the requirements listed in [“Hardware Requirements for vCenter Server, the vSphere Web Client, vCenter Inventory Service, and vCenter Single Sign-On,”](#) on page 17 and [“vCenter Server Software Requirements,”](#) on page 22, and that the required ports are open, as discussed in [“Required Ports for vCenter Server,”](#) on page 23.
- Before you install or upgrade any vSphere product, synchronize the clocks of all machines on the vSphere network. See [“Synchronizing Clocks on the vSphere Network,”](#) on page 59.
- Verify that the DNS name of the vCenter Server host machine matches the actual computer name.
- Verify that the host name of the machine that you are installing vCenter Server on complies with RFC 952 guidelines.
- The installation path of vCenter Server must be compatible with the installation requirements for Microsoft Active Directory Application Mode (ADAM/AD LDS). The installation path cannot contain any of the following characters: non-ASCII characters, commas (,), periods (.), exclamation points (!), pound signs (#), at signs (@), or percentage signs (%).
- Verify that the system on which you are installing vCenter Server is not an Active Directory domain controller.
- On each system that is running vCenter Server, verify that the domain user account has the following permissions:
 - **Member of the Administrators group**
 - **Act as part of the operating system**
 - **Log on as a service**
- vCenter Server requires the Microsoft .NET 3.5 SP1 Framework. If your system does not have it installed, the vCenter Server installer installs it. The .NET 3.5 SP1 installation might require Internet connectivity to download more files.

- 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 machine must be connected to a domain if you want to add Active Directory identity sources after the installation. To determine whether the system belongs to a workgroup or a domain, right-click **My Computer**. Click **Properties** and click the **Computer Name** tab. The **Computer Name** tab displays either a Workgroup label or a Domain label.
- Verify that the NETWORK SERVICE account has read permission on the folder in which vCenter Server is installed and on the HKLM registry.
- During the installation, verify that the connection between the machine and the domain controller is working.
- Before the vCenter Server installation, in the Administrative Tools control panel of the vCenter Single Sign-On instance that you will register vCenter Server to, verify that the following services are started: VMware Certificate Service, VMware Directory service, VMware Identity Manager Service, VMware KDC service, and tcruntime-C-ProgramData-VMware-cis-runtime-VMwareSTSService.
- You must log in as a member of the Administrators group on the host machine, with a user name that does not contain any non-ASCII characters.

Network Prerequisites

- Verify that the fully qualified domain name (FQDN) of the system where you will install vCenter Server is resolvable. To check that the FQDN is resolvable, type **nslookup your_vCenter_Server_fqdn** at a command line prompt. If the FQDN is resolvable, the **nslookup** command returns the IP and name of the domain controller machine.
- Verify that DNS reverse lookup returns a fully qualified domain name when queried with the IP address of the vCenter Server. When you install vCenter Server, 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 of the vCenter Server from its IP address. Reverse lookup is implemented using PTR records. To create a PTR record, see the documentation for your vCenter Server host operating system.
- Verify that no Network Address Translation (NAT) exists between the vCenter Server system and the hosts it will manage.
- Install vCenter Server, like any other network server, on a 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. Ensure that the ESXi host management interface has a valid DNS resolution from the vCenter Server and all vSphere Web Clients. Ensure that the vCenter Server has a valid DNS resolution from all ESXi hosts and all vSphere Web Clients. 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). Ping the computer name to test this connection. For example, if the computer name is *host-1.company.com*, run the following command in the Windows command prompt:

```
ping host-1.company.com
```

If you can ping the computer name, the name is updated in DNS.

- If you will use Active Directory as an identity source, verify that it is set up correctly. The DNS of the vCenter Single Sign-On Server host machine must contain both lookup and reverse lookup entries for the domain controller of the Active Directory. For example, pinging *mycompany.com* should return the domain controller IP address for *mycompany*. Similarly, the **ping -a** command for that IP address should return the domain controller hostname. Avoid trying to correct name resolution issues by editing the hosts file. Instead, make sure that the DNS server is correctly set up. For more information about configuring Active Directory, see the Microsoft Web site.

Database Prerequisites

- Verify that your vCenter Server database meets the database requirements. See [“vCenter Server Database Configuration Notes,”](#) on page 30 and [“Preparing vCenter Server Databases,”](#) on page 30.
- Create a vCenter Server database, unless you plan to install the bundled database.

How vCenter Single Sign-On Affects vCenter Server Installation

Starting with version 5.1, vSphere includes a vCenter Single Sign-On component as part of the vCenter Server management infrastructure. This change affects vCenter Server installation.

Authentication by vCenter Single Sign-On makes the VMware cloud infrastructure platform more secure by allowing the vSphere software components to communicate with each other through a secure token exchange mechanism.

For information about configuring vCenter Single Sign-On, see *vSphere Security*. For more information about vCenter Single Sign-On deployment modes, see [“vCenter Single Sign-On Deployment Modes,”](#) on page 52.

For the first installation of vCenter Server, you must install all components. In subsequent installations in the same environment, or if you add services, you do not have to install vCenter Single Sign-On. One vCenter Single Sign-On server can serve your entire vSphere environment. After you install vCenter Single Sign-On once, you can connect all new vCenter Server instances to the same vCenter Single Sign-On service. You must install an Inventory Service instance for each vCenter Server instance.

Simple Install

The Simple Install option installs vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server on the same host or virtual machine. Simple Install is appropriate for most deployments.

Custom Install

If you want to customize the location and setup of each component, you can install the components separately by performing a custom install and selecting the individual installation options, in the following order:

- 1 vCenter Single Sign-On
- 2 vSphere Web Client
- 3 vCenter Inventory Service
- 4 vCenter Server

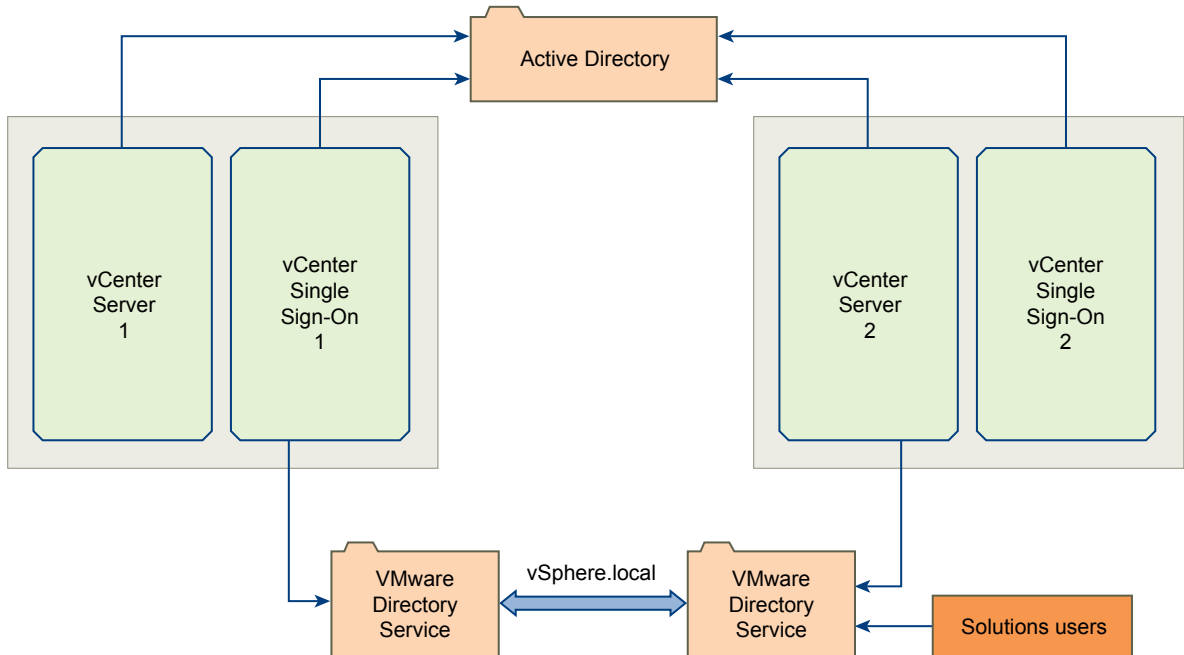
You can install each component on a different host or virtual machine.

If you decide on installing multiple vCenter Server systems, you can point to the same vCenter Single Sign-On service for each vCenter Server.

Installing in Multiple Locations

Unlike vCenter Single Sign-On version 5.1, vCenter Single Sign-On 5.5 synchronizes authentication data across locations.

If you install vCenter Server systems in multiple locations, you can install a vCenter Single Sign-On server in each location. When you install the second and subsequent instances of vCenter Single Sign-On, you can point it to the first vCenter Single Sign-On instance during installation. The two instances synchronize their VMware Directory Service instances. Changes to one instance are propagated to the other instance.

Figure 3-1. Installing vCenter Single Sign-On in Multiple Locations

vCenter Single Sign-On Deployment Modes

vCenter Server provides several ways to deploy vCenter Single Sign-On to best serve your vSphere environment

You can deploy vCenter Single Sign-On in one of three modes.

To choose the right mode for your environment, consider the way you use vCenter Server.

Table 3-2. Choosing a vCenter Single Sign-On Deployment Mode

vCenter Server Deployment	Single Sign-On Deployment Mode
Single vCenter Server	Basic vCenter Single Sign-On
Multiple local vCenter Servers	Basic vCenter Single Sign-On
Multiple remote vCenter Servers	Basic vCenter Single Sign-On
Multiple vCenter Servers in Linked Mode	Multisite vCenter Single Sign-On
vCenter Servers with high availability	<p>Basic vCenter Single Sign-On with VMware vSphere HA (provides high availability for vCenter Server and vCenter Single Sign-On)</p> <p>Basic vCenter Single Sign-On with vCenter Server Heartbeat (provides high availability for vCenter Server and vCenter Single Sign-On)</p> <p>See “vCenter Single Sign-On and High Availability,” on page 54.</p>

Basic

Basic vCenter Single Sign-On is the most common deployment mode, and meets the requirements of most vSphere 5.1 and 5.5 users. Typically, this deployment mode maintains the same architecture as previous vCenter Server environments. In most cases, you can use vCenter Simple Install to deploy vCenter Server with vCenter Single Sign-On in basic mode.

In Basic deployment mode, a single standalone instance of the vCenter Single Sign-On server supports the connectivity of Active Directory, OpenLDAP, Local Operating System, and vCenter Single Sign-On embedded users and groups. In most cases, the vCenter Single Sign-On instance is installed on the same host machine as vCenter Server, as with the vCenter Server Simple Install option, or the vCenter Server Appliance.

The Basic vCenter Single Sign-On deployment is appropriate in the following circumstances:

- If you have a single vCenter Server of any supported inventory size: up to 1,000 hosts or 10,000 virtual machines.
- If you have multiple geographically dispersed locations, each with a local vCenter Server and you do not require a single-pane-of-glass view as provided by vCenter Linked Mode.

Multiple Single Sign-On instances in the same location

For this deployment mode, you install a vCenter Single Sign-On primary instance and one or more additional vCenter Single Sign-On nodes. Both the primary and high availability instances are placed behind a third-party network load balancer (for example, Apache HTTPD or vCNS). Each vCenter Single Sign-On has its own VMware Directory Service that replicates information with other vCenter Single Sign-On servers. vCenter Single Sign-On administrator users, when connected to vCenter Server through the vSphere Web Client, will see the primary vCenter Single Sign-On instance.

This deployment mode has the following limitations:

- It provides provides failover only for the vCenter Single Sign-On service. It does not provide failover for the vCenter Single Sign-On host machine.
- It supports the connectivity of Active Directory, OpenLDAP and vCenter Single Sign-On embedded users and groups, but does not support the use of local operating system user accounts.

See [“vCenter Single Sign-On and High Availability,”](#) on page 54 for high availability options.

Multiple Single Sign-On instances in different locations

This mode is designed for vCenter Server deployments with multiple physical locations. Multisite deployment is required when a single administrator needs to administer vCenter Server instances that are deployed on geographically dispersed sites in Linked Mode.

Each site is represented by one vCenter Single Sign-On instance, with one vCenter Single Sign-On server, or a high-availability cluster. The vCenter Single Sign-On site entry point is the machine that other sites communicate with. This is the only machine that needs to be visible from the other sites. In a clustered deployment, the entry point of the site is the machine where the load balancer is installed.

NOTE This deployment mode is required if you have geographically dispersed vCenter Servers in Linked Mode. You might also consider this mode in the following cases:

- If multiple vCenter Servers require the ability to communicate with each other.
 - If you require one vCenter Single Sign-On server security domain for your organization.
-

This deployment mode has the following limitations:

- It supports the connectivity of Active Directory, OpenLDAP and vCenter Single Sign-On embedded users and groups, but does not support the use of local operating system user accounts.
- Secondary vCenter Single Sign-On instances must belong to the same Active Directory or OpenLDAP domain as the primary vCenter Single Sign-On server and must have a local domain controller available.

You can install the vCenter Single Sign-On nodes in this deployment in any order. Any node that is installed after the first node can point to any node that is already installed. For example, the third node can point to either the first or second node.

vCenter Single Sign-On and High Availability

vSphere provides several ways to ensure availability of your vSphere deployment with vCenter Single Sign-On.

vCenter Single Sign-On is merely an authentication component for vCenter Server. Single Sign-On protection does not provide any benefit without vCenter Server protection. Protecting one without the other does not provide an effective availability solution. The solution you choose to protect vCenter Server will provide the same protection for vCenter Single Sign-On without the additional complexity caused by including third-party technologies.

Options for Protecting vCenter Single Sign-On and vCenter Server

The following options vary in the level of protection afforded, and in the recovery time required.

Backup and restore	Backup and restore should be an essential part of any availability solution, providing a granular recovery method, by tape, disk, or snapshot. However, the recovery time is typically measured in hours or days and requires manual intervention. Any backup solution must be independent of vCenter Server. Solutions like VMware Data Protection require an operational vCenter Server with a functioning vCenter Single Sign-On server to restore a virtual machine.
vSphere HA	vSphere HA is an industry standard for maintaining uptime of virtual machines and for detection of ESXi host failure. Also, with vSphere HA, a failed response to a configured VMware Tools heartbeat automatically reboots the virtual machine onto another operational host within the vSphere cluster. This detection usually occurs within seconds. A virtual machine can be fully rebooted within minutes, providing redundancy for vSphere host failures and virtual machine operating system crashes. vSphere HA does not have any knowledge of the application running inside the virtual machine.
vCenter Server Heartbeat	This separately licensed vCenter Server plug-in provides vCenter Server protection (physical or virtual) and can protect against failure of hosts. vCenter Server Heartbeat also adds application-level monitoring and intelligence of all vCenter Server components. vCenter Server Heartbeat is installed directly onto the vCenter Server or vCenter Server component, and replicates changes to a cloned virtual machine. The cloned virtual machine can take over when a failure event is triggered. The recovery can be accomplished by restarting the component, by restarting the entire application, or by the entire failover of the component or application to one or more paired virtual machines. Recovery time is measured in minutes.

vCenter Single Sign-On Deployment Modes and High Availability

To determine the best deployment mode for vCenter Single Sign-On availability, consider the environment that vCenter Single Sign-On will serve.

Single vCenter Server with local vCenter Single Sign-On in Basic deployment mode	In the simplest deployment of vCenter Single Sign-On for high availability, you install vCenter Single Sign-On in Basic deployment mode, local to vCenter Server, and then add the availability solution. If the single machine that hosts vCenter Server and vCenter Single Sign-On is virtual, you can place it in a vSphere HA-enabled cluster and protect it with no further configuration. If you require protection at the application level, you can use vCenter Server Heartbeat. If vCenter Server and vCenter Single Sign-On are hosted on a physical server, vCenter Server Heartbeat is the only solution for availability.
Multiple vCenter Servers in a single location	In this environment, a dedicated, standalone vCenter Single Sign-On instance serves multiple vCenter Server instances in one physical location. If vCenter Single Sign-On is hosted on a virtual machine, you can place the standalone vCenter Single Sign-On server in a vSphere HA-enabled cluster and protect vCenter Single Sign-On with no further configuration. If you require application-level protection, you can use vCenter Server Heartbeat.

vCenter Server Heartbeat is the only solution for availability if vCenter Single Sign-On is on a physical server. With either vSphere HA or vCenter Server Heartbeat, this deployment provides complete protection of the centralized vCenter Single Sign-On environment.

Geographically dispersed vCenter Servers

If your vSphere deployment includes vCenter Servers in different locations, it is not advisable to use a remote centralized vCenter Single Sign-On environment for vCenter Server authentication. Instead, you can provide one or more vCenter Single Sign-On instances at each location. Depending on the deployment of vCenter Servers at each location, you can use one of the same availability strategies described above in the options "Single vCenter Server with local vCenter Single Sign-On in Basic deployment mode" and "Multiple vCenter Servers in a single location with one vCenter Single Sign-On server."

vCenter Single Sign-On Components

vCenter Single Sign-On includes the Security Token Service (STS), an administration server, and vCenter Lookup Service, as well as the VMware Directory Service (vmdir).

The components are deployed as part of installation.

STS (Security Token Service)

STS certificates enable a user who has logged on through vCenter Single Sign-On to use any vCenter service that vCenter Single Sign-On supports without authenticating to each one. The STS service issues Security Assertion Markup Language (SAML) tokens. These security tokens represent the identity of a user in one of the identity source types supported by vCenter Single Sign-On.

Administration server

The administration server allows users with administrator privileges to vCenter Single Sign-On to configure the vCenter Single Sign-On server and manage users and groups from the vSphere Web Client. Initially, only the user administrator@vsphere.local has these privileges.

vCenter Lookup Service

vCenter Lookup Service contains topology information about the vSphere infrastructure, enabling vSphere components to connect to each other securely. Unless you are using Simple Install, you are prompted for the Lookup Service URL when you install other vSphere components. For example, the Inventory Service and the vCenter Server installers ask for the Lookup Service URL and then contact the Lookup Service to find vCenter Single Sign-On. After installation, the Inventory Service and vCenter Server system are registered in vCenter Lookup Service so other vSphere components, like the vSphere Web Client, can find them.

VMware Directory Service

Directory service associated with the vsphere.local domain. This service is a multi-tenanted, multi-mastered directory service that makes an LDAP directory available on port 11711. In multisite mode, an update of VMware Directory Service content in one VMware Directory Service instance results in the automatic update of the VMware Directory Service instances associated with all other vCenter Single Sign-On nodes.

Setting the vCenter Server Administrator User

The way you set the vCenter Server administrator user depends on your vCenter Single Sign On deployment.

In vSphere versions before vSphere 5.1, vCenter Server administrators are the users that belong to the local operating system administrators group.

In vSphere 5.1.x and 5.5, when you install vCenter Server, you must provide the default (initial) vCenter Server administrator user or group. For deployments where vCenter Server and vCenter Single Sign-On are on the same host machine, you can designate the local operating system group Administrators as vCenter Server administrative users. This option is the default. This behavior is unchanged from vCenter Server 5.0.

For larger installations, where vCenter Single Sign-On and vCenter Server are deployed on different hosts, you cannot preserve the same behavior as in vCenter Server 5.0. Instead, assign the vCenter Server administrator role to a user or group from an identity source that is registered in the vCenter Single Sign-On server: Active Directory, OpenLDAP, or the system identity source.

Authenticating to the vCenter Server Environment

In vCenter Server versions 5.1 and later, users authenticate through vCenter Single Sign-On.

In vCenter Server versions earlier than vCenter Server 5.1, when a user connects to vCenter Server, vCenter Server authenticates the user by validating the user against an Active Directory domain or the list of local operating system users.

The user administrator@vsphere.local has vCenter Single Sign-On administrator privileges by default. When logged in to the vCenter Single Sign-On server from the vSphere Web Client, the administrator@vsphere.local user can assign vCenter Single Sign-On administrator privileges to other users. These users might be different from the users that administer vCenter Server.

Users can log in to vCenter Server with the vSphere Web Client. Users authenticate to vCenter Single Sign-On. Users can view all the vCenter Server instances that the user has permissions on. After users connect to vCenter Server, no further authentication is required. The actions users can perform on objects depend on the user's vCenter Server permissions on those objects.

For more information about vCenter Single Sign-On, see *vSphere Security*.

How vCenter Single Sign-On Affects Log In Behavior

vCenter Single Sign-On log in behavior depends on the domain the user belongs to and the identity sources that you have added to vCenter Single Sign-On.

When a user logs in to a vCenter Server system from the vSphere Web Client, the login behavior depends on whether the user is in the default domain.

- Users who are in the default domain can log in with their user name and password.
- Users who are in a domain that has been added to vCenter Single Sign-On as an identity source but is not the default domain can log in to vCenter Server but must specify the domain in one of the following ways.
 - Including a domain name prefix, for example, MYDOMAIN\user1
 - Including the domain, for example, user1@mydomain.com
- Users who are in a domain that is not a vCenter Single Sign-On identity source cannot log in to vCenter Server. If the domain that you add to vCenter Single Sign-On is part of a domain hierarchy, Active Directory determines whether users of other domains in the hierarchy are authenticated or not.

After installation on a Windows system, the user administrator@vsphere.local has administrator privileges to both the vCenter Single Sign-On server and to the vCenter Server system.

After you deploy the vCenter Virtual Appliance, the user administrator@vsphere.local has administrator privileges to both the vCenter Single Sign-On server and to the vCenter Server system. The user root@localos has administrative privileges on the vCenter Single Sign-On server and can authenticate to the vCenter Server system. Assign permissions to root@localos to allow that user access to the vCenter Server system.

Identity Sources for vCenter Server with vCenter Single Sign-On

Identity sources allow you to attach one or more domains to vCenter Single Sign-On. A domain is a repository for users and groups that the vCenter Single Sign-On server can use for user authentication.

An identity source is a collection of user and group data. The user and group data is stored in Active Directory, OpenLDAP, or locally to the operating system of the machine where vCenter Single Sign-On is installed. Upon installation, every instance of vCenter Single Sign-On has the Local OS identity source identity source `vpshere.local`. This identity source is internal to vCenter Single Sign-On.

A vCenter Single Sign-On administrator user can create vCenter Single Sign-On users and groups.

Types of Identity Sources

vCenter Server versions earlier than version 5.1 supported Active Directory and local operating system users as user repositories. As a result, local operating system users could always authenticate to the vCenter Server system. vCenter Server version 5.1 and version 5.5 uses vCenter Single Sign-On for authentication. See the vSphere 5.1 documentation for a list of supported identity sources with vCenter Single Sign-On 5.1. vCenter Single Sign-On 5.5 supports the following types of user repositories as identity sources, but supports only one default identity source.

- Active Directory versions 2003 and later. vCenter Single Sign-On allows you to specify a single Active Directory domain as an identity source. The domain can have child domains or be a forest root domain. Shown as **Active Directory (Integrated Windows Authentication)** in the vSphere Web Client.
- Active Directory over LDAP. vCenter Single Sign-On supports multiple Active Directory over LDAP identity sources. This identity source type is included for compatibility with the vCenter Single Sign-On service included with vSphere 5.1. Shown as **Active Directory as an LDAP Server** in the vSphere Web Client.
- OpenLDAP versions 2.4 and later. vCenter Single Sign-On supports multiple OpenLDAP identity sources. Shown as **OpenLDAP** in the vSphere Web Client.
- Local operating system users. Local operating system users are local to the operating system where the vCenter Single Sign-On server is running. The local operating system identity source exists only in basic vCenter Single Sign-On server deployments and is not available in deployments with multiple vCenter Single Sign-On instances. Only one local operating system identity source is allowed. Shown as **localos** in the vSphere Web Client.
- vCenter Single Sign-On system users. Exactly one system identity source named `vpshere.local` is created when you install vCenter Single Sign-On. Shown as **vpshere.local** in the vSphere Web Client.

NOTE At any time, only one default domain exists. If a user from a non-default domain logs in, that user must add the domain name (`DOMAIN\user`) to authenticate successfully.

vCenter Single Sign-On identity sources are managed by vCenter Single Sign-On administrator users.

You can add identity sources to a vCenter Single Sign-On server instance. Remote identity sources are limited to Active Directory and OpenLDAP server implementations.

For more information about vCenter Single Sign-On, see *vSphere Security*.

Login Behavior

When a user logs in to a vCenter Server system from the vSphere Web Client, the login behavior depends on whether the user is in the default domain.

- Users who are in the default domain can log in with their user name and password.

- Users who are in a domain that has been added to vCenter Single Sign-On as an identity source but is not the default domain can log in to vCenter Server but must specify the domain in one of the following ways.
 - Including a domain name prefix, for example, MYDOMAIN\user1
 - Including the domain, for example, user1@mydomain.com
- Users who are in a domain that is not a vCenter Single Sign-On identity source cannot log in to vCenter Server. If the domain that you add to vCenter Single Sign-On is part of a domain hierarchy, Active Directory determines whether users of other domains in the hierarchy are authenticated or not.

vCenter Single Sign-On does not propagate permissions that result from nested groups from dissimilar identity sources. For example, if you add the Domain Administrators group to the Local Administrators group, the permissions is not propagated because Local OS and Active Directory are separate identity sources.

Synchronizing Clocks on the vSphere Network

Before you install vCenter Single Sign-On, install the vSphere Web Client, or deploy the vCenter Server Appliance, make sure that all machines on the vSphere network have their clocks synchronized.

If the clocks on vCenter Server network machines 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 vSphere Web Client installation to fail or prevent the vCenter Server Appliance vpxd service from starting.

Make sure that any Windows host on which a vCenter component runs is synchronized with the NTP server. See the Knowledge Base article [Timekeeping best practices for Windows, including NTP](#).

Synchronize ESX and ESXi Clocks with a Network Time Server

Before you install vCenter Single Sign-On, the vSphere Web Client, or the vCenter Server appliance, make sure all machines on the vSphere network have their clocks synchronized.

Procedure

- 1 From the vSphere Web Client, connect to the vCenter Server.
- 2 Select the host in the inventory.
- 3 Select the **Manage** tab.
- 4 Select **Settings**.
- 5 In the System section, select **Time Configuration**.
- 6 Click **Edit** and set up the NTP server.
 - a Select **Use Network Time Protocol (Enable NTP client)**.
 - b Set the NTP Service Startup Policy.
 - c Enter the IP addresses of the NTP servers to synchronize with.
 - d Click **Start** or **Restart** in the NTP Service Status section.
- 7 Click **OK**.

The host synchronizes with the NTP server.

Synchronize the vCenter Server Appliance Clock with an NTP Server

Before you deploy the vCenter Server Appliance, make sure all machines on the network have their clocks synchronized. Unsynchronized clocks can cause installation and authentication errors.

On systems that are joined to a Windows domain, the vCenter Server Appliance clock is synchronized automatically with the domain controller. On other systems, you can enable synchronizing the clock through VMware Tools. As an alternative, you can use this procedure.

Procedure

- 1 Open a Web browser and navigate to the vCenter Server Appliance Management Interface (<https://vCenter-Appliance-Address:5480/>).
- 2 Log in as root.
- 3 From the vCenter Server tab, select the Time subtab.
- 4 Select one or more of the available options.

Option	Description
No synchronization	Does not perform synchronization.
NTP synchronization	Select this option and specify one or more NTP servers to configure the appliance to synchronize with an NTP server directly.
VMware Tools synchronization	Select this option to synchronize all virtual machines.
Active Directory synchronization	This option becomes available only if you add the appliance to an Active Directory domain. If you select this option, none of the other options is available.

- 5 Click **Save Settings**.

The vCenter Server Appliance clock is synchronized with the NTP server.

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.

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.

Installing vCenter Server on IPv6 Machines

vCenter Server 5.1 supports connection between vCenter Server and vCenter Server components by IP address only if the IP address is IPV4-compliant. To connect to vCenter Server system in an IPv6 environment you must use the fully qualified domain name (FQDN) or host name of the vCenter Server.

The best practice is to use the FQDN, which works in all cases, instead of the IP address, which can change if assigned by DHCP.

JDBC URL Formats for the vCenter Server Database

The vCenter Server installer generates and validates the JDBC URL for the vCenter Server database. If the installer fails to connect to the database using the generated JDBC URL, the installer will prompt you to specify the JDBC URL.

JDBC URL Note for All Databases

NOTE The domain name cannot contain the exclamation point character (!). Java interprets the exclamation point as a jar file separator.

JDBC URL Formats for Microsoft SQL Server Databases

For Microsoft SQL Server databases, you can use the following example JDBC URLs as a model:

- Connect to default (unnamed) SQL Server instance by host name:
`jdbc:sqlserver://host;databaseName=database`
- Connect to named instance by host name and instance name:
`jdbc:sqlserver://host;instanceName=instance;databaseName=database`
- Connect to SQL Server by host name and port:
`jdbc:sqlserver://host:port;databaseName=database`
- Connect by port:
`jdbc:sqlserver://localhost:1422;databaseName\=VIM_VCDB` (user name, password, and database type to be passed separately)
- Connect to local server with integrated security:
`jdbc:sqlserver://localhost\SQLEXP_VIM;databaseName=VIM_VCDB;integratedSecurity=true`
- Connect to local server without integrated security:
`jdbc:sqlserver://localhost\SQLEXP_VIM;databaseName\=VIM_VCDB` (user name, password, and database type to be passed separately)

VMware vCenter Server JDBC configuration for Microsoft SQL Server might not work by default with direct IPv6 addresses. You must use one of the following forms:

- Use the host name form for a standard Type-4 JDBC URL (recommended):
`jdbc:sqlserver://database-fully-qualified-host-name:port`
- Use direct IPv6 address format:
`jdbc:sqlserver://;serverName=[IPv6-address]`

For more information about JDBC URL formatting for MS SQL databases, including port and instance configuration options, see the [msdn.microsoft.com](http://msdn.microsoft.com/en-us/library/ms378428.aspx) Web site. At the time of this topic's publication, the information was available at <http://msdn.microsoft.com/en-us/library/ms378428.aspx>.

JDBC URL Formats for Oracle Databases

For Oracle databases, you can use the following example JDBC URLs as a model:

- This format requires host name and address, port (default 1521) and service name (for example, "oracle.world"):
`jdbc:oracle:thin:@host:port/service`

- This format requires host name and address, port (default 1521) and SID (for example, "ORCL"):

```
jdbc:oracle:thin:@host:port:SID
```

- This format is for a fully configured Oracle client with Oracle Net, which is useful for non-TCP configuration or Oracle RAC (real application clusters):

```
jdbc:oracle:thin:@tnsname
```

- The following example is for an Oracle RAC with a thin driver, without the full Oracle client installed:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=rac1-vip)(PORT=1521))
(ADDRESS=(PROTOCOL=TCP)(HOST=rac2-vip)(PORT=1521))(LOAD_BALANCE=yes)(FAILOVER=ON)
(CONNECT_DATA=(SERVER=DEDICATED)(SERVICE_NAME=RAC.DBTEAM)(FAILOVER_MODE=(BACKUP=rac1)
(TYPE=SELECT)(METHOD=BASIC))))
```

In this example, **rac1-vip** is first node virtual IP, **rac2-vip** is second node virtual IP, **RAC.DBTEAM** is RAC DB service name, and **rac1** is name of failover node.

For more information about JDBC URL formatting for Oracle databases, see the oracle.com Web site.

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 or Upgrading vCenter Single Sign-On, Inventory Service, vCenter Server, and the vSphere Web Client

Prepare for the vCenter Server installation by recording the values that vCenter Server and related components require.

The vCenter Single Sign-On, vSphere Web Client, vCenter Inventory Service, and vCenter Server installation wizards prompt you for the installation or upgrade information. Keep a record of the values entered, in case you must reinstall vCenter Server. You can print this topic as a worksheet to record the information that you need for the installation or upgrade of vCenter Single Sign-On, the vSphere Web Client, Inventory Service, and vCenter Server.

The following tables list the required information for installing or upgrading vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, vCenter Server.

- [Table 3-3.](#)
- [Table 3-4.](#)
- [Table 3-5.](#)
- [Table 3-6.](#)

NOTE Depending on the type of installation or upgrade you are doing, some entries might not be required.

Table 3-3. Information Required for vCenter Single Sign-On Installation.

Required Information	Default	Your Entry
Setup Language. This selection controls the language only for the installer.	English	
vCenter Single Sign-On HTTPS port.	7444	

Table 3-3. Information Required for vCenter Single Sign-On Installation. (Continued)

Required Information	Default	Your Entry
vCenter Single Sign-On deployment type. Choose from the following options: <ul style="list-style-type: none"> ■ vCenter Single Sign-On for your first vCenter Server. Select this option to create a new vCenter Single Sign-On server, which will become the first vCenter Single Sign-On server in a new domain.. ■ vCenter Single Sign-On for an additional vCenter Server in an existing site. Select this option to create an additional vCenter Single Sign-On server that replicates information from an existing vCenter Single Sign-On server in the domain. ■ vCenter Single Sign-On for an additional vCenter Server with a new site. Select this option to create an additional vCenter Single Sign-On server that replicates information from an existing vCenter Single Sign-On server in a different site. 		
Domain name.	vsphere.local	You cannot change the domain name from the default during installation.
User name.	administrator@vsphere.local	You cannot change the user name from the default during installation.
Password for the vCenter Single Sign-On administrator account in the default domain. You must use the same vCenter Single Sign-On password name when you install or upgrade vCenter Single Sign-On, Inventory Service, vCenter Server, and the vSphere Web Client. IMPORTANT Be sure to record the password. If you need to restore the vCenter Single Sign-On configuration from a backup, the restore process requires the password you enter for the original vCenter Single Sign-On installation, even if you change the password later. By default, the password must have at least eight characters, at least one lowercase character, one uppercase character, one number, and one special character. See the <i>vSphere Security</i> documentation for information about changing the password policy. The following characters are not supported in passwords: non-ASCII characters, semicolon (;), double quotation mark ("), single quotation mark ('), circumflex (^), and backslash (\).		
Site name. Your name for the vCenter Single Sign-On site.		
Partner host name. Required only if you are installing additional vCenter Single Sign-On servers. The partner host name is the DNS name of the existing vCenter Single Sign-On server to replicate from.		

Table 3-4. Information Required for the vSphere Web Client Installation

Required Information	Default	Your Entry
Setup Language. This selection controls the language only for the installer.	English	
Destination folder. The folder to install the vSphere Web Client in. The installation path cannot contain the following characters: non-ASCII characters, commas (,), periods (.), exclamation points (!), pound signs (#), at signs (@), or percentage signs (%). If 8.3 name creation is disabled on the host machine, do not install the vSphere Web Client in a directory that does not have an 8.3 short name or has a name that contains spaces. This situation will make the vSphere Web Client inaccessible.	C:\Program Files\VMware\Infrastructure	
vSphere Web Client HTTP port.	9090	
vSphere Web Client HTTPS port.	9443	
vCenter Single Sign-On administrator user name. The entry is case sensitive, and must match the administrator user name you enter when you install vCenter Single Sign-On	administrator@vsphere.local	
Password for the vCenter Single Sign-On administrator account in the default domain. You must use the same vCenter Single Sign-On password when you install or upgrade vCenter Single Sign-On, Inventory Service, vCenter Server, and the vSphere Web Client. IMPORTANT Be sure to record the password. If you need to restore the vCenter Single Sign-On configuration from a backup, the restore process requires the password you enter for the original vCenter Single Sign-On installation, even if you change the password later.		
Lookup Service URL. The Lookup Service URL takes the form <code>https://SSO_host_FQDN_or_IP:7444/lookupservice/sdk</code> , where 7444 is the default vCenter Single Sign-On HTTPS port number. This entry must match the URL you enter when you install vCenter Inventory Service.		

Table 3-5. Information Required for vCenter Inventory Service Installation or Upgrade

Required Information	Default	Your Entry
Setup Language. This selection controls the language only for the installer.	English	
Destination folder. The folder to install Inventory Service in. The installation path cannot contain the following characters: non-ASCII characters, commas (,), periods (.), exclamation points (!), pound signs (#), at signs (@), or percentage signs (%).	C:\Program Files\VMware\Infrastructure	
Fully Qualified Domain Name. The FQDN for the Inventory Service local system.		
vCenter Inventory Service HTTPS port.	10443	
vCenter Inventory Service management port.	10109	See “Required Ports for vCenter Server,” on page 23.
vCenter Inventory Service Linked Mode communication port.	10111	

Table 3-5. Information Required for vCenter Inventory Service Installation or Upgrade (Continued)

Required Information	Default	Your Entry
Inventory size. The inventory size of your vCenter Server deployment: <ul style="list-style-type: none"> ■ Small (less than 100 hosts or 1000 virtual machines. ■ Medium (100-400 hosts or 1000-4000 virtual machines. ■ Large (more than 400 hosts or 4000 virtual machines. This setting determines the maximum JVM heap settings for VMware VirtualCenter Management Webservices (Tomcat), Inventory Service, and Profile-Driven Storage Service. You can adjust this setting after installation if the number of hosts in your environment changes. See the recommendations in “Hardware Requirements for vCenter Server, the vSphere Web Client, vCenter Inventory Service, and vCenter Single Sign-On,” on page 17.		
User name for the vCenter Single Sign-On administrator user account. You must use the same vCenter Single Sign-On user name and password name when you install vCenter Single Sign-On, and install or upgrade Inventory Service, vCenter Server, and the vSphere Web Client.	administrator	
Lookup Service URL. The Lookup Service URL takes the form <code>https://SSO_host_FQDN_or_IP:7444/lookupservice/sdk</code> , where 7444 is the default vCenter Single Sign-On HTTPS port number. If you enter a different port number when you install vCenter Single Sign-On, use that port number.		

Table 3-6. Information Required for vCenter Server Installation or Upgrade

Required Information	Default	Your Entry
Setup Language. This selection controls the language only for the installer.	English	
vCenter Server license key. If you omit the license key, vCenter Server is installed in evaluation mode. After you install vCenter Server, you can enter the vCenter Server license in the vSphere Web Client.		
Data source name (DSN). Required if you use an existing database. Not required if you are using the bundled Microsoft SQL Server 2008 Express database. Leading and trailing spaces are not supported. Remove spaces from the beginning or end of the DSN.		
Database user name. Database password.	Required to use an existing database. Not required if you are using the bundled database. Non-ASCII characters are not supported.	
JDBC URL for database. Required if you use an existing database. The vCenter Server installer should generate and validate the JDBC URL for the vCenter Server database. If the installer fails to connect to the database by using the generated JDBC URL, the installer prompts you to specify the JDBC URL. The format of the JDBC URL depends on the database that you are using. See “JDBC URL Formats for the vCenter Server Database,” on page 61.v		

Table 3-6. Information Required for vCenter Server Installation or Upgrade (Continued)

Required Information	Default	Your Entry
vCenter Server Service account information. Can be the Microsoft Windows system account or a user-specified account. Use a user-specified account if you plan to use Microsoft Windows authentication for SQL Server.	Microsoft Windows system account	
Fully qualified domain name (FQDN) for the vCenter Server machine The FQDN of the system that you are installing vCenter Server on. The vCenter Server installer checks that the FQDN is resolvable. If not, a warning message appears. Change the entry to a resolvable FQDN. You must enter the FQDN, not the IP address.		
Standalone or join group. Join a Linked Mode group to enable the vSphere Web Client to view, search, and manage data across multiple vCenter Server systems.	Standalone	
Fully qualified domain name of Directory Services for the vCenter Server group. The FQDN of a remote instance of vCenter Server. Required if this instance of vCenter Server is joining a group. The local and remote instances will be members of a Linked Mode group.		
LDAP port for the Directory Services for the remote vCenter Server instance. The LDAP port of the remote instance. Required if this instance of vCenter Server is joining a Linked Mode group. See “Required Ports for vCenter Server,” on page 23.	389	
vCenter Server HTTPS port.	443	
vCenter Server HTTP port.	80	
Heartbeat port (UDP) used for sending data to ESX/ESXi hosts.	902	
VMware VirtualCenter Management Web Services HTTP port	8080	
VMware VirtualCenter Management Web Services HTTPS port.	8443	
Web Services change service notification port.	60099	
LDAP port for the Directory Services for the local vCenter Server instance.	389	

Table 3-6. Information Required for vCenter Server Installation or Upgrade (Continued)

Required Information	Default	Your Entry
SSL port for the Directory Services for the local vCenter Server instance.	636	
Ephemeral ports. Select Increase the number of available ephemeral ports if your vCenter Server manages hosts on which you will power on more than 2000 virtual machines simultaneously. This option prevents the pool of available ephemeral ports from being exhausted.		
Inventory size. The inventory size of your vCenter Server deployment: <ul style="list-style-type: none"> ■ Small (less than 100 hosts or 1000 virtual machines. ■ Medium (100-400 hosts or 1000-4000 virtual machines. ■ Large (more than 400 hosts or 4000 virtual machines. This setting determines the maximum JVM heap settings for VMware VirtualCenter Management Webservices (Tomcat), Inventory Service, and Profile-Driven Storage Service. You can adjust this setting after installation if the number of hosts in your environment changes. See the recommendations in “Hardware Requirements for vCenter Server, the vSphere Web Client, vCenter Inventory Service, and vCenter Single Sign-On,” on page 17.		
User name for the vCenter Single Sign-On administrator user account.	administrator	
Password for the vCenter Single Sign-On administrator user account.	You must use the same vCenter Single Sign-On user name and password name when you install vCenter Single Sign-On, and install or upgrade Inventory Service, vCenter Server, and the vSphere Web Client.	
Lookup Service URL. The Lookup Service URL takes the form <code>https://SSO_host_FQDN_or_IP:7444/lookupservice/sdk</code> , where 7444 is the default vCenter Single Sign-On HTTPS port number. If you enter a different port number when you install vCenter Single Sign-On, use that port number.		
Inventory Service URL. The inventory Service URL takes the form <code>https://Inventory_Service_host_FQDN_or_IP:10443</code> . 10443 is the default Inventory Service HTTPS port number. If you enter a different port number when you install Inventory Service, use that port number.		
Destination folder. The folder to install vCenter Server in. The installation path cannot contain the following characters: non-ASCII characters, commas (,), periods (.), exclamation points (!), pound signs (#), at signs (@), or percentage signs (%).	C:\Program Files\VMware\Infrastructure	

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?.

Installing vCenter Server

Install vCenter Server to manage your vSphere system.

This chapter includes the following topics:

- [“vCenter Server Installation and Sign-In Process,”](#) on page 69
- [“vCenter Server Components and Support Tools,”](#) on page 72
- [“Download the vCenter Server Installer,”](#) on page 73
- [“Install vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server by Using Simple Install,”](#) on page 73
- [“Use Custom Install to Install vCenter Server and Required Components,”](#) on page 76
- [“Add a vCenter Single Sign-On Identity Source,”](#) on page 85
- [“Install or Upgrade vCenter Server Java Components Separately,”](#) on page 90
- [“Install or Upgrade vCenter Server tc Server Separately,”](#) on page 91
- [“vCenter Single Sign-On Installation Fails,”](#) on page 91
- [“Download and Deploy the VMware vCenter Server Appliance,”](#) on page 92

vCenter Server Installation and Sign-In Process

As part of installation, you are prompted to make choices. Understand the complete installation, vCenter Single Sign-On setup, and permission assignment process before you start.

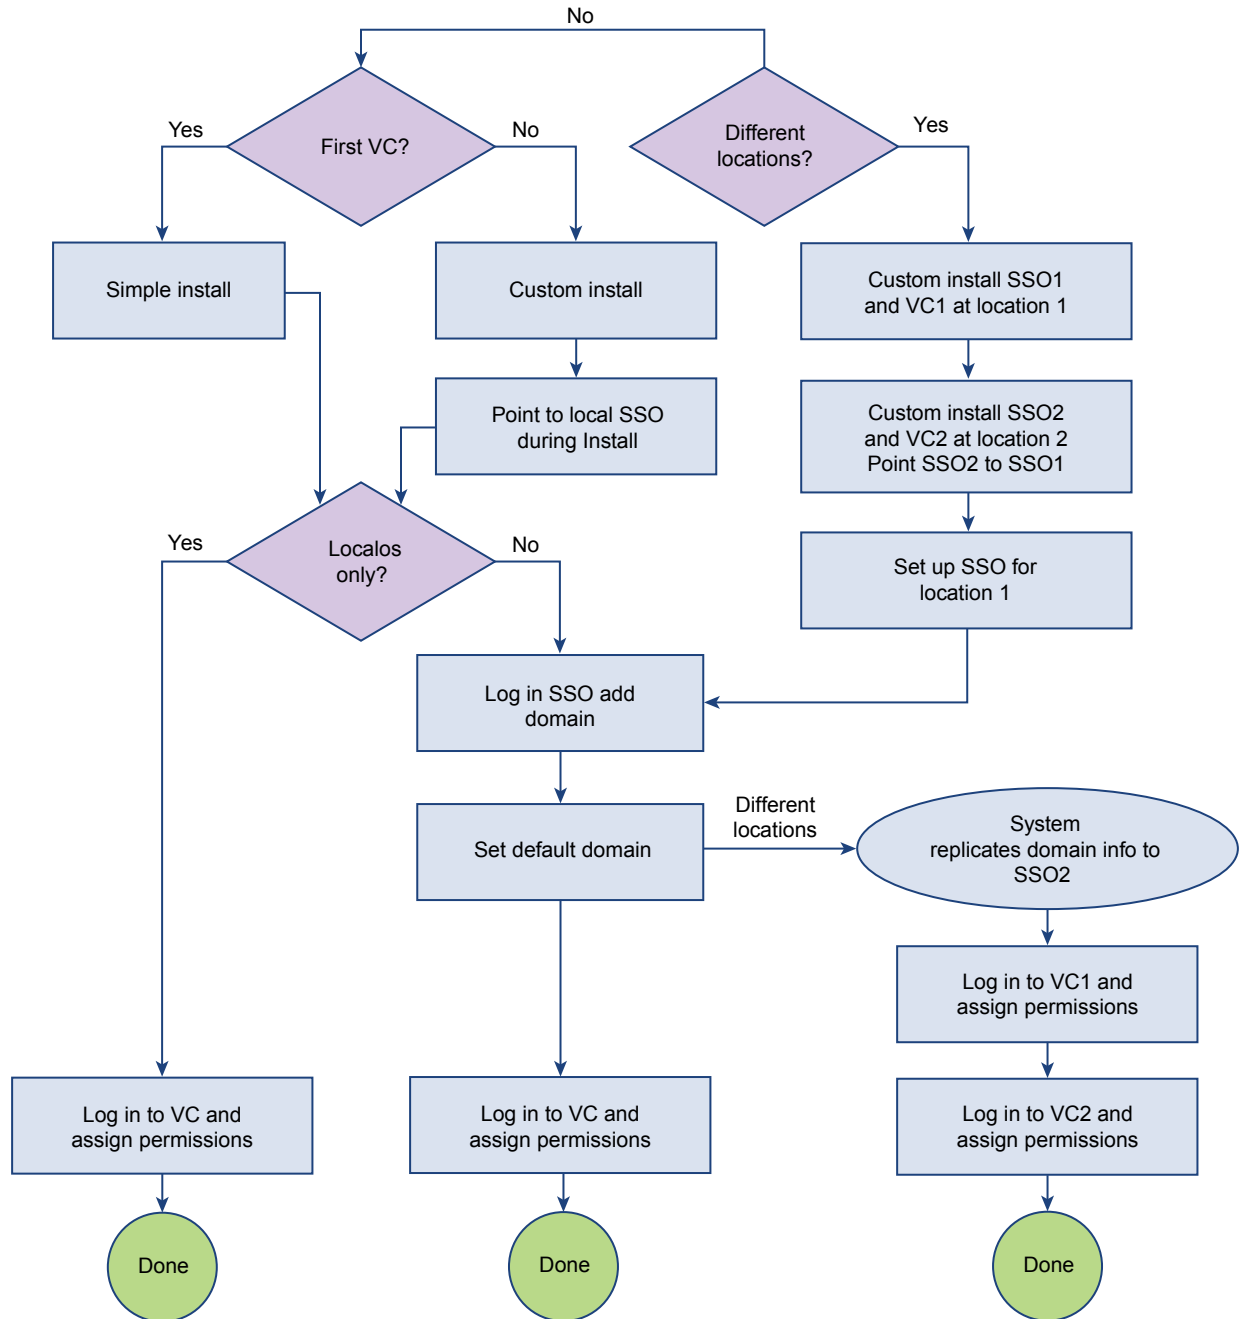
Before installing vCenter Server, consider your environment and requirements to make the following decisions:

- Decide how to set up vCenter Server services: on a single Windows platform or vCenter Server Appliance, or distributed across multiple Windows hosts or vCenter Server Appliances?
- Decide whether you want your vSphere deployment to authenticate and give permissions to Active Directory users and groups.
- Decide whether you want your vSphere deployment to authenticate and give permissions to Open LDAP users and groups.
- Decide whether to use an embedded or an external vCenter Server database
- If you have multiple vCenter Servers, decide whether they should be standalone instances or joined in a Linked Mode group.
- Decide whether your vCenter Server deployment requires high availability protection.
- Decide whether you will be deploying vSphere across multiple sites.

For guidance in making these decisions, review the sections in [Chapter 3, “Before You Install vCenter Server,”](#) on page 29.

After you verify your hardware and software environment, you can start installation. Your interaction with the installer and with vCenter Single Sign-On is shown in the figure below.

Figure 4-1. Flowchart of vCenter Installation and User Management Process



The interaction proceeds as follows.

- 1 You decide whether to install the vCenter components in different locations.
 - For a single location, the process differs if you perform a first installation or you are installing multiple vCenter Server systems at the same location. Continue with step 2.

- For multiple locations for the different components, you can perform a custom install (step 5) and set up an Active Directory or OpenLDAP identity source.
- 2 The installation type you select when you install everything in one location depends on whether this is the first vCenter Server system.
 - If this is the first or the only vCenter Server system that you install at this location, you can perform a Simple Install. See [“Install vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server by Using Simple Install,”](#) on page 73.

NOTE Use Simple Install unless you have good reasons to install the vCenter Server in different locations.

- If this is not the first vCenter Server system that you install at this location, you can perform a Custom Install and point to the vCenter Single Sign-On service that you installed earlier when prompted. See [“Use Custom Install to Install vCenter Server and Required Components,”](#) on page 76.
- 3 After an installation in one location, local operating system users and the user administrator@vsphere.local can authenticate.
 - If the localos identity source is sufficient, you can log in to vCenter Server as administrator@vsphere.local and assign permissions to other local operating system (localos) users. See User Management Tasks in *vSphere Security*.
 - If an Active Directory identity source is needed, continue to step 4.
 - 4 If you want to use an Active Directory or OpenLDAP identity source:
 - a Log in to the vCenter Single Sign-On server as administrator@vsphere.local and add the domain as an identity source.
 - b Add the Active Directory or OpenLDAP domain as an identity source.
 - c Log in to vCenter Server and assign permissions so users in your domain can access vCenter Server components.

See [“Add a vCenter Single Sign-On Identity Source,”](#) on page 85 and [“Assign Permissions in the vSphere Web Client,”](#) on page 87.

- 5 If you decide to install vCenter Server in multiple locations, you can install one or more vCenter Server components and a vCenter Single Sign-On component in each location. See [“Use Custom Install to Install vCenter Server and Required Components,”](#) on page 76.
 - a Custom Install vCenter Server and vCenter Single Sign-On in location 1.
 - b Custom Install vCenter Server and vCenter Single Sign-On in location 2.

As part of the installation of the second system, point the vCenter Single Sign-On system to the first vCenter Single Sign-On system you installed. If you want, you can also set up Linked Mode for the vCenter Server instances.
 - c Set up vCenter Single Sign-On for location 1 using the procedure in Steps 3 and 4.
- 6 After you complete the setup of vCenter Single Sign-On in one location, vCenter Single Sign-On settings in all other locations are updated immediately.
- 7 Log in to vCenter Server and assign permissions to users and groups.

You set permissions for each vCenter Server instance. See [“Assign Permissions in the vSphere Web Client,”](#) on page 87.

vCenter Server Components and Support Tools

When you install vCenter Server, other components are also installed.

The following components can also be installed when you install vCenter Server.

VMware vCenter Server	Windows service to manage ESXi and legacy ESX hosts.
vCenter Single Sign-On	The vCenter Single Sign-On authentication service makes the VMware cloud infrastructure platform more secure by allowing the various vSphere software components to 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. Available and required with vCenter Server 5.1.
vCenter Inventory Service	Inventory Service stores vCenter Server application and inventory data, enabling you to search and access inventory objects across linked vCenter Servers.
Microsoft .NET 3.5 SP1 Framework	Software used by the Database Upgrade wizard and the vSphere Web Client. Also used by vCenter Server if you are using the bundled database. If it is not installed on your system, the vCenter Server installer installs it.
Microsoft Windows Installer version 4.5	If you plan to use the Microsoft SQL Server 2008 R2 Express database that is bundled with vCenter Server, Microsoft Windows Installer version 4.5 (MSI 4.5) is required on your system. You can also install MSI 4.5 directly from the vCenter Server <code>autorun.exe</code> installer.
VMware vCenter Orchestrator	vCenter Server module that provides a set of tools to manage your virtual IT environment. vCenter Orchestrator module is not supported on IPv6-only operating systems. If you install vCenter Server in a mixed environment (both IPv4 and IPv6 enabled), the vCenter Orchestrator module can be configured using IPv4. See the <i>Administering vCenter Orchestrator</i> . This component is installed automatically with vCenter Server.
Microsoft SQL Server 2008 R2 Express (optional)	Free, bundled version of the Microsoft SQL Server database for smaller scale applications. If you choose to use an existing database, the installer does not install the bundled database.
vSphere Web Client	Server application that you can use to manage an ESXi host by Web browser through a vCenter Server.
vSphere Update Manager	vCenter Server component that provides security monitoring and patching support for hosts and virtual machines.
vSphere ESXi Dump Collector	vCenter Server support tool. You can configure ESXi to dump the vmkernel memory to a network server, rather than to a disk, when the system has encountered a critical failure. The Dump Collector collects such memory dumps over the network.
vSphere Syslog Collector	vCenter Server support tool that provides a unified architecture for system logging and enables network logging and combining of logs from multiple hosts.

vSphere Auto Deploy

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.

vSphere Authentication Proxy

vCenter Server support tool that enables ESXi hosts to join a domain without using Active Directory credentials. This tool enhances security for PXE-booted hosts and hosts that are provisioned using Auto Deploy, by removing the need to store Active Directory credentials in the host configuration.

Download the vCenter Server Installer

Download the installer for vCenter Server, the vSphere Web Client, and associated vCenter 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 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>.

Install vCenter Single Sign-On, the vSphere Web Client , vCenter Inventory Service, and vCenter Server by Using Simple Install

You can install vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server together on a single host machine by using the vCenter Server Simple Install option. This option is appropriate for most deployments.

Alternatively, you can use Custom Install to install vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server separately and customize the location and configuration of each component. See [“Use Custom Install to Install vCenter Server and Required Components,”](#) on page 76.

If vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, or vCenter Server is already installed on the computer, this procedure upgrades the existing version.

Prerequisites

- Review the topics in [Chapter 3, “Before You Install vCenter Server,”](#) on page 29.
- Review [“Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server,”](#) on page 49.

Procedure

- 1 [Install vCenter Single Sign-On, the vSphere Web Client, and vCenter Inventory Service as Part of a vCenter Server Simple Install](#) on page 74

Create the only node in a basic, Simple Install vCenter Single Sign-On installation, and install the vSphere Web Client and vCenter Inventory Service.

2 [Install vCenter Server as Part of a Simple Install](#) on page 75

In a simple Install, after you install vCenter Single Sign-On, the vSphere Web Client, and vCenter Inventory Service, the installer starts the vCenter Server installation.

Install vCenter Single Sign-On, the vSphere Web Client, and vCenter Inventory Service as Part of a vCenter Server Simple Install

Create the only node in a basic, Simple Install vCenter Single Sign-On installation, and install the vSphere Web Client and vCenter Inventory Service.

You can use Simple Install for the first vCenter Single Sign-On and vCenter Server in a deployment with multiple vCenter Servers. Succeeding instances of vCenter Single Sign-On and vCenter Server in the same deployment must be installed by using Custom Install. For more information about vCenter Single Sign-On, see [“How vCenter Single Sign-On Affects vCenter Server Installation,”](#) on page 51. and the *vSphere Security* documentation.

NOTE vCenter Server 5.5 supports connection between vCenter Server and vCenter Server components by IP address only if the IP address is IPv4-compliant. To connect to a vCenter Server system in an IPv6 environment, you must use the fully qualified domain name (FQDN) or host name of the vCenter Server. The best practice is to use the FQDN, which works in all cases, instead of the IP address, which can change if assigned by DHCP.

Prerequisites

- See [“Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server,”](#) on page 49
- Download the vCenter Server installer. See [“Download the vCenter Server Installer,”](#) on page 73.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vCenter™ Simple Install**, and click **Install**.
- 3 Follow the prompts in the installation wizard to choose the installer language, and agree to the end user patent and license agreements.
- 4 If the prerequisites check screen shows any problems, cancel the installation, correct the problems, and restart the installer.
- 5 Set the password for the vCenter Single Sign-On administrator account.

This is the password for the user `administrator@vsphere.local`. `vsphere.local` is a new domain that is created by vCenter Single Sign-On. After installation, you can log in to vCenter Single Sign-On and in to vCenter Server as `adminstrator@vsphere.local`.

By default, the password must have at least eight characters, at least one lowercase character, one uppercase character, one number, and one special character. See the *vSphere Security* documentation for information about changing the password policy. The following characters are not supported in passwords: non-ASCII characters, semicolon (;), double quotation mark ("), single quotation mark ('), circumflex (^), and backslash (\).

- 6 Enter the site name for vCenter Single Sign-On.
Choose your own name for the vCenter Single Sign-On site.
- 7 Accept or change the HTTPS port for vCenter Single Sign-On.
- 8 Select the folder in which to install vCenter Single Sign-On.

The installation path cannot contain any of the following characters: non-ASCII characters, commas (,), periods (.), exclamation points (!), pound signs (#), at signs (@), or percentage signs (%).

- 9 Review the installation options and click **Install**.

The vCenter Single Sign-On installation begins. When the vCenter Single Sign-On installation is complete, the installer proceeds with the vSphere Web Client and vCenter Inventory Service installations.

No input is required for a new Simple Install installation of the vSphere Web Client and vCenter Inventory Service.

NOTE After each component is installed, the installer might take a few minutes to start the installer for the next component.

Install vCenter Server as Part of a Simple Install

In a simple Install, after you install vCenter Single Sign-On, the vSphere Web Client, and vCenter Inventory Service, the installer starts the vCenter Server installation.

NOTE vCenter Server 5.5 supports connection between vCenter Server and vCenter Server components by IP address only if the IP address is IPv4-compliant. To connect to a vCenter Server system in an IPv6 environment, you must use the fully qualified domain name (FQDN) or host name of the vCenter Server. The best practice is to use the FQDN, which works in all cases, instead of the IP address, which can change if assigned by DHCP.

Prerequisites

- See [“Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server,”](#) on page 49

Procedure

- 1 (Optional) Enter your license key.

If you omit the license key, vCenter Server is in evaluation mode, which allows you to use the full feature set for a 60-day evaluation period. After installation, you can enter the license key to convert vCenter Server to licensed mode.

- 2 Choose the type of database that you want to use.

- To use the bundled database, click **Install a Microsoft SQL Server 2008 Express instance (for small-scale deployments)**.

This database is suitable for deployments of up to 5 hosts and 50 virtual machines.

- To use an existing database, click **Use an existing supported database** and select your database from the list of available DSNs. Enter the user name and password for the DSN.

If your database uses Windows NT authentication, the user name and password fields are disabled.

NOTE A warning might appear that the DSN points to an older version of a repository that must be upgraded. If you click **Yes**, the installer upgrades the database schema, making the database irreversibly incompatible with previous VirtualCenter versions. See the *vSphere Upgrade* documentation.

- 3 Set the vCenter Server service account information.

- If you are using a nonbundled database, enter the administrator name and password that you use when you log in to the system on which you are installing vCenter Server.
- If you are using the bundled SQL Server database, select **Use Windows Local System Account**.

You need the user name and password entered here to log in to vCenter Server after you have installed it.

The Fully Qualified Domain Name field displays the FQDN of the system that you are installing vCenter Server on. The vCenter Server installer checks that the FQDN is resolvable. If not, a warning message appears when you click **Next**. Change the entry to a resolvable FQDN. You must enter the FQDN, not the IP address.

- 4 For each component that you install, accept the default port numbers, or if another service is using the defaults, enter alternative ports.
- 5 (Optional) Select **Increase the number of available ephemeral ports**.
- 6 Select the size of your vCenter Server inventory to allocate memory for several Java services that are used by vCenter Server.

This setting determines the maximum JVM heap settings for VMware VirtualCenter Management Webservices (Tomcat), Inventory Service, and Profile-Driven Storage Service. You can adjust this setting after installation if the number of hosts in your environment changes. See the recommendations in the topic vCenter Server Hardware Requirements.

- 7 Click **Install**.

Multiple progress bars appear during the installation of the selected components.

The vCenter Simple Install is complete.

What to do next

See [Chapter 5, “After You Install vCenter Server,”](#) on page 99.

Use Custom Install to Install vCenter Server and Required Components

You can install vCenter Server and other vCenter components separately to customize the location and configuration of each component.

For most basic vCenter Single Sign-On deployments, if all components are on the same host machine, you can install vCenter Single Sign-On, the vSphere Web Client, Inventory Service, and vCenter Server together on a single host machine using the vCenter Server Simple Install option.

See [“Install vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server by Using Simple Install,”](#) on page 73.

NOTE vCenter Server 5.5 supports connection between vCenter Server and vCenter Server components by IP address only if the IP address is IPv4-compliant. To connect to a vCenter Server system in an IPv6 environment, you must use the fully qualified domain name (FQDN) or host name of the vCenter Server. The best practice is to use the FQDN, which works in all cases, instead of the IP address, which can change if assigned by DHCP.

Prerequisites

- Review [Chapter 3, “Before You Install vCenter Server,”](#) on page 29.
- Review [“Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server,”](#) on page 49

Procedure

- 1 [Install the First or Only vCenter Single Sign-On Instance in a vCenter Server Deployment](#) on page 77
Create the only vCenter Single Sign-On instance in a basic vCenter Single Sign-On installation or the first vCenter Single Sign-On instance in a deployment with multiple vCenter Single Sign-On instances.

- 2 [\(Optional\) Install an Additional vCenter Single Sign-On Node at an Existing Site](#) on page 78
Create an additional vCenter Single Sign-On node at an existing vCenter Single Sign-On installation. An additional vCenter Single Sign-On node might be useful if your deployment includes multiple vCenter Server instances.
- 3 [\(Optional\) Install an Additional vCenter Single Sign-On Node at a New Site](#) on page 79
Create an additional vCenter Single Sign-On node for a multisite vCenter Single Sign-On installation. An additional node can be useful if you need multiple vCenter Server instances in different locations. Authentication information is replicated between vCenter single Sign-On instances that are related.
- 4 [Install or Upgrade the vSphere Web Client](#) on page 80
The vSphere Web Client lets you connect to a vCenter Server system to manage your vSphere deployment through a browser.
- 5 [Install vCenter Inventory Service Separately by Using Custom Install](#) on page 81
You can use Custom Install to install vCenter Single Sign-On, vCenter Inventory Service, and vCenter Server separately to customize the location and configuration of the components.
- 6 [Install vCenter Server as Part of a Custom Install](#) on page 83
You can install vCenter Server separately from vCenter Single Sign-On and vCenter Inventory Service to customize the location and configuration of the components.

Install the First or Only vCenter Single Sign-On Instance in a vCenter Server Deployment

Create the only vCenter Single Sign-On instance in a basic vCenter Single Sign-On installation or the first vCenter Single Sign-On instance in a deployment with multiple vCenter Single Sign-On instances.

These instructions let you install vCenter Single Sign-On only. You must install vCenter Single Sign-On and Inventory Service before installing vCenter Server. For most deployments, you can install vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server together on a single host machine by using vCenter Server Simple Install. See [“vCenter Single Sign-On Deployment Modes,”](#) on page 52 and [“Install vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server by Using Simple Install,”](#) on page 73.

For more information about vCenter Single Sign-On, see [“How vCenter Single Sign-On Affects vCenter Server Installation,”](#) on page 51. and the *vSphere Security* documentation.

NOTE vCenter Server 5.5 supports connection between vCenter Server and vCenter Server components by IP address only if the IP address is IPv4-compliant. To connect to a vCenter Server system in an IPv6 environment, you must use the fully qualified domain name (FQDN) or host name of the vCenter Server. The best practice is to use the FQDN, which works in all cases, instead of the IP address, which can change if assigned by DHCP.

Prerequisites

- Review [“vCenter Single Sign-On Deployment Modes,”](#) on page 52.
- Review [“How vCenter Single Sign-On Affects vCenter Server Installation,”](#) on page 51. .
- Review [“Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server,”](#) on page 49
- Download the vCenter Server installer. See [“Download the vCenter Server Installer,”](#) on page 73.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vCenter Single Sign-On** and click **Install**.

- 3 Follow the prompts in the installation wizard to choose the installer language, and agree to the end user patent and license agreements.
- 4 If the prerequisites check screen shows any problems, cancel the installation, correct the problems, and restart the installer.
- 5 Set the password for the vCenter Single Sign-On administrator account.

This is the password for the user administrator@vsphere.local. vsphere.local is a new domain that is created by vCenter Single Sign-On. After installation, you can log in to vCenter Single Sign-On and in to vCenter Server as administrator@vsphere.local.

By default, the password must have at least eight characters, at least one lowercase character, one uppercase character, one number, and one special character. See the *vSphere Security* documentation for information about changing the password policy. The following characters are not supported in passwords: non-ASCII characters, semicolon (;), double quotation mark ("), single quotation mark ('), circumflex (^), and backslash (\).

- 6 Accept or change the HTTPS port for vCenter Single Sign-On.
- 7 Select the deployment mode **vCenter Single Sign-On for your first vCenter Server**.
- 8 Enter the site name for vCenter Single Sign-On.
Choose your own name for the vCenter Single Sign-On site.
- 9 Review the installation options and click **Install**.

vCenter Single Sign-On is installed.

After vCenter Single Sign-On is installed or upgraded, the following default identity sources and users are available:

localos	All local operating system users. These users can be granted permissions to vCenter Server. If you are upgrading, those users who already have permissions keep those permissions.
vsphere.local	Contains all users who have administrator access to the vCenter Single Sign-On server. Initially, only the user administrator is defined.

What to do next

To deploy vCenter Server with multiple vCenter Single Sign-On instances, install an additional vCenter Single Sign-On at an existing or new site. See [“\(Optional\) Install an Additional vCenter Single Sign-On Node at an Existing Site,”](#) on page 78 or [“\(Optional\) Install an Additional vCenter Single Sign-On Node at a New Site,”](#) on page 79. If your vCenter Server deployment requires only one vCenter Single-Sign-On instance, install the vSphere Web Client. See [“Install or Upgrade the vSphere Web Client,”](#) on page 101.

To add other identity sources, such as a native Active Directory (Integrated Windows Authentication) domain or an OpenLDAP directory service, see [“Add a vCenter Single Sign-On Identity Source,”](#) on page 85.

(Optional) Install an Additional vCenter Single Sign-On Node at an Existing Site

Create an additional vCenter Single Sign-On node at an existing vCenter Single Sign-On installation. An additional vCenter Single Sign-On node might be useful if your deployment includes multiple vCenter Server instances.

Prerequisites

- Review [“vCenter Single Sign-On Deployment Modes,”](#) on page 52.
- See [“Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server,”](#) on page 49.

- Install the first node in the vCenter Single Sign-On installation. See [“Install the First or Only vCenter Single Sign-On Instance in a vCenter Server Deployment,”](#) on page 77.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vCenter Single Sign-On** and click **Install**.
- 3 Follow the prompts in the installation wizard to choose the installer language, and agree to the end user patent and license agreements.
- 4 Accept or change the HTTPS port for vCenter Single Sign-On.
- 5 Select the deployment mode **vCenter Single Sign-On for an additional vCenter Server in an existing site**.
- 6 Enter the information to point this additional node to the first vCenter Single Sign-On server.

NOTE If the primary node is in a high-availability cluster, enter the address of the primary node load balancer.

- a Enter the **Partner host name**.
The partner host name is the DNS name of the existing vCenter Single Sign-On server to replicate from.
 - b Enter the password for the vCenter Single Sign-On administrator account of the existing vCenter Single Sign-On server (`administrator@vsphere.local`).
- 7 Select an existing site as the partner or enter a new site.
 - 8 Click **Install**.

(Optional) Install an Additional vCenter Single Sign-On Node at a New Site

Create an additional vCenter Single Sign-On node for a multisite vCenter Single Sign-On installation. An additional node can be useful if you need multiple vCenter Server instances in different locations. Authentication information is replicated between vCenter single Sign-On instances that are related.

Prerequisites

- Review [“vCenter Single Sign-On Deployment Modes,”](#) on page 52.
- See [“Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server,”](#) on page 49.
- Install the first node in the vCenter Single Sign-On installation. See [“Install the First or Only vCenter Single Sign-On Instance in a vCenter Server Deployment,”](#) on page 77.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vCenter Single Sign-On** and click **Install**.
- 3 Follow the prompts in the installation wizard to choose the installer language, and agree to the end user patent and license agreements.
- 4 Accept or change the HTTPS port for vCenter Single Sign-On.
- 5 Select the deployment mode **vCenter Single Sign-On for an additional vCenter Server with a new site**.

- 6 Enter the information to point this additional node to the first vCenter Single Sign-On server.

NOTE If the primary node is in a high-availability cluster, enter the address of the primary node load balancer.

- a Enter the **Partner host name**.

The partner host name is the DNS name of the existing vCenter Single Sign-On server to replicate from.

- b Enter the password for the vCenter Single Sign-On administrator account of the existing vCenter Single Sign-On server (administrator@vsphere.local).

- 7 Select an existing site as the partner or enter a new site.

- 8 Click **Install**.

The additional vCenter Single Sign-On server is installed.

What to do next

Repeat this procedure for each additional node.

Install or Upgrade the vSphere Web Client

The vSphere Web Client lets you connect to a vCenter Server system to manage your vSphere deployment through a browser.

If an earlier version of the vSphere Web Client is installed, this procedure upgrades the vSphere Web Client.

NOTE vCenter Server 5.5 supports connection between vCenter Server and vCenter Server components by IP address only if the IP address is IPv4-compliant. To connect to a vCenter Server system in an IPv6 environment, you must use the fully qualified domain name (FQDN) or host name of the vCenter Server. The best practice is to use the FQDN, which works in all cases, instead of the IP address, which can change if assigned by DHCP.

Prerequisites

- Download the vCenter Server installer. See [“Download the vCenter Server Installer,”](#) on page 73.
- Verify that the system has an Internet connection.
- Verify that the system meets the software requirements for the vSphere Web Client. See [“vSphere Web Client Software Requirements,”](#) on page 22.
- Before you install or upgrade any vSphere product, synchronize the clocks of all machines on the vSphere network. See [“Synchronizing Clocks on the vSphere Network,”](#) on page 59.
- Install vCenter Single Sign-On, or upgrade to the current version.
- Verify that the vSphere Web Client and vCenter Server are registered to the same vCenter Single Sign-On server, to ensure that the vSphere Web Client can access the vCenter Server inventory.
- Close all browsers before installing or uninstalling the vSphere Web Client.
- Log in as a member of the Administrators group on the host machine, with a user name that does not contain any non-ASCII characters.

Procedure

- 1 In the software installer directory, double-click the autorun.exe file to start the installer.
- 2 Select **vSphere Web Client** and click **Install**.

- 3 Follow the prompts in the installation wizard to choose the installer language, and agree to the end user patent and license agreements.
- 4 Either accept the default destination folder or click **Change** to select another location.

The installation path cannot contain any of the following characters: non-ASCII characters, commas (,), periods (.), exclamation points (!), pound signs (#), at signs (@), or percentage signs (%).

If 8.3 name creation is disabled on the host machine, do not install the vSphere Web Client in a directory that does not have an 8.3 short name or has a name that contains spaces. This situation will make the vSphere Web Client inaccessible.

- 5 Accept or change the default port settings.
- 6 Enter the information to register the vSphere Web Client with vCenter Single Sign-On.

The vCenter Single Sign-On administrator user name is administrator@vsphere.local, and the password must match the password you entered for the administrator user when you installed vCenter Single Sign-On. The Lookup Service URL takes the form https://SSO_host_FQDN_or_IP:7444/lookupservice/sdk, where 7444 is the default vCenter Single Sign-On HTTPS port number. Your entry should match the entry you made when you installed vCenter Single Sign-On. If you entered a different port number when you installed vCenter Single Sign-On, use that port number.

- 7 Click **Install**.
- 8 Start the vSphere Web Client by taking one of the following actions.
 - If you are starting the vSphere Web Client for the first time, open a supported browser, and go to https://vSphere_Web_Client_host_name_or_IP:9443/vsphere-client.
 - In subsequent sessions, you can start the vSphere Web Client from the Windows Start menu, by selecting **Programs > VMware > VMware vSphere Web Client > vSphere Web Client**.

NOTE After you upgrade the vSphere Web Client, when you log in for the first time, you may see the error message *Failed to navigate to desired location*. This can happen when a vSphere Web Client session from the previous version remains open when you upgrade. In this case, refresh the browser and log in again.

What to do next

Install the Client Integration Plug-In in the vSphere Web Client. See [“Install the Client Integration Plug-In in the vSphere Web Client,”](#) on page 100

Install vCenter Inventory Service Separately by Using Custom Install

You can use Custom Install to install vCenter Single Sign-On, vCenter Inventory Service, and vCenter Server separately to customize the location and configuration of the components.

These instructions install vCenter Inventory Service only. You must install vCenter Single Sign-On before installing Inventory Service and vCenter Server. For most deployments, you can install vCenter Single Sign-On, the vSphere Web Client, Inventory Service, and vCenter Server together on a single host machine using the vCenter Server Simple Install option. See [“vCenter Single Sign-On Deployment Modes,”](#) on page 52 and [“Install vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server by Using Simple Install,”](#) on page 73.

NOTE vCenter Server 5.5 supports connection between vCenter Server and vCenter Server components by IP address only if the IP address is IPv4-compliant. To connect to a vCenter Server system in an IPv6 environment, you must use the fully qualified domain name (FQDN) or host name of the vCenter Server. The best practice is to use the FQDN, which works in all cases, instead of the IP address, which can change if assigned by DHCP.

Prerequisites

- Review “[vCenter Single Sign-On Deployment Modes](#),” on page 52.
- Review “[How vCenter Single Sign-On Affects vCenter Server Installation](#),” on page 51. .
- Review “[Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server](#),” on page 49
- Download the vCenter Server Installer.
- Install vCenter Single Sign-On.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vCenter Inventory Service** and click **Install**.
- 3 Follow the prompts in the installation wizard to choose the installer language, and agree to the end user patent and license agreements.
- 4 Accept or change the default installation folder.

The installation path cannot contain any of the following characters: non-ASCII characters, commas (,), periods (.), exclamation points (!), pound signs (#), at signs (@), or percentage signs (%).

- 5 Enter the fully qualified domain name for the Inventory Service host machine.
- 6 If you are upgrading or reinstalling an existing instance of Inventory Service, choose whether to keep the existing database or replace it with a new empty database.
- 7 Accept or change the default values for Inventory Service port numbers.
- 8 Select the size of your vCenter Server inventory to allocate memory for several Java services that are used by vCenter Server.

This setting determines the maximum JVM heap settings for VMware VirtualCenter Management Webservices (Tomcat), Inventory Service, and Profile-Driven Storage Service. You can adjust this setting after installation if the number of hosts in your environment changes. See the recommendations in the topic [vCenter Server Hardware Requirements](#).

- 9 Enter the information to register Inventory Service with vCenter Single Sign-On.

The vCenter Single Sign-On administrator user name is `administrator@vsphere.local`, and the password must match the password you entered when you installed vCenter Single Sign-On. The Lookup Service URL takes the form `https://SSO_host_FQDN_or_IP:7444/lookupservice/sdk`, where 7444 is the default vCenter Single Sign-On HTTPS port number. Your entry should match the entry you made when you installed vCenter Single Sign-On. If you entered a different port number when you installed vCenter Single Sign-On, use that port number.

NOTE If you installed vCenter Single Sign-On in a vCenter Server Appliance, you can also enter the Single Sign-On administrator user as `root@localos`. In this case, the password is the root password of the vCenter Server Appliance. The Lookup Service URL takes the form `https://vCenter_Appliance_IP_or_host_name:{7444}/lookupservice/sdk`.

- 10 Click **Install Certificates**.
- 11 Click **Install**.

Inventory Service is installed.

Install vCenter Server as Part of a Custom Install

You can install vCenter Server separately from vCenter Single Sign-On and vCenter Inventory Service to customize the location and configuration of the components.

These instructions let you install vCenter Server only. For most deployments, you can install vCenter Server, vCenter Single Sign-On, and Inventory Service together on a single host machine using the vCenter Server Simple Install option. See [“vCenter Single Sign-On Deployment Modes,”](#) on page 52 and [“Install vCenter Single Sign-On, the vSphere Web Client, vCenter Inventory Service, and vCenter Server by Using Simple Install,”](#) on page 73.

If you do not enter a license key, vCenter Server will be in evaluation mode, which allows you to use the full feature set for a 60-day evaluation period. After installation, you can enter the license key to convert vCenter Server to licensed mode.

NOTE vCenter Server 5.5 supports connection between vCenter Server and vCenter Server components by IP address only if the IP address is IPv4-compliant. To connect to a vCenter Server system in an IPv6 environment, you must use the fully qualified domain name (FQDN) or host name of the vCenter Server. The best practice is to use the FQDN, which works in all cases, instead of the IP address, which can change if assigned by DHCP.

Prerequisites

- Review [“vCenter Single Sign-On Deployment Modes,”](#) on page 52.
- Review [“How vCenter Single Sign-On Affects vCenter Server Installation,”](#) on page 51.
- Review [“Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server,”](#) on page 49
- Install vCenter Single Sign-On and Inventory Service.
- To install the vCenter Server on a drive other than C:, verify that there is enough space in the C: drive to install the Microsoft Windows Installer .msi file.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vCenter Server** and click **Install**.
- 3 Follow the prompts in the installation wizard to choose the installer language, agree to the end user patent and license agreements, and enter your license key.
- 4 Choose the type of database that you want to use.
 - To use the bundled database, click **Install a Microsoft SQL Server 2008 Express instance (for small-scale deployments: up to 5 hosts and 50 virtual machines)**.
 - To use an existing database, click **Use an existing supported database** and select your database from the list of available DSNs. Enter the user name and password for the DSN.

If your database uses Windows NT authentication, the user name and password fields are disabled.

NOTE You might get a warning that the DSN points to an older version of a repository that must be upgraded. If you click **Yes**, the installer upgrades the database schema, making the database irreversibly incompatible with previous VirtualCenter versions. See the *vSphere Upgrade* documentation.

- 5 If the installer prompts you, enter the JDBC URL for your existing vCenter Server database.

The installer should generate and validate the JDBC URL for the vCenter Server database. If the installer fails to connect to the database by using the generated JDBC URL, the installer prompts you to specify the JDBC URL.

- 6 Set the vCenter Server service account information.

- If you are using a nonbundled database, enter the administrator name and password that you use when you log in to the system on which you are installing vCenter Server.
- If you are using the bundled SQL Server database, select **Use Windows Local System Account**.

You need the user name and password to log in to vCenter Server after you install it.

The Fully Qualified Domain Name text box displays the FQDN of the system that you are installing vCenter Server on. The vCenter Server installer checks that the FQDN is resolvable. If not, a warning message appears when you click **Next**. Change the entry to a resolvable FQDN. You must enter the FQDN, not the IP address.

- 7 Select **Create a standalone VMware vCenter Server instance** or **Join a VMware vCenter Group using Linked Mode to share information**.

Joining a Linked Mode group enables the vSphere Web Client to view, search, and manage data across multiple vCenter Server systems.

NOTE This option does not appear if you are upgrading the VirtualCenter or vCenter Server database schema. You can join a Linked Mode group after the installation is complete.

- 8 If you join a group, enter the fully qualified domain name and LDAP port number of any remote vCenter Server system.
- 9 Enter the port numbers that you want to use or accept the default port numbers.
- 10 (Optional) Select **Increase the number of available ephemeral ports**.

- 11 Select the size of your vCenter Server inventory to allocate memory for several Java services that are used by vCenter Server.

This setting determines the maximum JVM heap settings for VMware VirtualCenter Management Webservices (Tomcat), Inventory Service, and Profile-Driven Storage Service. You can adjust this setting after installation if the number of hosts in your environment changes. See the recommendations in the topic vCenter Server Hardware Requirements.

- 12 Enter the information to register vCenter Server with vCenter Single Sign-On.

The vCenter Single Sign-On administrator user name is administrator@vsphere.local, and the password must match the password you entered when you installed vCenter Single Sign-On. The Lookup Service URL takes the form `https://SSO_host_FQDN_or_IP:7444/lookupservice/sdk`, where 7444 is the default vCenter Single Sign-On HTTPS port number. Your entry should match the entry you made when you installed vCenter Single Sign-On. If you entered a different port number when you installed vCenter Single Sign-On, use that port number.

NOTE If you installed vCenter Single Sign-On in a vCenter Server Appliance, you can enter the vCenter Single Sign-On administrator user as root@localos. In this case, the password is the root password of the vCenter Server Appliance. The Lookup Service URL takes the form `https://vCenter_Appliance_IP_or_host_name:{7444}/lookupservice/sdk`.

- 13 If prompted to install or overwrite a certificate, follow the prompt.
- 14 Enter the vCenter Single Sign-On user or group to add as a vCenter Server administrator.

The administrator or group you enter here is granted the necessary privileges to administer the vCenter Server instance that you are installing.

- 15 Enter the Inventory Service URL.

The Inventory Service URL takes the form `https://Inventory_Service_host_FQDN_or_IP:10443`. 10443 is the default Inventory Service HTTPS port number. If you entered a different port number when you installed Inventory Service, use that port number here.

- 16 Either accept the default destination folder or click **Change** to select another location.

The installation path cannot contain any of the following characters: non-ASCII characters, commas (,), periods (.), exclamation points (!), pound signs (#), at signs (@), or percentage signs (%).

- 17 Click **Install**.

Multiple progress bars appear during the installation of the selected components.

- 18 Click **Finish**.

The vCenter Server installation is complete.

What to do next

After you install vCenter Server, you can display the vCenter Server welcome page by typing the IP address of the vCenter Server machine or by typing **localhost** from a browser installed on the vCenter Server machine. You can install the vSphere Web Client to access vCenter Server.

Review the topics in [Chapter 5, “After You Install vCenter Server,”](#) on page 99 for other postinstallation actions you might want to take.

Add a vCenter Single Sign-On Identity Source

Users can log in to vCenter Server only if they are in a domain that has been added as a vCenter Single Sign-On identity source. vCenter Single Sign-On administrator users can add identity sources from the vSphere Web Client.

An identity source can be a native Active Directory (Integrated Windows Authentication) domain or an OpenLDAP directory service. For backward compatibility, Active Directory as an LDAP Server is also available.

Immediately after installation, the following default identity sources and users are available:

locals	All local operating system users. These users can be granted permissions to vCenter Server. If you are upgrading, those users who already have permissions keep those permissions.
vsphere.local	Contains the vCenter Single Sign-On internal users.

Procedure

- 1 Log in to the vSphere Web Client as `administrator@vsphere.local` or as another user with vCenter Single Sign-On administrator privileges.
- 2 Browse to **Administration > Single Sign-On > Configuration**.
- 3 On the **Identity Sources** tab, click the **Add Identity Source** icon.
- 4 Select the type of identity source and enter the identity source settings.

Option	Description
Active Directory (Integrated Windows Authentication)	Use this option for native Active Directory implementations. See “Active Directory Identity Source Settings,” on page 86.
Active Directory as an LDAP Server	This option is available for backward compatibility. It requires that you specify the domain controller and other information. See “Active Directory LDAP Server and OpenLDAP Server Identity Source Settings,” on page 87.

Option	Description
OpenLDAP	Use this option for an OpenLDAP identity source. See “Active Directory LDAP Server and OpenLDAP Server Identity Source Settings,” on page 87.
LocalOS	Use this option to add the local operating system as an identity source. You are prompted only for the name of the local operating system. If you select this option, all users on the specified machine are visible to vCenter Single Sign-On, even if those users are not part of another domain.

NOTE If the user account is locked or disabled, authentications and group and user searches in the Active Directory domain will fail. The user account must have read-only access over the User and Group OU, and must be able to read user and group attributes. This is the default Active Directory domain configuration for user permissions. VMware recommends using a special service user.

- 5 If you configured an Active Directory as an LDAP Server or an OpenLDAP identity source, click **Test Connection** to ensure that you can connect to the identity source.
- 6 Click **OK**.

What to do next

When an identity source is added, all users can be authenticated but have the **No access** permission. A user with vCenter Server **Modify.permissions** privileges can assign permissions to users or groups of users to enable them to log in to vCenter Server. See [“Assign Permissions in the vSphere Web Client,”](#) on page 87.

Active Directory Identity Source Settings

If you select the Active Directory (Integrated Windows Authentication) identity source type, you can either use the local machine account as your SPN (Service Principal Name) or specify an SPN explicitly.

Select **Use machine account** to speed up configuration. If you expect to rename the local machine on which vCenter Single Sign-On runs, specifying an SPN explicitly is preferable.

Table 4-1. Add Identity Source Settings

Field	Description
Domain name	FDQN of the domain. Do not provide an IP address in this field.
Use machine account	Select this option to use the local machine account as the SPN. When you select this option, you specify only the domain name. Do not select this option if you expect to rename this machine.
Use SPN	Select this option if you expect to rename the local machine. You must specify an SPN, a user who can authenticate with the identity source, and a password for the user.
Service Principal	SPN that helps Kerberos to identify the Active Directory service. Include the domain in the name, for example, STS/example.com. You might have to run <code>setspn -S</code> to add the user you want to use. See the Microsoft documentation for information on <code>setspn</code> . The SPN must be unique across the domain. Running <code>setspn -S</code> checks that no duplicate is created.

Table 4-1. Add Identity Source Settings (Continued)

Field	Description
User Principal Name	Name of a user who can authenticate with this identity source. Use the email address format, for example, jchin@mydomain.com. You can verify the User Principal Name with the Active Directory Service Interfaces Editor (ADSI Edit).
Password	Password for the user who is used to authenticate with this identity source, which is the user who is specified in User Principal Name. Include the domain name, for example, jdoe@example.com.

Active Directory LDAP Server and OpenLDAP Server Identity Source Settings

The Active Directory as an LDAP Server identity source is available for backward compatibility. Use the Active Directory (Integrated Windows Authentication) option for a setup that requires less input. The OpenLDAP Server identity source is available for environments that use OpenLDAP.

Table 4-2. Active Directory as an LDAP Server and OpenLDAP Settings

Field	Description
Name	Name of the identity source.
Base DN for users	(Optional) Base domain name for users.
Domain name	FDQN of the domain, for example, example.com. Do not provide an IP address in this field.
Domain alias	The domain's NetBIOS name. Add the NetBIOS name of the Active Directory domain as an alias of the identity source if you are using SSPI authentications.
Base DN for groups	(Optional) The base domain name for groups.
Primary Server URL	Primary domain controller LDAP server for the domain. Use the format ldap://hostname:port or ldaps://hostname:port. The port is typically 389 for ldap: connections and 636 for ldaps: connections. For Active Directory multi-domain controller deployments, the port is typically 3268 for ldap: connections and 3269 for ldaps: connections. A certificate that establishes trust for the LDAPS endpoint of the Active Directory server is required when you use ldaps:// in the primary or secondary LDAP URL.
Secondary server URL	(Optional) Address of a secondary domain controller LDAP server that is used for failover.
Username	ID of a user in the domain who has a minimum of read-only access to Base DN for users and groups.
Password	Password of the user who is specified by Username.

Assign Permissions in the vSphere Web Client

After you create users and groups and define roles, you must assign the users and groups and their roles to the relevant inventory objects. You can assign the same permissions at one time on multiple objects by moving the objects to a folder and setting the permissions on the folder.

Permissions assigned from the vSphere Web Client must match permissions, including case, in ActiveDirectory precisely. If you upgraded from earlier versions of vSphere, check for case inconsistencies if you experience problems with groups.

Prerequisites

Permissions. **Modify permission** on the parent object of the object whose permissions you want to modify.

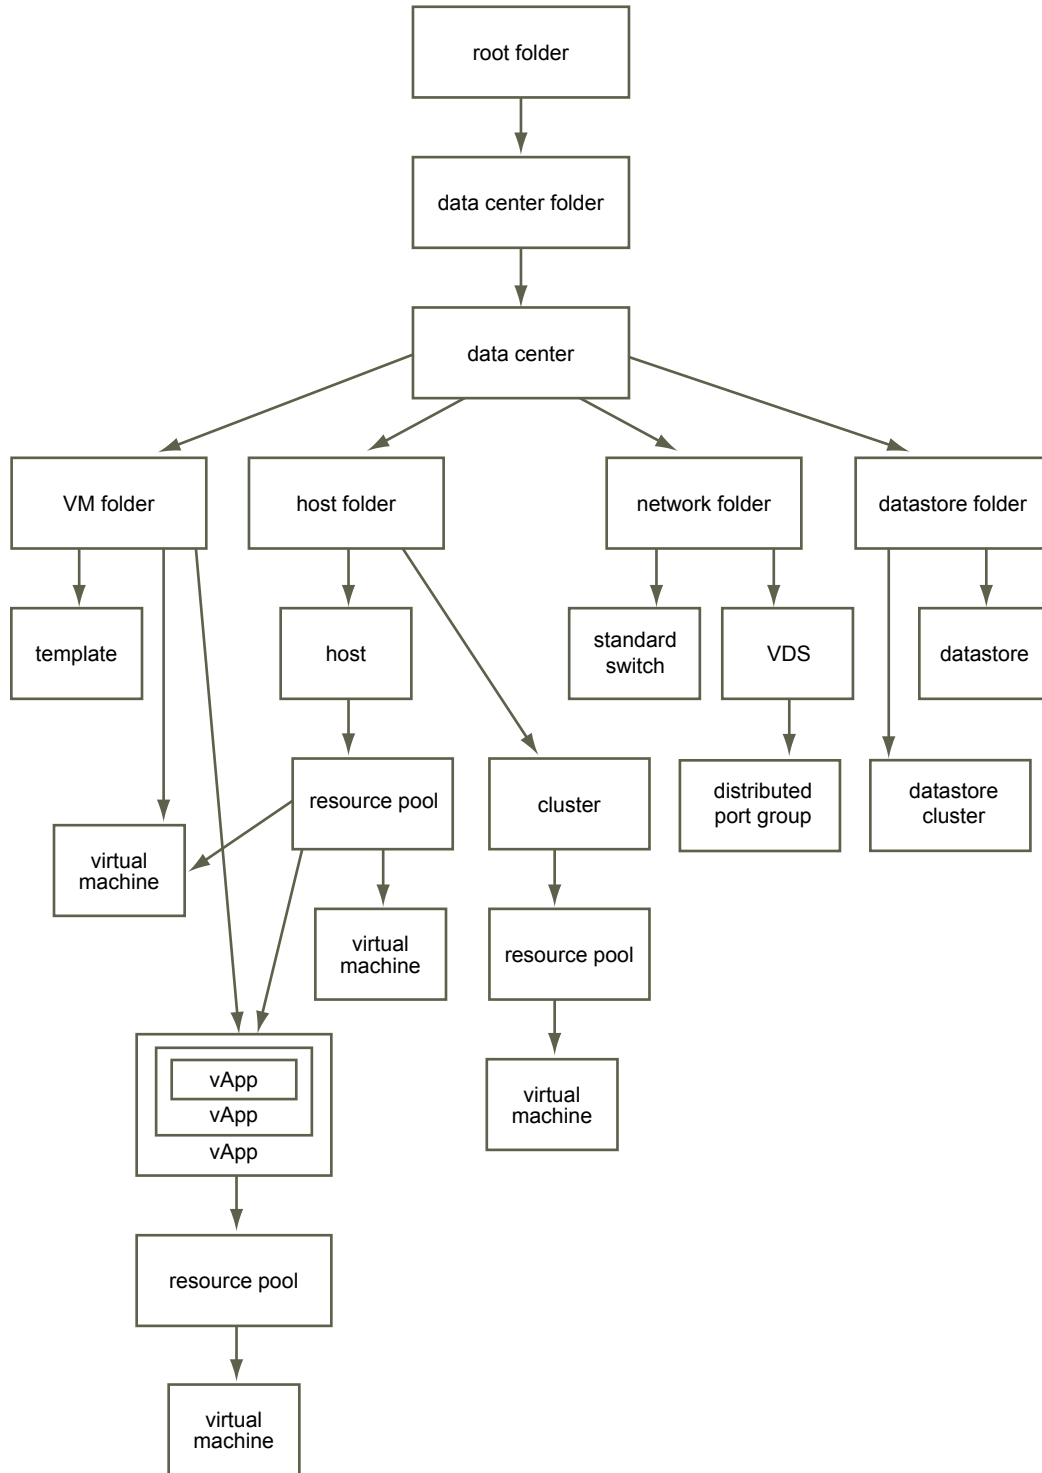
Procedure

- 1 Browse to the object in the vSphere Web Client object navigator.
- 2 Click the **Manage** tab and select **Permissions**.
- 3 Click **Add Permission**.
- 4 Click **Add**.
- 5 Identify the user or group that will have the permission.
 - a Select the domain where the user or group is located from the **Domain** drop-down menu.
 - b Type a name in the Search box or select a name from the list.
The system searches user names, group names, and descriptions.
 - c Select the user or group and click **Add**.
The name is added to either the **Users** or **Groups** list.
 - d (Optional) Click **Check Names** to verify that the user or group exists in the database.
 - e Click **OK**.
- 6 Select a role from the **Assigned Role** drop-down menu.
The roles that are assigned to the object appear in the menu. The privileges contained in the role are listed in the section below the role title.
- 7 (Optional) Deselect the **Propagate to Child Objects** check box.
The role is applied only to the selected object and does not propagate to the child objects.
- 8 Verify that the users and groups are assigned to the appropriate permissions and click **OK**.
The server adds the permission to the list of permissions for the object.
The list of permissions references all users and groups that have roles assigned to the object and indicates where in the vCenter Server hierarchy the role is assigned.

Hierarchical Inheritance of Permissions

When you assign a permission to an object, you can choose whether the permission propagates down the object hierarchy. You set propagation for each permission. Propagation is not universally applied. Permissions defined for a child object always override the permissions that are propagated from parent objects.

The figure illustrates inventory hierarchy and the paths by which permissions can propagate.

Figure 4-2. vSphere Inventory Hierarchy

Most inventory objects inherit permissions from a single parent object in the hierarchy. For example, a datastore inherits permissions from either its parent datastore folder or parent datacenter. Virtual machines inherit permissions from both the parent virtual machine folder and the parent host, cluster, or resource pool simultaneously. To restrict a user's privileges on a virtual machine, you must set permissions on both the parent folder and the parent host, cluster, or resource pool for that virtual machine.

To set permissions for a distributed switch and its associated distributed port groups, set permissions on a parent object, such as a folder or datacenter. You must also select the option to propagate these permissions to child objects.

Permissions take several forms in the hierarchy:

Managed entities

You can define permissions on managed entities.

- Clusters
- Datacenters
- Datastores
- Datastore clusters
- Folders
- Hosts
- Networks (except vSphere Distributed Switches)
- Distributed port groups
- Resource pools
- Templates
- Virtual machines
- vSphere vApps

Global entities

Global entities derive permissions from the root vCenter Server system.

- Custom fields
- Licenses
- Roles
- Statistics intervals
- Sessions

Install or Upgrade vCenter Server Java Components Separately

The required vCenter Server Java Components (JRE) are installed or upgraded silently when you install or upgrade vCenter Server. You can also install or upgrade vCenter Server Java Components separately.

By using the separate installer, you can update or upgrade JRE to a version that is released asynchronously from vCenter Server releases. If an earlier version of JRE is present on the system, this procedure upgrades the existing JRE version.

Prerequisites

- Verify that Microsoft Windows Installer 3.0 or later is present on your system.
- Download the vCenter Server installer from the VMware downloads page at <http://www.vmware.com/support/> and extract the zip archive.

Procedure

- 1 In Windows Explorer, double-click the file `vCenter_Server_installation_directory/vJRE/VMware-jre.exe`.

The VMware vCenter Server - Java Components installer wizard opens.

- 2 Accept the license agreement.

- 3 Accept or change the default installation folder.
- 4 Click **Install**.

The vCenter Server Java Components (JRE) are installed or upgraded.

Install or Upgrade vCenter Server tc Server Separately

The required vCenter Server component tc Server is installed or upgraded silently when you install or upgrade vCenter Server. You can also install or upgrade vCenter Server tc Server separately.

By using the separate installer, you can update or upgrade vCenter Server tc Server to a version that is released asynchronously from vCenter Server releases. If an earlier version of vCenter Server tc Server is present on the system, this procedure upgrades the existing JRE version.

Prerequisites

- Verify that Microsoft Windows Installer 3.0 or later is present on your system.
- Download the vCenter Server installer from the VMware downloads page at <http://www.vmware.com/support/> and extract the zip archive.

Download the vCenter Server installer from the VMware downloads page at <http://www.vmware.com/support/> and extract the zip archive.

Procedure

- 1 In Windows Explorer, double-click the file `vCenter_Server_installation_directory/vtcServer/VMware-tcserver.exe`.
The VMware vCenter Server - tc Server installer wizard opens.
- 2 Accept the license agreement.
- 3 Accept or change the default installation folder.
- 4 Click **Install**.

vCenter Server tc Server is installed or upgraded.

vCenter Single Sign-On Installation Fails

In a Windows environment, vCenter Single Sign-On installation might fail for several reasons.

Problem

The vCenter Single Sign-On installation fails in a Windows environment.

Cause

Multiple causes of an installation failure.

Solution

- 1 Verify that all installation setup prerequisites are met.
At the time the installation fails, the installer displays a message similar to ####: Installation failed due to....
- 2 At a command line, run the following command to gather a vCenter Single Sign-On support bundle.
`C:\Windows\System32\cscript.exe "SSO_Server\scripts\sso-support.wsf" /z`
- 3 Click **OK**

- 4 View the logs in %TEMP%\vminst.log for details about the failure and possible solutions.

For a complete list of logs, see VMware Knowledge Base article [2033430](#).

Download and Deploy the VMware vCenter Server Appliance

As an alternative to installing vCenter Server on a Windows machine, you can download the VMware vCenter Server Appliance. The vCenter Server Appliance is a preconfigured Linux-based virtual machine optimized for running vCenter Server and associated services.

The vCenter Server Appliance has the following default user names

- **root@localos** with the password **vmware**.
- **administrator@vsphere.local** with the password that you entered during installation for the vCenter Single Sign-On administrator account in the built-in domain.

You can also create a custom password that the vCenter Server Appliance reads on first boot. See “[Create a Custom Password on the First Boot for the vCenter Server Appliance](#),” on page 94.

NOTE On a newly deployed version 5.5 vCenter Server Appliance, the root user does not have vCenter Single Sign-On Administrator privileges. For information about administering vCenter Single Sign-On, see the *vSphere Security* documentation.

For external databases, the vCenter Server Appliance supports only Oracle databases, in the same versions shown in the VMware Product Interoperability Matrix for the version of the vCenter Server Appliance that you are deploying. See the VMware Product Interoperability Matrix at http://www.vmware.com/resources/compatibility/sim/interop_matrix.php. Versions 5.0.1 and later of the vCenter Server Appliance use PostgreSQL for the embedded database instead of IBM DB2, which is used in vCenter Server Appliance 5.0.

The vCenter Server Appliance does not support Linked Mode configuration.

Version 5.5 of the vCenter Server Appliance is deployed with virtual hardware version 7, which supports eight virtual CPUs per virtual machine in ESXi. Depending on the hosts that you will manage with the vCenter Server Appliance, you might want to upgrade the ESXi hosts and update the hardware version of the vCenter Server Appliance to support more virtual CPUs:

- ESXi 4.x supports up to virtual hardware version 7 with up to 8 virtual CPUs per virtual machine.
- ESXi 5.0.x supports up to virtual hardware version 8 with up to 32 virtual CPUs per virtual machine.
- ESXi 5.1.x supports up to virtual hardware version 9 with up to 64 virtual CPUs per virtual machine.



CAUTION If you update the vCenter Server appliance to hardware version 10, you cannot edit the virtual machine settings for the appliance using the vSphere Client. This might cause difficulties in managing the vCenter Server Appliance, because you cannot use the vSphere Web Client to connect directly to the host on which the vCenter Server Appliance resides to manage it. Do not upgrade the vCenter Server Appliance to hardware version 10.

For inventory and other configuration limits in the vCenter Server Appliance, see the *Configuration Maximums* documentation.

For information about upgrading ESXi hosts, see the *vSphere Upgrade* documentation. To update the virtual hardware version of a virtual machine, see the information about virtual machine compatibility levels in the *vSphere Virtual Machine Administration* documentation.

For information about configuring the vCenter Server Appliance, see the *vCenter Server and Host Management* documentation.

To configure networking for the vCenter Server Appliance, you must use the `vami_config_net` tool, located at `/opt/vmware/share/vami/vami_config_net`. Instructions for using this tool are in the document *User's Guide to Deploying vApps and Virtual Appliances*, located at https://www.vmware.com/support/developer/studio/studio26/va_user.pdf.

NOTE vCenter Server 5.5 supports connection between vCenter Server and vCenter Server components by IP address only if the IP address is IPv4-compliant. To connect to a vCenter Server system in an IPv6 environment, you must use the fully qualified domain name (FQDN) or host name of the vCenter Server. The best practice is to use the FQDN, which works in all cases, instead of the IP address, which can change if assigned by DHCP.

The ESXi Dump Collector service, the vSphere Syslog service, and vSphere Auto Deploy must use an IPv4 address to communicate with the vCenter Server Appliance.

Prerequisites

- Verify that the host machine meets the hardware requirements for the vCenter Server Appliance listed in “[Hardware Requirements for vCenter Server, the vSphere Web Client, vCenter Inventory Service, and vCenter Single Sign-On](#),” on page 17.
- Verify that the hosts are running ESX version 4.x or ESXi version 4.x or later.
- Synchronize the clocks of all machines on the vSphere network. See “[Synchronizing Clocks on the vSphere Network](#),” on page 59.
- If you plan to configure the vCenter Server Appliance from a configuration file, prepare the file and store it in a location that is accessible from the vCenter Server Appliance host machine. See “[Format for the vCenter Server Appliance Configuration File](#),” on page 96.
- See the *vSphere Virtual Machine Administration* documentation for instructions on deploying OVA files and OVF templates.

Procedure

- 1 From the vSphere 5 download page on the VMware Web site, download the .OVA file or the .OVF and VMDK System and Data disks for the vCenter Server appliance onto your system.

If you use the .OVF and VMDK System and Data disks, download them all into the same folder.

- 2 Using the vSphere Client or vSphere Web Client, deploy the .OVA file or the .OVF and VMDK system and data disks as an OVF template.

If you do not want to commit to using the maximum 80GB of disk space at deployment, deploy the vCenter Server Appliance with thin provisioning. In the Disk Format panel of the Deploy OVF template wizard, select **Thin provisioned format**.

NOTE The vSphere Web Client connects to vSphere through the vCenter Server, but cannot connect directly to an ESXi instance. If your vSphere environment does not already have vCenter Server installed, use the vSphere Client to deploy the vCenter Server Appliance to an ESXi instance. Alternatively, you can use the OVF Tool, which is available, with documentation, from VMware Communities at <https://communities.vmware.com/community/vmtn/automationtools/ovf>. Once your vSphere environment includes a vCenter Server instance, you can use the vSphere Web Client to deploy the vCenter Server Appliance to an ESXi instance through the vCenter Server.

- 3 Power on the vCenter Server Appliance.
- 4 Open a console view.
- 5 Follow the instructions on the welcome screen to open a browser window to the URL shown.
- 6 Log in to the vCenter Server Appliance and accept the license agreement.

When you log in, the vCenter Server Setup wizard starts.

- 7 Select the configuration option for your installation.

Option	Description
Configure with default settings	Sets up embedded vCenter Server database in the vCenter Server Appliance and configures the database and Active Directory with default settings.
Upload configuration file	To configure the vCenter Server Appliance from a prepared configuration file.
Set custom configuration	To customize the configuration of the vCenter Server Appliance. The setup wizard displays separate panels for you to connect the appliance to embedded or external vCenter Server database, and to configure custom Active Directory settings.

- 8 Follow the prompts to complete the wizard.

If you uploaded a configuration file, enter any settings that were not included in the file as you complete the wizard.

The vCenter Server Appliance is deployed and set up.

What to do next

See the *vCenter Server and Host Management* documentation for information about using vCenter Server and the vCenter Server Appliance.

Create a Custom Password on the First Boot for the vCenter Server Appliance

The vCenter Server Appliance has the default user name **root@localos** and password **vmware**. You can also create a custom password that the vCenter Server Appliance reads the first time you boot the appliance.

By creating a custom password the first time you start the vCenter Server Appliance, you ensure that the default password cannot be used.

NOTE When you change the root vCenter Server Appliance password using the vCenter Server Appliance Web interface, the GRUB password is changed automatically. However, when you change the vCenter Server Appliance password using the vCenter Server Appliance console, the GRUB password is not changed. You must update the vCenter Server Appliance root password using the Web interface to simultaneously change the GRUB password.

Procedure

- 1 From a terminal window on your Linux host machine, type the following command to create an MD5 hash of the custom password.

```
grub-md5-crypt
```

- 2 At the prompt, type the new password and press Enter.

The system returns the MD5 hash of the password.

- 3 Create a CD directory.

```
mkdir cd
```

- 4 Add the MD5 hash to the vCenter Server Appliance configuration file.

```
echo 'rootPwdHashMD5=hash_password' > cd/vcva.cfg
```

hash_password is the MD5 hash of the password returned in [Step 2](#). Use single quotes for the echo command, as shown, because the *hash_password* contains \$ characters that must be escaped otherwise.

- 5 Create an ISO file containing the password.

```
mkisofs -R -o rootpass.iso cd
```

- 6 Attach the CD/DVD drive of the vCenter Server Appliance virtual machine to the `rootpass.iso` file, and make sure that **Connected at power on** is selected.

When you turn on the vCenter Server Appliance, it reads and applies the custom password you created for the root user.

Configure a vCenter Server Appliance to Use the vCenter Single Sign-On of a Different Virtual Machine

You can direct one or more vCenter Server Appliances to use a vCenter Single Sign-On instance running on another virtual machine. This action makes all the vCenter Server instances accessible by each vSphere Web Client configured with the same vCenter Single Sign-On instance.

You can also set a new vCenter Server Appliance to use an external vCenter Single Sign-On instance when you run the vCenter Server Setup wizard for a newly deployed vCenter Server Appliance. See [“Download and Deploy the VMware vCenter Server Appliance,”](#) on page 92.

Prerequisites

- Verify that you have the URL of the Lookup Service for the target instance of vCenter Single Sign-On.

Procedure

- 1 Connect to the vCenter Server Appliance from a Web browser.
- 2 On the **vCenter Server** tab, click **SSO**.
- 3 Set the **SSO deployment type** to **external**.
- 4 Under Lookup service location, type the **URL** of the Lookup Service for the target instance of vCenter Single-Sign On.

The new vCenter Server Appliance is configured to use the existing vCenter Single Sign-On instance, and both vCenter Server instances are accessible by the vSphere Web Client of either vCenter Server Appliance.

What to do next

You can repeat this procedure for multiple vCenter Server Appliances.

Format for the vCenter Server Appliance Configuration File

When you deploy a new vCenter Server Appliance, you can configure the appliance by uploading a configuration file, instead of entering the settings manually in the vCenter Server Setup wizard.

Configuration File Format

Each line of the configuration file supplies the setting for the corresponding entry in the vCenter Center Setup wizard. The values shown here are variables that describe acceptable entries.

NOTE You must enter the header exactly as shown below: **vCenter Server Preseed Config v1.0**. Otherwise, the file will be rejected.

Passwords in the configuration file are not used. You must enter passwords manually when you complete the vCenter Setup wizard.

```
#vCenter Server Preseed Config v1.0
```

```
#Database options
```

```
VC_DB_TYPE=embedded or oracle
```

```
#Additional database options if VC_DB_TYPE=oracle
```

```
VC_DB_SERVER=vCenter Server database server IP address
```

```
VC_DB_SERVER_PORT=vCenter Server database server port number
```

```
VC_DB_INSTANCE=vCenter Server database server instance name
```

```
VC_DB_USER=vCenter Server database user name
```

```
VC_DB_PASSWORD=
```

```
#vCenter Single Sign-On options
```

```
SSO_TYPE=embedded or external
```

```
#Additional vCenter Single Sign-On options if SSO_TYPE=external
```

```
SSO_LS_LOCATION=vCenter Single Sign-On Lookup Service URL
```

```
SSO_LS_CERT_THUMBPRINT=vCenter Single Sign-On Lookup Service certificate thumbprint
```

```
SSO_REG_ADMIN_USERNAME=fully qualified name of a user with vCenter Single Sign-On administrative privileges
```

```
SSO_REG_ADMIN_PASSWORD=password for above user with vCenter Single Sign-On administrative privileges
```

```
SSO_VC_ADMIN_USERNAME=name of the default principal administrator for vCenter Server. If not qualified, it is
```

```
assumed to be an Operating System principal
```

```
SSO_VC_ADMIN_IS_GROUP=true or false. True if above default principal administrator for vCenter Server is a group.
```

```
#Additional vCenter Single Sign-On options if SSO_TYPE=embedded
```

```
SSO_ADMIN_PASS=password for the administrator@vsphere.local newly installed embedded vCenter Single Sign-On service
```

```
SSO_ADMIN_PASS2=set to the same string as SSO_ADMIN_PASS
```

```
SSO_CALLTYPE=set to "typed" for SSO_ADMIN_PASS to take effect.
```

```
#Microsoft Active Directory options
```

```
VC_AD_STATUS=0 to disable, 1 to enable
```

```
#Additional Microsoft Active Directory option if VC_AD_STATUS=1
```

```
VC_AD_DOMAIN=Active Directory fully qualified domain name
```


Example: Example Configuration File

This example shows a configuration file for a vCenter Server Appliance that uses an external Oracle vCenter Server database and an external vCenter Single Sign-On instance.

```
#vCenter Server Preseed Config v1.0
```

```
#Database options
```

```
VC_DB_TYPE=oracle
```

```
VC_DB_SERVER=10.111.11.111
```

```
VC_DB_SERVER_PORT=1521
```

```
VC_DB_INSTANCE=orcl
```

```
VC_DB_USER=VCA-6
```

```
VC_DB_PASSWORD=
```

```
#vCenter Single Sign-On options
```

```
SSO_TYPE=external
```

```
SSO_LS_LOCATION=https://machinename.corp.com:7444/lookupservice/sdk
```

```
SSO_LS_CERT_THUMBPRINT=
```

```
#Microsoft Active Directory options
```

```
VC_AD_STATUS=0
```

```
VC_AD_DOMAIN=
```


After You Install vCenter Server

After you install vCenter Server, consider these postinstallation options before adding inventory for the vCenter Server to manage.

General Requirements

- Install the vSphere Web Client and make sure that you can access the vCenter Server instance.
- When vCenter Server and the database are installed on the same machine, after you reboot the machine, you might need to restart the VMware VirtualCenter Management Webservices service. See [“VMware vCenter Management Webservices Service Fails to Start,”](#) on page 116.
- Review the subtopics in this section for other postinstallation options.

Oracle Database Requirements

- For the Oracle Instant client, copy ojdbc14.jar to the vCenter Server tomcat directory *vCenter install location\Infrastructure\tomcat\lib*.
- The Oracle 10g client and Oracle 11g client include ojdbc14.jar at: *Install location\oracle\product\10.2.0\instance_name\jdbc\lib* or *Install location\app\Administrator\product\11.1.0\instance_name\sqldeveloper\jdbc\lib*. The vCenter Server installer copies the file from the Oracle client installation location to the vCenter Server tomcat directory *vCenter install location\Infrastructure\tomcat\lib*.
- If neither the ojdbc14.jar file nor the ojdbc5.jar file is found in the Oracle 10g or Oracle 11g client location, the vCenter Server installer prompts you to copy the file manually. You can download the file from the Oracle.com Web site.
- After you install or upgrade the Oracle database and, if necessary, copy the ojdbc14.jar file or ojdbc5.jar file manually, restart the VMware VirtualCenter Management Webservices service. The VMware VirtualCenter Management Webservices service is located in the Services section of the Windows Administrative Tools control panel.

Security Requirement

For environments that require strong security, VMware recommends that you replace the default certificates on your vCenter Server system with certificates signed by a commercial Certificate Authority (CA). See the information on increasing security for session information in the *vSphere Examples and Scenarios* documentation.

This chapter includes the following topics:

- [“Install vCenter Server Components,”](#) on page 100
- [“Creating vCenter Server Linked Mode Groups,”](#) on page 109
- [“Configuring VMware vCenter Server - tc Server Settings in vCenter Server,”](#) on page 114
- [“VMware vCenter Management Webservices Service Fails to Start,”](#) on page 116
- [“Back Up the Inventory Service Database on Windows,”](#) on page 116
- [“Restore an Inventory Service Database Backup on Windows,”](#) on page 116
- [“Back Up the Inventory Service Database on Linux,”](#) on page 117
- [“Restore an Inventory Service Database Backup on Linux,”](#) on page 117
- [“Reset the vCenter Inventory Service Database,”](#) on page 118
- [“Enable IPv6 Support for vCenter Inventory Service,”](#) on page 119

Install vCenter Server Components

You can install vCenter Server components on the same machine that hosts vCenter Server or on remote machines.

Install the Client Integration Plug-In in the vSphere Web Client

The Client Integration Plug-in provides access to a virtual machine's console in the vSphere Web Client, and provides access to other vSphere infrastructure features.

You use the Client Integration Plug-in to deploy OVF or OVA templates and transfer files with the datastore browser. You can also use the Client Integration Plug-in to connect virtual devices that reside on a client computer to a virtual machine.

Install the Client Integration Plug-in only once to enable all the functionality the plug-in delivers. You must close the Web browser before installing the plug-in.

If you install the Client Integration 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 Client Integration 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.

You cannot launch the virtual machine console in Internet Explorer without the Client Integration Plug-in. In other supported browsers, the virtual machine console can run without the plug-in.

The Client Integration Plug-in also lets you log in to the vSphere Web Client by using Windows session credentials.

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 In the vSphere Web Client, navigate to a link to download the Client Integration Plug-in.

Option	Description
vSphere Web Client login page	<ol style="list-style-type: none"> a Open a Web browser and type the URL for the vSphere Web Client. b At the bottom of the vSphere Web Client login page, click Download Client Integration Plug-in.
Guest OS Details panel	<p>This option is not available for browsers that run on a Mac OS.</p> <ol style="list-style-type: none"> a Select a virtual machine in the inventory and click the Summary tab. b Click Download Plug-in.
OVF deployment wizard	<ol style="list-style-type: none"> a Select a host in the inventory and select Actions > All vCenter Actions > Deploy OVF Template. b Click Download Client Integration Plug-in.
Virtual machine console	<p>This option is not available for Microsoft Internet Explorer, and for browsers that run on a Mac OS.</p> <ol style="list-style-type: none"> a Select a virtual machine in the inventory, click the Summary tab, and click Launch Console. b At the top right corner of the virtual machine console window, click Download Client Integration Plug-in.

- 2 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.

Install or Upgrade the vSphere Web Client

The vSphere Web Client lets you connect to a vCenter Server system to manage your vSphere deployment through a browser.

If an earlier version of the vSphere Web Client is installed, this procedure upgrades the vSphere Web Client.

NOTE vCenter Server 5.5 supports connection between vCenter Server and vCenter Server components by IP address only if the IP address is IPv4-compliant. To connect to a vCenter Server system in an IPv6 environment, you must use the fully qualified domain name (FQDN) or host name of the vCenter Server. The best practice is to use the FQDN, which works in all cases, instead of the IP address, which can change if assigned by DHCP.

Prerequisites

- Download the vCenter Server installer. See [“Download the vCenter Server Installer,”](#) on page 73.
- Verify that the system has an Internet connection.
- Verify that the system meets the software requirements for the vSphere Web Client. See [“vSphere Web Client Software Requirements,”](#) on page 22.
- Before you install or upgrade any vSphere product, synchronize the clocks of all machines on the vSphere network. See [“Synchronizing Clocks on the vSphere Network,”](#) on page 59.
- Install vCenter Single Sign-On, or upgrade to the current version.
- Verify that the vSphere Web Client and vCenter Server are registered to the same vCenter Single Sign-On server, to ensure that the vSphere Web Client can access the vCenter Server inventory.
- Close all browsers before installing or uninstalling the vSphere Web Client.
- Log in as a member of the Administrators group on the host machine, with a user name that does not contain any non-ASCII characters.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vSphere Web Client** and click **Install**.
- 3 Follow the prompts in the installation wizard to choose the installer language, and agree to the end user patent and license agreements.
- 4 Either accept the default destination folder or click **Change** to select another location.

The installation path cannot contain any of the following characters: non-ASCII characters, commas (,), periods (.), exclamation points (!), pound signs (#), at signs (@), or percentage signs (%).

If 8.3 name creation is disabled on the host machine, do not install the vSphere Web Client in a directory that does not have an 8.3 short name or has a name that contains spaces. This situation will make the vSphere Web Client inaccessible.

- 5 Accept or change the default port settings.
- 6 Enter the information to register the vSphere Web Client with vCenter Single Sign-On.

The vCenter Single Sign-On administrator user name is `administrator@vsphere.local`, and the password must match the password you entered for the administrator user when you installed vCenter Single Sign-On. The Lookup Service URL takes the form `https://SSO_host_FQDN_or_IP:7444/lookupservice/sdk`, where 7444 is the default vCenter Single Sign-On HTTPS port number. Your entry should match the entry you made when you installed vCenter Single Sign-On. If you entered a different port number when you installed vCenter Single Sign-On, use that port number.

- 7 Click **Install**.
- 8 Start the vSphere Web Client by taking one of the following actions.
 - If you are starting the vSphere Web Client for the first time, open a supported browser, and go to `https://vSphere_Web_Client_host_name_or_IP:9443/vsphere-client`.
 - In subsequent sessions, you can start the vSphere Web Client from the Windows Start menu, by selecting **Programs > VMware > VMware vSphere Web Client > vSphere Web Client**.

NOTE After you upgrade the vSphere Web Client, when you log in for the first time, you may see the error message `Failed to navigate to desired location`. This can happen when a vSphere Web Client session from the previous version remains open when you upgrade. In this case, refresh the browser and log in again.

What to do next

Install the Client Integration Plug-In in the vSphere Web Client. See [“Install the Client Integration Plug-In in the vSphere Web Client,”](#) on page 100

Install a Local Copy of vSphere Web Client Help

If you do not have internet access from the system you use to access the vSphere Web Client, you can download and deploy a local copy of the online Help.

By default, vSphere Web Client accesses online Help on the Web. This allows the client to access the most up-to-date version of the Help content.

If you download and deploy Help locally, the local copy is not updated when new Help is published to the Web. If you deploy local Help, check the download location periodically for updates.

For instructions for downloading and deploying vSphere Web Client online Help locally, see <http://kb.vmware.com/kb/2030344>.

Install the Update Manager Server

The Update Manager installation requires a connection with a single vCenter Server instance. You can install Update Manager on the same computer on which vCenter Server is installed or on a different computer.

Prerequisites

See Update Manager installation prerequisites in *Installing and Administering VMware vSphere Update Manager*.

Check the compatibility and interoperability of the vCenter Server server with SRM. You should use caution when connecting the Update Manager server to a vCenter Server instance to which the SRM server is connected. Connecting the Update Manager server to the same vCenter Server instance as SRM might cause problems when you upgrade SRM or vSphere, and when you perform daily operations.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file and select **vSphere Update Manager**.

If you cannot run `autorun.exe`, browse to the `UpdateManager` folder and run `VMware-UpdateManager.exe`.

- 2 Select a language for the installer and click **OK**.
- 3 Review the Welcome page and click **Next**.
- 4 Read the patent agreement and click **Next**.
- 5 Accept the terms in the license agreement and click **Next**.
- 6 Review the support information, select whether to download updates from the default download sources immediately after installation, and click **Next**.

If you deselect **Download updates from default sources immediately after installation**, Update Manager downloads updates once daily according to the default download schedule or immediately after you click the **Download Now** button on the Download Settings page. You can modify the default download schedule after the installation is complete.

If you deselect **Download updates from default sources immediately after installation**, the update download task runs after installation, but it does not download any updates.

- 7 Type the vCenter Server IP address or name, HTTP port, and the administrative account that the Update Manager server will use to connect to the vCenter Server system, and click **Next**.
- 8 Select the type of database that you want to use.
 - If you do not have an existing database, select **Install a Microsoft SQL Server 2008 R2 Express instance (for small scale deployments)** and click **Next**.
This database is suitable for deployments of up to 5 hosts and 50 virtual machines.
 - If you have a supported database, select **Use an existing supported database** and select a DSN from the drop-down menu. If the DSN does not use Windows NT authentication, type the user name and password for the DSN and click **Next**.

IMPORTANT The DSN must be a 32-bit DSN.

- 9 (Optional) Select the database options.
 - If the system DSN you specify points to an existing Update Manager database with the current schema, you can either retain your existing database or replace it with an empty one.
 - If the system DSN you specify points to an existing Update Manager database with different schema, on the Database Upgrade page, select **Yes, I want to upgrade my Update Manager database** and **I have taken a backup of the existing Update Manager database**, and click **Next**.
- 10 From the drop-down menu, select the IP address or the host name of your Update Manager instance.
 If the computer on which you install Update Manager has one NIC, the Update Manager installer automatically detects the IP address. If the computer has multiple NICs, you must select the correct IP address or use a DNS name. The DNS name must be resolved from all hosts that this Update Manager instance will manage.
- 11 Specify the Update Manager port settings, select whether you want to configure the proxy settings, and click **Next**.

NOTE Use caution when you specify the Update Manager port settings, as you cannot modify them after installation.

For the SOAP port, you have no limitations to the range of ports used, as long as there are no conflicts.

For the Server port, you can use the following range: 80, 9000-9100. Update Manager automatically opens ESX/ESXi firewall ports in this range to allow outbound HTTP traffic to the patch store.

- 12 (Optional) Provide information about the proxy server, the port, and whether the proxy should be authenticated, and click **Next**.
- 13 Select the Update Manager installation and patch download directories, and click **Next**.
 If you do not want to use the default locations, you can click **Change** to browse to a different directory.
- 14 (Optional) In the warning message about the disk free space, click **OK**.
 This message appears when you try to install Update Manager on a computer that has less than 120GB free space.
- 15 Click **Install** to begin the installation.
- 16 Click **Finish**.

The Update Manager server component is installed, and the client component appears as an available plug-in in the Plug-in Manager of the vSphere Client.

What to do next

In the vSphere Client, select **Plug-ins > Manage Plug-ins** to install and enable the Update Manager Client plug-in.

Install or Upgrade vSphere ESXi Dump Collector

You can configure ESXi to dump the vmkernel memory to a network server, rather than to a disk, when the system has encountered a critical failure. Install vSphere ESXi Dump Collector to collect such memory dumps over the network.

If an earlier version of the vSphere ESXi Dump Collector is installed on your system, this procedure upgrades the vSphere ESXi Dump Collector to the current version.

NOTE In the vCenter Server Appliance, the vSphere ESXi Dump Collector is installed and enabled by default. These instructions apply to Windows-based deployments.

For instructions on configuring ESXi to dump kernel memory to the network server, see [“Configure ESXi Dump Collector with ESXCLI,”](#) on page 187.

The vSphere ESXi Dump Collector is most useful for datacenters where ESXi hosts are configured using the Auto Deploy process, so the ESXi hosts might not have local storage. You can also install the vSphere ESXi Dump Collector for ESXi hosts that do have local storage, as an additional location where vmkernel memory dumps can be redirected when critical failures occur.

You can install the vSphere ESXi Dump Collector on the same machine as the associated vCenter Server, or on a different machine that has network connection to the vCenter Server.

The vSphere ESXi Dump Collector service binds to an IPv4 address for communication with vCenter Server, and does not support IPv6. The vCenter Server can be on a host machine in an IPv4-only, IPv4/IPv6 mixed-mode, or IPv6-only network environment, but the machine that connects to the vCenter Server through the vSphere Web Client must have an IPv4 address for the vSphere ESXi Dump Collector service to work.

Prerequisites

- Verify that you have administrator privileges
- Verify that the host machine has Windows Installer 3.0 or later.
- Verify that the host machine has a supported processor and operating system. The vSphere ESXi Dump Collector supports the same processors and operating systems as vCenter Server. See [“vCenter Server Software Requirements,”](#) on page 22 and [“Hardware Requirements for vCenter Server, the vSphere Web Client, vCenter Inventory Service, and vCenter Single Sign-On,”](#) on page 17.
- Verify that the host machine has a valid IPv4 address. You can install the vSphere ESXi Dump Collector on a machine in an IPv4-only or IPv4/IPv6 mixed-mode network environment, but you cannot install the vSphere ESXi Dump Collector on a machine in an IPv6-only environment.
- If you are using a network location for the Dump Collector repository, make sure the network location is mounted.

Gather the following information to complete the installation or upgrade:

- The location to install the vSphere ESXi Dump Collector to, if you are not using the default location.
- The location for the vSphere ESXi Dump Collector repository where the dump files will be stored.
- (Optional) The maximum size for the vSphere ESXi Dump Collector repository. The specified network location must have at least that much free space.
- Whether to install the vSphere ESXi Dump Collector as a standalone instance or to integrate the vSphere ESXi Dump Collector with a vCenter Server. The vSphere ESXi Dump Collector is not supported for integration with vCenter Server versions earlier than version 5.0.
- If the vSphere ESXi Dump Collector is integrated with a vCenter Server, the address and credentials for the vCenter Server: IP address or name, HTTP port, user name, and password.
- The vSphere ESXi Dump Collector server port, if you are not using the default setting.
- The host name or IP address to identify the vSphere ESXi Dump Collector on the network.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vSphere ESXi Dump Collector** and click **Install**.
- 3 Follow the wizard prompts to complete the installation or upgrade.

Install or Upgrade vSphere Syslog Collector

Install the vSphere Syslog Collector to enable ESXi system logs to be directed to a server on the network, rather than to a local disk.

If an earlier version of vSphere Syslog Collector is installed on your system, this procedure upgrades vSphere Syslog Collector to the current version.

You can install vSphere Syslog Collector on the same machine as the associated vCenter Server, or on a different machine that has network connection to the vCenter Server. The vSphere Syslog Collector service binds to an IPv4 address for communication with vCenter Server, and does not support IPv6. The vCenter Server can be on a host machine in an IPv4-only, IPv4/IPv6 mixed-mode, or IPv6-only network environment, but the machine that connects to the vCenter Server through the vSphere Web Client must have an IPv4 address for the vSphere Syslog Collector service to work.

Prerequisites

- Verify that you have administrator privileges.
- Verify that the host machine has Windows Installer 3.0 or later.
- Verify that the host machine has a supported processor and operating system. vSphere Syslog Collector supports the same processors and operating systems as vCenter Server. See [“vCenter Server Software Requirements,”](#) on page 22 and [“Hardware Requirements for vCenter Server, the vSphere Web Client, vCenter Inventory Service, and vCenter Single Sign-On,”](#) on page 17.
- Determine whether to install vSphere Syslog Collector as a standalone instance or to integrate vSphere Syslog Collector with a vCenter Server. vSphere Syslog Collector is not supported for integration with vCenter Server versions earlier than version 5.0.
- Verify that the host machine has a valid IPv4 address. You can install vSphere Syslog Collector on a machine in an IPv4-only or IPv4/IPv6 mixed-mode network environment, but you cannot install vSphere Syslog Collector on a machine in an IPv6-only environment.

Gather the following information to complete the installation or upgrade:

- The location to install vSphere Syslog Collector, if you are not using the default location.
- The location for the vSphere Syslog Collector repository where the syslog files will be stored.
- (Optional) The maximum size for the vSphere Syslog Collector repository. The specified network location must have at least that much free space.
- (Optional) The maximum number of vSphere Syslog Collector log rotations to keep.
- If vSphere Syslog Collector is integrated with a vCenter Server, the address and credentials for the vCenter Server: IP address or name, HTTP port, user name, and password.
- The vSphere Syslog Collector server port, if you are not using the default setting, and whether to use TCP and UDP protocols for this port.
- The vSphere Syslog Collector server SSL port, if you are not using the default setting, and whether to use secure connection (SSL) for this port.
- The host name or IP address to identify vSphere Syslog Collector on the network.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vSphere Syslog Collector** and click **Install**.
- 3 Follow the wizard prompts to complete the installation or upgrade.

Install or Upgrade vSphere Auto Deploy

Install vSphere Auto Deploy to provision and customize physical hosts by loading the ESXi image directly into memory. You can provision and reprovision hundreds of ESXi hosts efficiently with vCenter Server.

If an earlier version of vSphere Auto Deploy is installed on your system, this procedure upgrades vSphere Auto Deploy to the current version.

You must install the vSphere Auto Deploy server separately for each instance of vCenter Server that you plan to use the vSphere Auto Deploy with. vSphere Auto Deploy is not supported with vCenter Server versions earlier than version 5.0. You must upgrade vSphere Auto Deploy when you upgrade vCenter Server. vSphere Auto Deploy is supported only in the same version as the corresponding vCenter Server.

vSphere Auto Deploy supports both IPv4 and IPv6. However, vSphere Auto Deploy uses a PXE boot infrastructure that supports only IPv4. You can use vSphere Auto Deploy in a mixed IPv4-IPv6 environment or an IPv4-only environment, but not in an IPv6-only environment.

Prerequisites

- Verify that you have administrator privileges
- Verify that the host machine has Windows Installer 3.0 or later.
- Verify that the host machine has a supported processor and operating system. vSphere Auto Deploy supports the same processors and operating systems as vCenter Server.

See [“vCenter Server Software Requirements,”](#) on page 22 and [“Hardware Requirements for vCenter Server, the vSphere Web Client, vCenter Inventory Service, and vCenter Single Sign-On,”](#) on page 17.

Gather the following information to complete the installation or upgrade:

- The location to install vSphere Auto Deploy in, if you are not using the default location.
- The location for the vSphere Auto Deploy repository. Do not use a network share for the repository.
- (Optional) The maximum size for the vSphere Auto Deploy repository. Best practice is to allocate 2GB to have enough room for four image profiles and some extra space. Each image profile requires approximately 350MB. Determine how much space to reserve for the vSphere Auto Deploy repository by considering how many image profiles you expect to use. The specified disk must have at least that much free space.
- The address and credentials of the vCenter Server that you are installing the vSphere Auto Deploy feature for: IP address or name, HTTP port, user name, and password.
- The vSphere Auto Deploy server port, if you are not using the default setting.
- The host name or IP address to identify vSphere Auto Deploy on the network.

Procedure

- 1 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 2 Select **vSphere Auto Deploy** and click **Install**.
- 3 Follow the wizard prompts to complete the installation or upgrade.

Install or Upgrade vSphere Authentication Proxy

Install vSphere Authentication Proxy to enable ESXi hosts to join a domain without using Active Directory credentials. vSphere Authentication Proxy enhances security for PXE-booted hosts and hosts that are provisioned using Auto Deploy, by removing the need to store Active Directory credentials in the host configuration.

If an earlier version of the vSphere Authentication Proxy is installed on your system, this procedure upgrades the vSphere Authentication Proxy to the current version.

You can install vSphere Authentication Proxy on the same machine as the associated vCenter Server, or on a different machine that has network connection to the vCenter Server. The vSphere Authentication Proxy is not supported with vCenter Server versions earlier than version 5.0.

The vSphere Authentication Proxy service binds to an IPv4 address for communication with vCenter Server, and does not support IPv6. The vCenter Server can be on a host machine in an IPv4-only, IPv4/IPv6 mixed-mode, or IPv6-only network environment, but the machine that connects to the vCenter Server through the vSphere Web Client must have an IPv4 address for the vSphere Authentication Proxy service to work.

Prerequisites

- Install vSphere Auto Deploy. See [“Install or Upgrade vSphere Auto Deploy,”](#) on page 107.
- Verify that you have administrator privileges.
- Verify that the host machine has Windows Installer 3.0 or later.
- Verify that the host machine has a supported processor and operating system. vSphere Authentication Proxy supports the same processors and operating systems as vCenter Server. See [“vCenter Server Software Requirements,”](#) on page 22 and [“Hardware Requirements for vCenter Server, the vSphere Web Client, vCenter Inventory Service, and vCenter Single Sign-On,”](#) on page 17.
- Verify that the host machine has a valid IPv4 address. You can install vSphere Authentication Proxy on a machine in an IPv4-only or IPv4/IPv6 mixed-mode network environment, but you cannot install vSphere Authentication Proxy on a machine in an IPv6-only environment.
- If you are installing vSphere Authentication Proxy on a Windows Server 2008 R2 host machine, download and install the Windows hotfix described in Windows KB Article 981506 on the support.microsoft.com Web site. If this hotfix is not installed, the vSphere Authentication Proxy Adapter fails to initialize. This problem is accompanied by error messages in `camadapter.log` similar to `Failed to bind CAM website with CTL` and `Failed to initialize CAMAdapter`.

Gather the following information to complete the installation or upgrade:

- The location to install vSphere Authentication Proxy, if you are not using the default location.
- The address and credentials for the vCenter Server that vSphere Authentication Proxy will connect to: IP address or name, HTTP port, user name, and password.
- The host name or IP address to identify vSphere Authentication Proxy on the network.

Procedure

- 1 On the host machine where you will install the vSphere Authentication Proxy service, install the .NET Framework 3.5.
- 2 Install vSphere Auto Deploy.
You do not have to install Auto Deploy on the same host machine as the vSphere Authentication Proxy service.
- 3 Add the host machine where you will install the authentication proxy service to the domain.
- 4 Use the Domain Administrator account to log in to the host machine.

- 5 In the software installer directory, double-click the `autorun.exe` file to start the installer.
- 6 Select **vSphere Authentication Proxy** and click **Install**.
- 7 Follow the wizard prompts to complete the installation or upgrade.

During installation, the authentication service registers with the vCenter Server instance where Auto Deploy is registered.

When you install the vSphere Authentication Proxy service, the installer creates a domain account with appropriate privileges to run the authentication proxy service. The account name begins with the prefix `CAM-` and has a 32-character, randomly generated password associated with it. The password is set to never expire. Do not change the account settings.

What to do next

Configure ESXi to use vSphere Authentication Proxy to join a domain. See the *vSphere Security* documentation.

Uninstall VMware vSphere Components

The VMware vSphere components and support tools are uninstalled separately, even if they are on the same machine. You must have administrator privileges to uninstall VMware vCenter Server



CAUTION Uninstalling a vCenter Server system while it is running disrupts the vSphere Web Client connections, which can cause data loss.

vCenter Single Sign-On and vCenter Server depend on vCenter Server Java Components (JRE) and vCenter Server tc Server. Inventory Service and the vSphere Web Client depend on vCenter Server Java Components (JRE). Uninstall vCenter Single Sign-On, Inventory Service, vCenter Server, and the vSphere Web Client before you uninstall vCenter Server Java Components (JRE) and vCenter Server tc Server.

Uninstalling the vCenter Server system or the vSphere Web Client does not uninstall any of the other components, such as the bundled database or Microsoft .NET Framework. Do not uninstall the other components if other applications on your system depend on them.

Procedure

- 1 If you are uninstalling the vCenter Server system, remove the hosts from the Hosts and Clusters inventory.
- 2 As Administrator on the Microsoft Windows system, select **Start > Settings > Control Panel > Add/Remove Programs**.
- 3 Select the component to remove from the list and click **Remove**.
- 4 Click **Yes** to confirm that you want to remove the program.
- 5 Click **Finish**.

Creating vCenter Server Linked Mode Groups

A Linked Mode group allows you to log in to any single instance of vCenter Server and view and manage the inventories of all the vCenter Server systems in the group.

You can join multiple vCenter Server systems to form a Linked Mode group. You can configure a Linked Mode group during vCenter Server installation or after vCenter Server is installed.

To join a vCenter Server group, you enter the fully qualified domain name (or IP address) of a remote machine on which vCenter Server is running. The remote machine can be any vCenter Server instance that is, or will become, a member of the Linked Mode group.

You must also provide the LDAP port number of the remote vCenter Server instance.

vCenter Server instances in a group replicate shared global data to the LDAP directory. The global data includes the following information for each vCenter Server instance:

- Connection information (IP and ports)
- Certificates
- Licensing information
- User roles

The vSphere Web Client can connect to a Linked Mode vCenter Server environment only if the vSphere Web Client is logged in to vCenter Server as a domain user. For example, if the Linked Mode vCenter Servers are installed with user "abc" added to domain "xyz," the vSphere Web Client user should log in using **xyz\abc**.

NOTE vCenter Server 5.1 can be joined in a Linked Mode group only with other instances of vCenter Server 5.1. Do not join a version 5.1 vCenter Server instance to a version 5.0 or earlier vCenter Server instance.

Linked Mode Considerations for vCenter Server

Consider several issues before you configure a Linked Mode group.

Before you configure a Linked Mode group, consider the following issues.

- If you are upgrading a version 5.x vCenter Server that is part of a Linked Mode group, it will not be removed from the group. If you are upgrading a pre-5.0 vCenter Server that is part of a Linked Mode group, it will be removed from the group. vCenter Server does not support Linked Mode groups that contain both version 5.x and pre-5.0 versions of vCenter Server. Similarly, vCenter Server does not support Linked Mode groups that contain different 5.x versions of vCenter Server or different versions of vCenter Single Sign-On. After all 5.x vCenter Server and vCenter Single Sign-On instances in a Linked Mode group are upgraded to the same 5.x version, you can rejoin them.
- Each vCenter Server user sees the vCenter Server instances on which they have valid permissions.
- When you set up your vCenter Server Linked Mode group, you must install the first vCenter Server as a standalone instance because you do not yet have a remote vCenter Server machine to join. Subsequent vCenter Server instances can join the first vCenter Server or other vCenter Server instances that have joined the Linked Mode group.
- If you join a vCenter Server to a standalone instance that is not part of a domain, you must add the standalone instance to a domain and add a domain user as an administrator.
- The vCenter Server instances in a Linked Mode group do not need to have the same domain user login. The instances can run under different domain accounts. By default, they run as the LocalSystem account of the machine on which they are running, which means that they are different accounts.
- During vCenter Server installation, if you enter an IP address for the remote instance of vCenter Server, the installer converts it into a fully qualified domain name.



CAUTION If you need to uninstall and reinstall vCenter Server on more than one member of a Linked Mode group, do so with a single vCenter Server at a time. Uninstalling and reinstalling multiple linked vCenter Servers at the same time is not supported, and can cause errors that prevent vCenter Server from connecting to vCenter Inventory Service. If it is necessary to uninstall and reinstall multiple linked vCenter Servers at the same time, isolate them from the Linked Mode group first, and rejoin them to the Linked Mode group after the reinstallation is complete.

Linked Mode Prerequisites for vCenter Server

Prepare the vCenter Server system for joining a Linked Mode group.

Before joining a vCenter Server to a Linked Mode group, review [“Linked Mode Considerations for vCenter Server,”](#) on page 110.

All the requirements for standalone vCenter Server systems apply to Linked Mode systems.

The following requirements apply to each vCenter Server system that is a member of a Linked Mode group:

- vCenter Server does not support Linked Mode groups that contain both version 5.x and pre-5.0 versions of vCenter Server. Similarly, vCenter Server does not support Linked Mode groups that contain different 5.x versions of vCenter Server or different versions of vCenter Single Sign-On. After all 5.x vCenter Server and vCenter Single Sign-On instances in a Linked Mode group are upgraded to the same 5.x version, you can rejoin them.
- Make sure that all vCenter Servers in a Linked Mode group are registered to the same vCenter Single Sign-On server.
- To join a vCenter Server to another vCenter Server in Linked Mode, the currently logged-in user who is performing the join operation must have access to the vCenter Server database of each vCenter Server.
- When you join a vCenter Server instance to a Linked Mode group, the installer must be run by a domain user who is an administrator on both the machine where vCenter Server is installed and the target machine of the Linked Mode group.
- When you join a vCenter Server instance to a Linked Mode group, if User Account Control (UAC) is enabled on the machine, the join operation requires elevated privileges through the **Run as administrator** option. This is true even if the logged-in user is a domain administrator user.
- To join a Linked Mode group the vCenter Server must be in evaluation mode or licensed as a Standard edition. vCenter Server Foundation and vCenter Server Essentials editions do not support Linked Mode.
- DNS must be operational for Linked Mode replication to work.
- The vCenter Server instances in a Linked Mode group can be in different domains if the domains have a two-way trust relationship. Each domain must trust the other domains on which vCenter Server instances are installed.
- All vCenter Server instances must have network time synchronization. The vCenter Server installer validates that the machine clocks are not more than five minutes apart. See [“Synchronizing Clocks on the vSphere Network,”](#) on page 59.

Joining a Linked Mode Group During and After Installation

You can join a system to a Linked Mode group during the vCenter Server installation or after you install vCenter Server.

For example, suppose you have three machines on which you want to install vCenter Server. You want the three instances to be members of a Linked Mode group.

- 1 On Machine 1, you install vCenter Server as a standalone instance because you do not have a remote vCenter Server machine to join.
- 2 On Machine 2, you install vCenter Server, choose to join a Linked Mode group, and provide the fully qualified domain name of Machine 1.
- 3 On Machine 3, you upgrade to vCenter Server 5.x. After the upgrade, you configure Machine 3 to join either Machine 1 or Machine 2. Machine 1, Machine 2, and Machine 3 are now members of a Linked Mode group.

Join a Linked Mode Group After Installation

After installing vCenter Server, you can join a vCenter Server to a Linked Mode group.

Prerequisites

See [“Linked Mode Prerequisites for vCenter Server,”](#) on page 111.

NOTE vCenter Server does not support Linked Mode groups that contain both version 5.x and pre-5.0 versions of vCenter Server. Similarly, vCenter Server does not support Linked Mode groups that contain different 5.x versions of vCenter Server or different versions of vCenter Single Sign-On. After all 5.x vCenter Server and vCenter Single Sign-On instances in a Linked Mode group are upgraded to the same 5.x version, you can rejoin them.

Procedure

- 1 Select **Start > All Programs > VMware > vCenter Server Linked Mode Configuration**.
- 2 Click **Next**.
- 3 Select **Modify linked mode configuration** and click **Next**.
- 4 Click **Join this vCenter Server instance to an existing linked mode group or another instance** and click **Next**.
- 5 Enter the server name and LDAP port number of a remote vCenter Server instance that is a member of the group and click **Next**.

If you enter an IP address for the remote server, the installer converts it into a fully qualified domain name.

- 6 If the vCenter Server installer detects a role conflict, select how to resolve the conflict.

Option	Action
Yes, let VMware vCenter Server resolve the conflicts for me	Click Next . The role on the joining system is renamed to <i>vcenter_name</i> role_name, where <i>vcenter_name</i> is the name of the vCenter Server system that is joining the Linked Mode group, and <i>role_name</i> is the name of the original role.
No, I'll resolve the conflicts myself	To resolve the conflicts manually: <ol style="list-style-type: none"> a Using the vSphere Web Client, log in to one of the vCenter Server systems using an account with Administrator privileges. b Rename the conflicting role. c Close the vSphere Web Client session and return to the vCenter Server installer. d Click Back and click Next. The installation continues without conflicts.

A conflict results if the joining system and the Linked Mode group each contain a role with the same name but with different privileges.

- 7 Click **Finish**.

vCenter Server restarts. Depending on the size of your inventory, the change to Linked Mode might take from a few seconds to a few minutes to complete.

The vCenter Server instance is now part of a Linked Mode group. After you form a Linked Mode group, you can log in to any single instance of vCenter Server and view and manage the inventories of all the vCenter Servers in the group. It might take several seconds for the global data (such as user roles) that are changed on one machine to be visible on the other machines. The delay is usually 15 seconds or less. It might take a few minutes for a new vCenter Server instance to be recognized and published by the existing instances, because group members do not read the global data very often.

For information about configuring and using your Linked Mode group, see the *vCenter Server and Host Management* documentation.

Isolate a vCenter Server Instance from a Linked Mode Group

You can isolate a vCenter Server instance from a Linked Mode group.

Procedure

- 1 Select **Start > All Programs > VMware > vCenter Server Linked Mode Configuration**.
- 2 Click **Modify linked mode configuration** and click **Next**.
- 3 Click **Isolate this vCenter Server instance from linked mode group** and click **Next**.
- 4 Click **Continue** and click **Finish**.

vCenter Server restarts. Depending on the size of your inventory, the change to Linked Mode configuration might take from a few seconds to a few minutes to complete.

The vCenter Server instance is no longer part of the Linked Mode group.

Set the IP Address for a Linked Mode vCenter Server with Multiple Network Interfaces

If a vCenter Server in a Linked Mode group has multiple network interfaces, you must set the IP address that the vCenter Server advertises to the other vCenter Servers in the Linked Mode group.

Unless you set the IP address, the vSphere Web Client will not work properly with the vCenter Server.

Procedure

- 1 On the vCenter Server host machine, using a text editor, open the file `Inventory_Service_installation_directory/lib/server/config/query-server-config.xml`.

- 2 Find the following line.

```
!-- <property name="externalAddress" value="192.168.0.1" /> /--
```

- 3 Change the line as follows.

```
<property name="externalAddress" value="IP_address" />
```

IP_address is the IP address of this machine on a subnet that will be used to communicate with other vCenter Server instances in the Linked Mode group.

- 4 Restart the vCenter Inventory Service.
 - a In the **Administrative Tools** control panel, select **Services**.
 - b Right-click vCenter Inventory Service and select **Start**.

The status changes to Started.

The IP address is configured.

Linked Mode Troubleshooting

If you are having trouble with your Linked Mode group, consider the following points.

When you have multiple vCenter Server instances, each instance must have a working relationship with the domain controller and not conflict with another machine that is in the domain. Conflicts can occur, for example, when you clone a vCenter Server instance that is running in a virtual machine and you do not use sysprep or a similar utility to ensure that the cloned vCenter Server instance has a globally unique identifier (GUID).

If the domain controller is unreachable, vCenter Server might be unable to start. You might be unable to change the Linked Mode configuration of the affected vCenter Server system. If this occurs, resolve the problem with the domain controller and restart vCenter Server. If resolving the problem with the domain controller is impossible, you can restart vCenter Server by removing the vCenter Server system from the domain and isolating the system from its current Linked Mode group.

The DNS name of the machine must match with the actual machine name. Symptoms of machine names not matching the DNS name are data replication problems, ticket errors when trying to search, and missing search results from remote instances.

NOTE Make sure your Windows and network-based firewalls are configured to allow Linked Mode.

Configure Firewall Access by Opening Selected Ports

vCenter Server uses Microsoft ADAM/AD LDS to enable Linked Mode, which uses the Windows RPC port mapper to open RPC ports for replication. When you install vCenter Server in Linked Mode, the firewall configuration on any network-based firewalls must be modified.

Incorrect configuration of firewalls can cause licenses and roles to become inconsistent between instances.

Procedure

- ◆ Configure Windows RPC ports to generically allow selective ports for machine-to-machine RPC communication.

Choose one of the following methods.

- Change the registry settings. See <http://support.microsoft.com/kb/154596/en-us>.
- Use Microsoft's RPPCCfg.exe tool. See <http://support.microsoft.com/kb/908472/en-us>.

Configuring VMware vCenter Server - tc Server Settings in vCenter Server

Starting with vCenter Server 5.1, VMware Tomcat Server settings can no longer be configured through the Windows user interface. vCenter Server versions 5.1 and later use VMware vCenter Server - tc Server, an enterprise version of Apache Tomcat 7. Tomcat version 7 does not provide a control panel in the Windows user interface. Instead, you configure Tomcat by editing configuration files manually.

You can adjust the JVM maximum heap size for vCenter Server, vCenter Inventory Service, and Profile-Driven Storage Service. For JVM heap size recommendations, see “[Hardware Requirements for vCenter Server, the vSphere Web Client, vCenter Inventory Service, and vCenter Single Sign-On](#),” on page 17.

Settings for Java options are stored in the following files.

- vCenter Server. *installation_directory\VMware\Infrastructure\tomcat\conf\wrapper.conf*
- vCenter Inventory Service. *installation_directory\VMware\Infrastructure\Inventory Service\conf\wrapper.conf*

- Profile-Driven Storage Service. *installation_directory\VMware\Infrastructure\Profile-Driven Storage\conf\wrapper.conf*
- The vSphere Web Client.
installation_directory\VMware\vSphereWebClient\server\bin\service\conf\wrapper.conf

Table 5-1. Inventory Service and Profile-Driven Storage Service Java Maximum JVM Heap Size Setting in the *wrapper.conf* Files

Java Option	Setting and Default Value
maxmemorysize The maximum JVM heap size, in megabytes. This setting controls the maximum size of the Java heap. Tuning this parameter can reduce the overhead of garbage collection, improving server response time and throughput. For some applications, the default setting for this option is too low, resulting in a high number of minor garbage collections.	Inventory Service: <code>wrapper.java.maxmemory=2048</code> Profile-Driven Storage Service: <code>wrapper.java.maxmemory=1024</code> The vSphere Web Client: For large deployments you might need to set this option to <code>wrapper.java.maxmemory=2048</code>
ping.timeoutduration	The vSphere Web Client: For large deployments you might need to set this option to <code>wrapper.ping.timeout=120</code>

vCenter Server security and port settings are stored in the following files.

- *installation_directory\VMware\Infrastructure\tomcat\conf\server.xml* and
- *installation_directory\VMware\Infrastructure\tomcat\conf\catalina.properties*

Table 5-2. vCenter Server Port and Security Settings in the *server.xml* and *catalina.properties* Files

vCenter Server Port or Security Setting	Setting and Default Value
Base shutdown port	<code>base.shutdown.port=8003</code>
Base JMX port. The listener implemented by the <code>com.springsource.tcserver.serviceability.rmi.JmxSocketListener</code> class is specific to tc Server. This listener enables JMX management of tc Server, and is the JMX configuration that the AMS management console uses to manage tc Server instances. The port attribute specifies the port of the JMX server that management products, such as AMS, connect to. The variable <code>{jmx.port}</code> is set to 6969 in the default <i>catalina.properties</i> file. The bind attribute specifies the host of the JMX server. By default, this attribute is set to the localhost (127.0.0.1). The default -1 setting disables the port.	<code>base.jmx.port=-1</code>
Web services HTTPS	<code>bio-vmssl.http.port=8080</code>
Web services HTTPS	<code>bio-vmssl.https.port=8443</code>
SSL certificate	<code>bio-vmssl.keyFile.name=C:\ProgramData\VMware\VMware VirtualCenter\SSL\rui.pfx</code>
SSL certificate password	<code>bio-vmssl.SSL.password=testpassword</code>
AJP port	<code>bio-vmssl.ajp.port=8009</code>

See *Getting Started with vFabric tc Server* and *vFabric tc Server Administration* at <https://www.vmware.com/support/pubs/vfabric-tcserver.html>.

You can manage the Windows services for vCenter Server from the Administrative Tools control panel, under Services. The Windows service for vCenter Server is listed as VMware VirtualCenter Management Webservices.

VMware vCenter Management Webservices Service Fails to Start

When you reboot the vCenter Server machine after installing vCenter Server, the VMware VirtualCenter Management Webservices service does not start.

Problem

The VMware VirtualCenter Management Webservices service does not start automatically.

Cause

This problem can occur when vCenter Server and the database are installed on the same machine.

Solution

- ◆ Start the service manually.

Select **Settings > Control Panel > Administrative Tools > Services > VMware VirtualCenter Management Webservices** and start the service. The machine might require several minutes to start the service.

Back Up the Inventory Service Database on Windows

You should back up the Inventory Service database as part of your regular vCenter Server database administration.

To move the Inventory Service database to a different host machine, back up the database on the source machine and restore the database on the destination machine. Your vCenter Server database administration should also include regular backups of your vCenter Server database. See the vendor documentation for your vCenter server database type.

Prerequisites

- Consult your database administrator about backing up and restoring databases.
- Verify that you have system administrator privileges to perform backup or restore operations.

Procedure

- 1 On the source machine, stop the Inventory Service.
 - a From the Windows Administrative Tools control panel, select **Services**.
 - b Right-click **VMware vCenter Inventory Service** and select **Stop**.

The Status changes from Started to blank.
- 2 On the source machine, open the command prompt in the vCenter Server and change the directory to `vCenter_Server_installation_directory\Infrastructure\Inventory Service\scripts`.
- 3 Run the following command at the prompt to back up the Inventory Service database.

```
backup.bat -file backup_file_name
```

When the backup operation finishes, the message `Backup completed successfully` appears.

Restore an Inventory Service Database Backup on Windows

You can restore a backup of your Inventory Service database for disaster recovery, after a vCenter Server upgrade, or to move the database to a new machine.

The machine that you back up the database from is the source machine. The machine that you restore the database to is the destination machine.

Prerequisites

- Consult your database administrator about backing up and restoring databases.
- Verify that you have system administrator privileges to perform backup or restore operations.

Procedure

- 1 On the destination machine, stop the Inventory Service.
 - a From the Windows Administrative Tools control panel, select **Services**.
 - b Right-click **VMware vCenter Inventory Service** and select **Stop**.

The Status changes from Started to blank.
- 2 On the destination machine, open the command prompt in the vCenter Server and change the directory to `vCenter Server install location\Infrastructure\Inventory Service\scripts`.
- 3 Run the following command at the command prompt to restore the Inventory Service database.


```
restore -backup backup_file_name
```

When the restore operation finishes, the message `The Restore completed successfully` message appears.

Back Up the Inventory Service Database on Linux

You should back up the Inventory Service database as part of your regular vCenter Server database administration.

To move the Inventory Service database to a different host machine, back up the database on the source machine and restore the database on the destination machine.

Prerequisites

- Consult your database administrator about backing up and restoring databases.
- Verify that you have system administrator privileges to perform backup or restore operations.

Procedure

- 1 On the source machine, open a console and run the `service vmware-inventoryservice stop` command to stop the Inventory Service before you restore the Inventory Service database.
- 2 On the source machine, open a command prompt in the vCenter Server and change the directory to `/usr/lib/vmware-vpx/inventoryservice/scripts/`.
- 3 Run the following command to back up the Inventory Service database.


```
./backup.sh -file backup_file_name
```

When the backup operation finishes, the message `Backup completed successfully` appears.

What to do next

See [“Restore an Inventory Service Database Backup on Linux,”](#) on page 117.

Restore an Inventory Service Database Backup on Linux

You can restore a backup of your Inventory Service database for disaster recovery, after a vCenter Server upgrade, or to move the database to a new machine.

The machine that you back up the database from is the source machine. The machine that you restore the database to is the destination machine.

Prerequisites

- Consult your database administrator about backing up and restoring databases.
- Verify that you have system administrator privileges to perform backup or restore operations.

Procedure

- 1 On the destination machine, open a console and run the service `vmware-inventoryservice stop` command to stop the Inventory Service before you restore the Inventory Service database.
- 2 On the destination machine, open a command prompt in the vCenter Server and change the directory to `/usr/lib/vmware-vpx/inventoryservice/scripts/`.
- 3 Run the following command at the prompt to restore the Inventory Service database.

```
./restore.sh -backup backup_file_name
```

When the restore operation finishes, the message `The Restore completed successfully` message appears at the command prompt.

Reset the vCenter Inventory Service Database

If the vCenter Inventory Service Database is corrupted or otherwise inoperable, you can reset it. You should also reset the vCenter Inventory Service Database if you reset the vCenter Server database.



CAUTION Resetting the vCenter Inventory Service Database can cause data loss. Perform this procedure only with VMware Technical Support.

Procedure

- 1 Stop the vCenter Inventory Service.
 - a From the Windows Administrative Tools control panel, select **Services**.
 - b Right-click **VMware vCenter Inventory Service** and select **Stop**.
- 2 Open a command prompt.
- 3 Delete the entire contents of the `C:\Program Files\VMware\Infrastructure\Inventory_Service\data` directory.

If you installed vCenter Server in a different location from the default `C:\Program Files\`, adjust the path accordingly.
- 4 Change directory to `C:\Program Files\VMware\Infrastructure\Inventory_Service\scripts`

If you installed vCenter Server in a different location from the default `C:\Program Files\`, adjust the path accordingly.
- 5 Run the `createDB.bat` command, with no arguments, to reset the vCenter Inventory Service database.
- 6 Start the vCenter Inventory Service.
 - a From the Windows Administrative Tools control panel, select **Services**.
 - b Right-click **VMware vCenter Inventory Service** and select **Start**.
- 7 Change directory to `C:\Program Files\VMware\Infrastructure\VirtualCenter Server\isregtool`.

If you installed vCenter Server in a different location from the default `C:\Program Files\`, adjust the path accordingly.

- 8 Run the `register-is.bat` command to update the stored configuration information of the Inventory Service.

```
register-is.bat vCenter_Server_URL Inventory_Service_URL Lookup_Service_URL
```

Use the following example as a model.

```
register-is.bat https://machinename.corp.com:443/sdk https://machinename.corp.com:10443
https://machinename.corp.com:7444/lookupservice/sdk
```

In this example, 443, 10443, and 7444 are the default HTTPS port numbers for vCenter Server, Inventory Service, and vCenter Single Sign On respectively. If you use custom ports, replace the port numbers in the example with the port numbers you use.

- 9 Restart vCenter Server.

The vCenter Inventory Service database is reset.

Enable IPv6 Support for vCenter Inventory Service

vCenter Inventory Service does not support binding on IPv6 interfaces by default. When you install vCenter Server, vCenter Inventory Service supports only IPv4 by default. You can enable IPv6 support for vCenter Inventory Service by modifying the `Inventory Service dataservice.properties` file.

Procedure

- 1 Stop the vCenter Inventory Service.
 - a From the Administrative Tools control panel, select **Services**.
 - b Right-click **vCenter Inventory Service** and select **Stop**.
- 2 In a text editor, open the
file:`Inventory_Service_installation_directory/lib/server/config/dataservice.properties`.
- 3 Change the line `dataservice.nio.enabled = true` to `dataservice.nio.enabled = false`
- 4 Restart the vCenter Inventory Service.

IPv6 support for vCenter Inventory Service is enabled.

Before You Install ESXi

Before you install ESXi, understand the installation process and options.

This chapter includes the following topics:

- [“Options for Installing ESXi,”](#) on page 121
- [“Media Options for Booting the ESXi Installer,”](#) on page 123
- [“Using Remote Management Applications,”](#) on page 134
- [“Required Information for ESXi Installation,”](#) on page 134
- [“Download the ESXi Installer,”](#) on page 135

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.

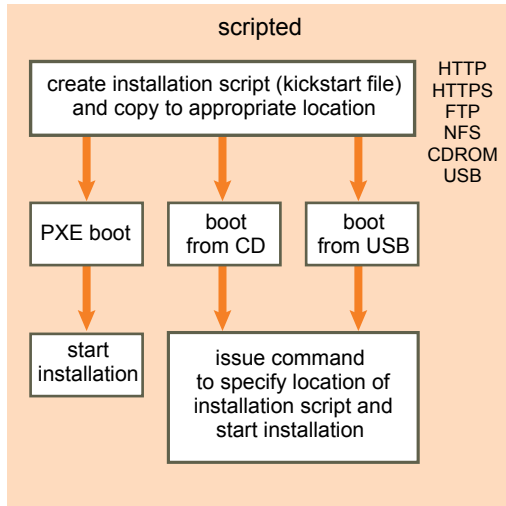
You boot the installer from a CD or DVD, from a bootable USB device, or by PXE booting the installer from a location on the network. You follow the prompts in the installation wizard to install ESXi to disk. See [“Installing ESXi Interactively,”](#) on page 137.

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, Upgrading, or Migrating Hosts Using a Script,”](#) on page 140.

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 6-1. Scripted Installation

Auto Deploy ESXi Installation

vSphere 5.x provides several ways to install ESXi with Auto Deploy.

These topics describe Auto Deploy options for ESXi installation.

Customizing Installations with ESXi Image Builder CLI

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

ESXi Image Builder CLI is a PowerShell CLI command set that you can use 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 Image Builder in either of the following ways:

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

See [“Using vSphere ESXi Image Builder CLI,”](#) on page 223 and [“Installing ESXi Using vSphere Auto Deploy,”](#) on page 155.

About ESXi Evaluation and Licensed Modes

After you purchase vSphere licenses, VMware provides a serial number that you use to license ESXi hosts. You can use evaluation mode to explore the entire set of features that are available for ESXi hosts, including features that are not included in the license that you have.

For example, in evaluation mode, you can use vMotion, HA, DRS, and other features, even if you have not licensed those features.

The installable version of ESXi is always installed in evaluation mode. ESXi Embedded is preinstalled on an internal USB 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, even if you start in licensed mode rather than evaluation mode. Any time during the 60-day evaluation period, you can convert from licensed mode to evaluation mode. To take full advantage of the 60-day evaluation period, you should convert to evaluation mode as soon as possible after you first power on the host.

For information about managing licensing and setting an ESXi host to evaluation mode, see the *vCenter Server and Host Management* documentation.

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,”](#) on page 123.
- Boot from a USB flash drive. See [“Format a USB Flash Drive to Boot the ESXi Installation or Upgrade,”](#) on page 123.
- PXE boot from the network. [“PXE Booting the ESXi Installer,”](#) on page 127
- Boot from a remote location using a remote management application. See [“Using Remote Management Applications,”](#) on page 134

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,”](#) on page 126.

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.

These instructions assume that you are performing the procedure on a Linux machine and that the USB flash drive is detected by the operating system as `/dev/sdb`.

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

Prerequisites

- From the VMware Web site, download the ESXi ISO image `VMware-VMvisor-Installer-5.x.x-XXXXXX.x86_64.iso`, including the file `isolinux.cfg`, where `5.x.x` is the version of ESXi you are installing, and `XXXXXX` is the build number of the installer ISO image.

- Verify that the machine on which you are performing this procedure has access to syslinux version 3.86. This procedure requires syslinux version 3.86.

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 In a terminal window, run the following command.

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

This command displays current log messages in the terminal window.

- b Plug in your USB flash drive.

The terminal window displays several messages identifying 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, without the brackets, in place of sdb, in this procedure.

- 2 Create a partition table on the USB flash device.

```
/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 `a` to set the active flag on partition 1.
 - e 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)
```

- f 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 Run the following commands.

```
/path_to_syslinux-3.86_directory/syslinux-3.86/bin/syslinux /dev/sdb1
cat /path_to_syslinux-3.86_directory/syslinux-3.86/usr/share/syslinux/mbr.bin > /dev/sdb
```

- 5 Mount the USB flash drive.

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

- 6 Mount the ESXi installer ISO image.

```
mount -o loop VMware-VMvisor-Installer-5.x.x-XXXXXX.x86_64.iso /esxi_cdrom
```

- 7 Copy the contents of the ISO image to `/usbdisk`.

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

- 8 Rename the `isolinux.cfg` file to `syslinux.cfg`.
`mv /usbdisk/isolinux.cfg /usbdisk/syslinux.cfg`
- 9 In the file `/usbdisk/syslinux.cfg`, change the line `APPEND -c boot.cfg` 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 now 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 The `ks` file containing the installation or upgrade script cannot be located 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,”](#) on page 140 and [“About PXE Configuration Files,”](#) on page 130.

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 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,”](#) on page 142 and [“About the boot.cfg File,”](#) on page 152.

Prerequisites

- Linux machine.
- The ESXi ISO image `VMware-VMvisor-Installer-5.x.x-XXXXXX.x86_64.iso`, where `5.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 into a folder:
mount -o loop VMware-VMvisor-Installer-5.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 using the `kernelopt` option.

This step makes the installation or upgrade completely automatic, without the need to specify the kickstart file during the installation or upgrade.
- 6 Recreate the ISO image:
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 /esxi_cdrom

The ISO image now includes your custom installation or upgrade script.

What to do next

Install ESXi from the ISO image.

PXE Booting the ESXi Installer

You use the preboot execution environment (PXE) to boot a host and launch the ESXi installer from a network interface.

ESXi 5.x 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 using PXE.

PXE uses 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 are capable of running ESXi have network adapters that are able to PXE boot.

NOTE Ensure that the Auto Deploy server has an IPv4 address. PXE booting is supported only with IPv4.

About the TFTP Server, PXELINUX, and gPXE

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.

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.

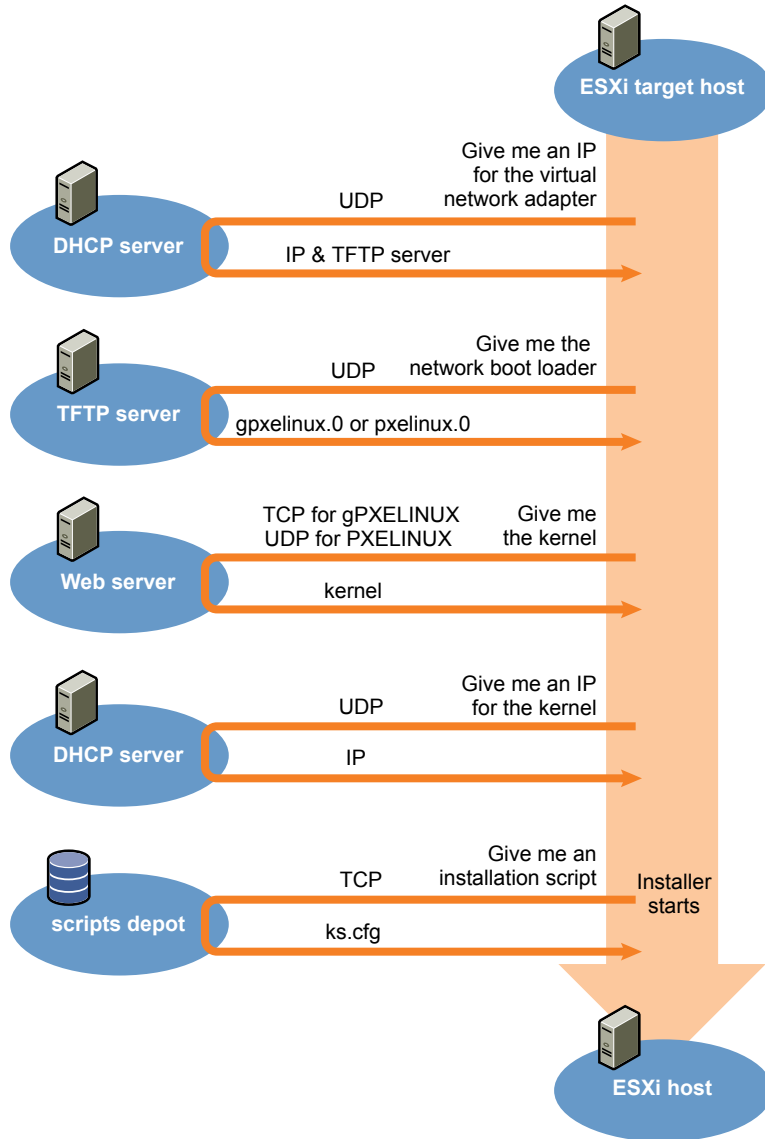
If your TFTP server will run on a Microsoft Windows host, use tftpd32 version 2.11 or later. See <http://tftpd32.jounin.net/>. Earlier versions of tftpd32 were incompatible with PXELINUX and gPXE.

You can also acquire a TFTP server from one of the packaged appliances on the VMware Marketplace.

The PXELINUX and gPXE environments allow your target machine to boot the ESXi installer. PXELINUX is part of the SYSLINUX package, which can be found at <http://www.kernel.org/pub/linux/utils/boot/syslinux/>, although many Linux distributions include it. Many versions of PXELINUX also include gPXE. Some distributions, such as Red Hat Enterprise Linux version 5.3, include earlier versions of PXELINUX that do not include gPXE.

If you do not use gPXE, you might experience problems while booting the ESXi installer on a heavily loaded network. TFTP is sometimes unreliable for transferring large amounts of data. If you use PXELINUX without gPXE, the `pxelinux.0` binary file, the configuration file, the kernel, and other files are transferred by TFTP. If you use gPXE, only the `gpxelinux.0` binary file and configuration file are transferred by TFTP. With gPXE, you can use a Web server to transfer the kernel and other files required to boot the ESXi installer.

NOTE VMware tests PXE booting with PXELINUX version 3.86. 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.

Figure 6-2. Overview of PXE Boot Installation Process

Sample DHCP Configuration

To PXE boot the ESXi installer, the DHCP server must send the address of the TFTP server and a pointer to the `pxelinux.0` or `gpxelinux.0` directory.

The DHCP server is used by the target machine to obtain an IP address. The DHCP server must be able to determine whether the target machine is allowed to boot and the location of the PXELINUX binary (which usually resides on a TFTP server). When the target machine first boots, it broadcasts a packet across the network requesting this information to boot itself. The DHCP server responds.



CAUTION Do not set up a new 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.

gPXE Example

This example shows how to configure a ISC DHCP version 3.0 server to enable gPXE.

```
allow booting;
allow bootp;
# gPXE options
option space gppe;
option gppe-encap-opts code 175 = encapsulate gppe;
option gppe.bus-id code 177 = string;
class "pxeclients" {
    match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
    next-server TFTP server address;
    if not exists gppe.bus-id {
        filename "/gpxelinux.0";
    }
}
subnet Network address netmask Subnet Mask {
    range Starting IP Address Ending IP Address;
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `gpxelinux.0` binary file on the TFTP server. The IP address assigned is in the range defined in the subnet section of the configuration file.

PXELINUX (without gPXE) Example

This example shows how to configure a ISC DHCP version 3.0 server to enable PXELINUX.

```
#
# DHCP Server Configuration file.
# see /usr/share/doc/dhcp*/dhcpd.conf.sample
#
ddns-update-style ad-hoc;
allow booting;
allow bootp;
class "pxeclients" {
    match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
    next-server xxx.xxx.xx.xx;
    filename = "pxelinux.0";
}
subnet 192.168.48.0 netmask 255.255.255.0 {
    range 192.168.48.100 192.168.48.250;
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `pxelinux.0` binary file on the TFTP server. The IP address assigned is in the range defined in the subnet section of the configuration file.

About PXE Configuration Files

The PXE configuration file defines the menu displayed to the target ESXi host as it boots up and contacts the TFTP server. You need a PXE configuration file to PXE boot the ESXi installer.

The TFTP server constantly listens for PXE clients on the network. When it detects that a PXE client is requesting PXE services, it sends the client a network package that contains a boot menu.

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,”](#) on page 152

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

File Location for the PXE Configuration File

Save the file in `var/lib/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 the ESXi Installer by Using PXELINUX and a PXE Configuration File

You can use a TFTP server to PXE boot the ESXi installer, using PXELINUX and a PXE configuration file.

See also [“About Installation and Upgrade Scripts,”](#) on page 142 and [“About the boot.cfg File,”](#) on page 152

Prerequisites

Verify that your environment has the following components:

- The ESXi installer ISO image downloaded from the VMware Web site.
- TFTP server that supports PXE booting with gPXE. See [“About the TFTP Server, PXELINUX, and gPXE,”](#) on page 127.
- DHCP server configured for PXE booting. See [“Sample DHCP Configuration,”](#) on page 128.
- PXELINUX
- Server with a hardware configuration that is supported with your version of ESXi. See the Hardware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php>.
- Network security policies to allow TFTP traffic (UDP port 69)
- (Optional) Installation script, the kickstart file. See [“About Installation and Upgrade Scripts,”](#) on page 142.

- Network adapter with PXE support on the target ESXi host
- IPv4 networking. IPv6 is not supported for PXE booting.

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.

Procedure

- 1 Create the `/tftpbboot/pxelinux.cfg` directory on your TFTP server.
- 2 On the Linux machine, install PXELINUX.
 PXELINUX is included in the SYSLINUX package. Extract the files, locate the `pxelinux.0` file and copy it to the `/tftpbboot` directory on your TFTP server.
- 3 Configure the DHCP server to send the following information to each client host:
 - The name or IP address of your TFTP server.
 - The name of your initial boot file. This is `pxelinux.0`.
- 4 Copy the contents of the ESXi installer image to the `/var/lib/tftpbboot` directory on the TFTP server.
- 5 (Optional) For a scripted installation, in the `boot.cfg` file, add the `kernelopt` option on the line following 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 containing the `ks.cfg` file.

```
kernelopt=ks=http://XXX.XXX.XXX.XXX/esxi_ksFiles/ks.cfg
```
- 6 Create a PXE configuration file.
 This file defines how the host boots when no operating system is present. The PXE configuration file references the boot files. Use the following code as a model, where `XXXXXX` is the build number of the ESXi installer image.

```
DEFAULT menu.c32
MENU TITLE ESXi-5.x.x-XXXXXX-full Boot Menu
NOHALT 1
PROMPT 0
TIMEOUT 80
LABEL install
    KERNEL mboot.c32
    APPEND -c location of boot.cfg
MENU LABEL ESXi-5.x.x-XXXXXX-full ^Installer
LABEL hddboot
    LOCALBOOT 0x80
MENU LABEL ^Boot from local disk
```
- 7 Name the file with the MAC address of the target host machine: `01-mac_address_of_target_ESXi_host`.
 For example, `01-23-45-67-89-0a-bc`.
- 8 Save the PXE configuration file in `/tftpbboot/pxelinux.cfg` on the TFTP server.
- 9 Boot the machine with the network adapter.

PXE Boot the ESXi Installer by Using PXELINUX and an isolinux.cfg PXE Configuration File

You can PXE boot the ESXi installer using PXELINUX, and use the `isolinux.cfg` file as the PXE configuration file.

See also [“About Installation and Upgrade Scripts,”](#) on page 142 and [“About the boot.cfg File,”](#) on page 152

Prerequisites

Verify that your environment has the following components:

- The ESXi installer ISO image downloaded from the VMware Web site.
- TFTP server that supports PXE booting with PXELINUX. See “[About the TFTP Server, PXELINUX, and gPXE,](#)” on page 127.
- DHCP server configured for PXE booting. See “[Sample DHCP Configuration,](#)” on page 128.
- PXELINUX
- Server with a hardware configuration that is supported with your version of ESXi. See the *Hardware Compatibility Guide* at <http://www.vmware.com/resources/compatibility/search.php>.
- Network security policies to allow TFTP traffic (UDP port 69)
- (Optional) Installation script, the kickstart file. See “[About Installation and Upgrade Scripts,](#)” on page 142.
- Network adapter with PXE support on the target ESXi host
- IPv4 networking. IPv6 is not supported for PXE booting.

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.

Procedure

- 1 Create the `/tftpbboot/pxelinux.cfg` directory on your TFTP server.

- 2 On the Linux machine, install PXELINUX.

PXELINUX is included in the SYSLINUX package. Extract the files, locate the file `pxelinux.0` and copy it to the `/tftpbboot` directory on your TFTP server.

- 3 Configure the DHCP server.

The DHCP server sends the following information to your client hosts:

- The name or IP address of your TFTP server.
- The name of your initial boot file. This is `pxelinux.0`.

- 4 Copy the contents of the ESXi installer image to the `/var/lib/tftpbboot` directory on the TFTP server.

- 5 (Optional) For a scripted installation, in the `boot.cfg` file, add the `kernelopt` option on the next line after the `kernel` command, to specify the location fo the installation script.

In the following example, `XXX.XXX.XXX.XXX` is the IP address of the server where the installation script resides.

```
kernelopt=ks=http://XXX.XXX.XXX.XXX/esxi_ksFiles/ks.cfg
```

- 6 Copy the `isolinux.cfg` file from the ESXi installer ISO image to the `/tftpbboot/pxelinux.cfg` directory.

The `isolinux.cfg` file contains the following code, where `XXXXXX` is the build number of the ESXi installer image:

```
DEFAULT menu.c32
MENU TITLE ESXi-5.x.x-XXXXXX-full Boot Menu
NOHALT 1
PROMPT 0
TIMEOUT 80
LABEL install
    KERNEL mboot.c32
    APPEND -c location of boot.cfg
```

```

MENU LABEL ESXi-5.x.x-XXXXXX-full ^Installer
LABEL hddboot
  LOCALBOOT 0x80
MENU LABEL ^Boot from local disk

```

- 7 Rename the `isolinux.cfg` file with the MAC address of the target host machine: `01-mac_address_of_target_ESXi_host`. For example, `01-23-45-67-89-0a-bc`
- 8 Boot the machine with the network adapter.

PXE Boot the ESXi Installer Using gPXE

You can PXE boot the ESXi installer using gPXE.

See also “[About Installation and Upgrade Scripts](#),” on page 142 and “[About the boot.cfg File](#),” on page 152

Prerequisites

Verify that your environment has the following components:

- The ESXi installer ISO image downloaded from the VMware Web site
- HTTP Web server that is accessible by your target ESXi hosts
- DHCP server configured for PXE booting: `/etc/dhcpd.conf` is configured for client hosts with a TFTP server and the initial boot file set to `gpxelinux.0/undionly.kpxe`. See “[Sample DHCP Configuration](#),” on page 128.
- Server with a hardware configuration that is supported with your version of ESXi. See the Hardware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php>.
- gPXELINUX
- (Optional) ESXi installation script. See “[About Installation and Upgrade Scripts](#),” on page 142.

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.

Procedure

- 1 Copy the contents of the ESXi installer ISO image to the `/var/www/html` directory on the HTTP server.
- 2 Modify the `boot.cfg` file with the information for the HTTP server.

Use the following code as a model, where `XXX.XXX.XXX.XXX` is the HTTP server IP address. The `kernelopt` line is optional. Include that option to specify the location of the installation script for a scripted installation.

```

title=Loading ESX installer
kernel=http://XXX.XXX.XXX.XXX/tboot.b00
kernelopt=ks=http://XXX.XXX.XXX.XXX/esxi_ksFiles/ks.cfg
modules=http://XXX.XXX.XXX.XXX/b.b00 --- http://XXX.XXX.XXX.XXX/useropts.gz ---
http://XXX.XXX.XXX.XXX/k.b00 --- http://XXX.XXX.XXX.XXX/a.b00 ---
http://XXX.XXX.XXX.XXX/s.v00 --- http://XXX.XXX.XXX.XXX/weaselin.v00 ---
http://XXX.XXX.XXX.XXX/tools.t00 --- http://XXX.XXX.XXX.XXX/imgdb.tgz ---
http://XXX.XXX.XXX.XXX/imgpayld.tgz

```

- 3 gPXE boot the host and press Ctrl+B to access the GPT menu.

- 4 Enter the following commands to boot with the ESXi installer, where XXX.XXX.XXX.XXX is the HTTP server IP address.

```
dhcp net0 ( if dhcp is not set)
kernel -n mboot.c32 http://XXX.XXX.XXX.XXX/mboot.c32
imgargs mboot.c32 -c http://XXX.XXX.XXX.XXX/boot.cfg
boot mboot.c32
```

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 Minimum Firmware Versions,”](#) on page 28. 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.

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 6-1. 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	

Table 6-1. Required Information for ESXi Installation (Continued)

Information	Required or Optional	Default	Comments
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 5GB if you install the components on a single disk.
Migrate existing ESX or ESXi settings. Preserve existing VMFS datastore.	Required if you are installing ESXi on a drive with an existing ESXi or ESX installation.	None	If you have an existing ESX/ESXi 4.x or ESXi 5.0 installation, the ESXi installer offers a choice between preserving or overwriting the VMFS datastore during installation
Root password	Optional	None	The root password must contain between 6 and 64 characters.

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>.

Installing ESXi

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

This chapter includes the following topics:

- [“Installing ESXi Interactively,”](#) on page 137
- [“Installing, Upgrading, or Migrating Hosts Using a Script,”](#) on page 140
- [“Installing ESXi Using vSphere Auto Deploy,”](#) on page 155
- [“Using vSphere ESXi Image Builder CLI,”](#) on page 223

Installing ESXi Interactively

Use the interactive installation option for small deployments of less 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 located 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 page 123
 - On a USB flash drive. See [“Format a USB Flash Drive to Boot the ESXi Installation or Upgrade,”](#) on page 123.

NOTE You can also PXE boot the ESXi installer to launch an interactive installation or a scripted installation. See [“PXE Booting the ESXi Installer,”](#) on page 127.

- 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 will be installed. Alternatively, use a remote management application. See [“Using Remote Management Applications,”](#) on page 134.

- Consider disconnecting your network storage. This action decreases the time it takes the installer to search for available disk drives. Note that 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,”](#) on page 134.
- 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 Virtual SAN 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 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 Virtual SAN 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 leave the password blank, but to secure the system from the first boot, enter a password. 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,”](#) on page 250

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 [Chapter 9, “After You Install and Set Up ESXi,”](#) on page 265.

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, reenable the option to boot to the iSCSI target, so the system will boot from the LUN you installed ESXi on.

Installing, Upgrading, or Migrating Hosts Using a Script

You can quickly deploy ESXi hosts using scripted, unattended installations or upgrades. Scripted installations, upgrades, or migrations 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, upgrade, or migration, you must use the supported commands to create a script, and 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
- HTTP/HTTPS
- NFS
- USB flash drive
- CDROM

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 7-1. 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,”](#) on page 143.

Enter Boot Options to Start an Installation or Upgrade Script

You can start an installation or upgrade script by typing boot command-line 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,”](#) on page 152 and [“PXE Booting the ESXi Installer,”](#) on page 127.

A `ks=...` option must be given, to specify the location of the installation script. Otherwise, a scripted installation or upgrade will not start. If `ks=...` is omitted, the text installer will proceed.

Supported boot options are listed in “[Boot Options](#),” on page 141.

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 7-2. Boot Options for ESXi Installation

Boot Option	Description
BOOTIF= <i>hwtype-MAC address</i>	Similar to the <code>netdevice</code> option, except in the PXELINUX format as described in the IPAPPEND option under SYSLINUX at the syslinux.zytor.com site.
<i>gateway=ip address</i>	Sets this network gateway as the default gateway to be used for downloading the installation script and installation media.
<i>ip=ip address</i>	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 IPAPPEND option under SYSLINUX at the syslinux.zytor.com site.
<i>ks=cdrrom://path</i>	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.
<i>ks=file://path</i>	Performs a scripted installation with the script at <i>path</i> .

Table 7-2. Boot Options for ESXi Installation (Continued)

Boot Option	Description
<code>ks=protocol://serverpath</code>	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.
<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,”](#) on page 140.

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,”](#) on page 126.
- USB Flash drive. See [“Create a USB Flash Drive to Store the ESXi Installation Script or Upgrade Script,”](#) on page 125.
- 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,”](#) on page 142.

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,”](#) on page 140.

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. The `install` command replaces the `autopart` command that was used for scripted ESXi 4.1 installations.

accepteula or vmaccepteula (required)

Accepts the ESXi license agreement. This command functions as it did in ESXi 4.1.

clearpart (optional)

Compared to kickstart, the behavior of the ESXi `clearpart` command is different. Carefully edit the `clearpart` command in your existing scripts.

Clears any existing partitions on the disk. Requires `install` command to be specified.

--drives=	Remove partitions on the specified drives.
--alldrives	Ignores the <code>--drives=</code> requirement and allows clearing of partitions on every drive.
--ignoredrives=	Removes partitions on all drives except those specified. Required unless the <code>--drives=</code> or <code>--alldrives</code> flag is specified.
--overwritevmfs	Permits overwriting of VMFS partitions on the specified drives. By default, overwriting VMFS partitions is not allowed.
--firstdisk= disk-type1 [disk-type2,...]	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>.</p>

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 <code>--disk=diskname</code> , the <i>diskname</i> can be in any of the forms shown in the following examples: <ul style="list-style-type: none"> ■ Path: <code>--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0</code> ■ MPX name: <code>--disk=mpx.vmhba1:C0:T0:L0</code> ■ VML name: <code>--disk=vm1.000000034211234</code>
----------------------------	--

- vmkLUN UID: `--disk=vmkLUN_UID`

For accepted disk name formats, see “Disk Device Names,” on page 152.

`--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`.

`--ignoressd`

Excludes solid-state disks (*S--firstdiskSDs*) 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 Virtual SAN disk group. If you use this option and there is no Virtual SAN partition on the selected disk, the installation will fail. When you install ESXi on a disk that is in Virtual SAN disk group, the result depends on the disk 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 Virtual SAN 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.

`--novmfsondisk`

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=vml.000000034211234`
- vmkLUN UID: `--disk=vmkLUN_UID`

For accepted disk name formats, see [“Disk Device Names,”](#) on page 152.

**--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`.

--overwritevsan

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

--forcemigrate

If a version 4.x host contains customizations, such as third-party VIBs or drivers, that are not included in the installer .ISO, the installer exits with an error describing the problem. The `forcemigrate` option overrides the error and forces the upgrade.

If you are upgrading a 5.0.x host, supported custom VIBs on the host that are not included in the ESXi installer ISO are migrated. If the host or the installer .ISO contains a VIB that creates a conflict and prevents the upgrade, an error message identifies the offending VIB. You can remove the VIB and retry the upgrade, or use ESXi Image Builder to create a custom installer .ISO that resolves the conflict. The `forcemigrate` option is not available.

See the *vSphere Upgrade* documentation for information about upgrading hosts that have third-party custom VIBs.



CAUTION Using the `forcemigrate` option might cause the upgraded host to not boot properly, to exhibit system instability, or to lose functionality.

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
- Default
- Estonian
- Finnish
- French
- German
- Greek
- Icelandic
- Italian
- Japanese
- Latin American
- Norwegian
- Polish
- Portuguese
- Russian
- Slovenian
- Spanish
- Swedish
- Swiss French

- Swiss German
- Turkish
- US Dvorak
- Ukranian
- United Kingdom

serialnum or vmserialnum (optional)

Deprecated in ESXi 5.0.x. Supported in ESXi 5.1. 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)

Specify a network address for the system.

--bootproto=[dhcp|static] Specify 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>.</p>

reboot (optional)

Reboots the machine after the scripted installation is complete.

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

rootpw (required)

Sets the root password for the system.

<code>--iscripted</code>	Specifies that the password is encrypted.
<i>password</i>	Specifies the password value.

upgrade

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

<code>--disk=</code> or <code>--drive=</code>	Specifies the disk to partition. In the command <code>--disk=iskname</code> , the <i>iskname</i> can be in any of the forms shown in the following examples: <ul style="list-style-type: none"> ■ Path: <code>--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0</code> ■ MPX name: <code>--disk=mpx.vmhba1:C0:T0:L0</code> ■ VML name: <code>--disk=vml.000000034211234</code> ■ vmkLUN UID: <code>--disk=vmkLUN_UID</code> <p>For accepted disk name formats, see “Disk Device Names,” on page 152.</p>
<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 (*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.`

`--deletecosvmdk`

If the system is being upgraded from ESX, remove the directory that contains the old Service Console VMDK file, *cos.vmdk*, to reclaim unused space in the VMFS datastore.

`--forcemigrate`

If a version 4.x host contains customizations, such as third-party VIBs or drivers, that are not included in the installer .ISO, the installer exits with an error describing the problem. The *forcemigrate* option overrides the error and forces the upgrade. If you are upgrading a 5.0.x host, supported custom VIBs that are not included in the ESXi installer ISO are migrated. You do not need to use the *forcemigrate* option.

See the *vSphere Upgrade* documentation for information about upgrading hosts that have third-party custom VIBs.



CAUTION Using the *forcemigrate* option might cause the upgraded host to not boot properly, to exhibit system instability, or to lose functionality.

%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.

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

`=[python|busybox]`

`--timeout=secs` Specifies a timeout for running the script. If the script is not finished when the timeout expires, the script is forcefully terminated.

`--ignorefailure` If true, the installation is considered a success even if the *%post* script terminated with an error.

=[true|false]

%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.

--interpreter Specifies an interpreter to use. The default is busybox.
=[python|busybox]

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.

Differences Between ESXi 4.x and ESXi 5.x Scripted Installation and Upgrade Commands

Before you perform a scripted ESXi installation or upgrade, if you are familiar with ESXi version 4.x scripted installation, note the differences between ESXi 4.x and ESXi 5.x scripted installation and upgrade commands.

In ESXi 5.x, because the installation image is loaded directly into the host RAM when the host boots, you do not need to include the location of the installation media in the installation script.

ESXi 5.x supports scripted upgrades in addition to scripted installation.

Command differences are noted in the following summary.

accepteula OR vmaccepteula	Only in ESXi
autopart	Deprecated and replaced with <code>install</code> , <code>upgrade</code> , or <code>installorupgrade</code> .
auth OR authconfig	Not supported in ESXi 5.x.
bootloader	Not supported in ESXi 5.x.
esxlocation	Deprecated and unused in ESXi.
firewall	Not supported in ESXi 5.x.
firewallport	Not supported in ESXi 5.x.
install , installorupgrade , upgrade	These commands replace the deprecated <code>autopart</code> command. Use one of these command to specify the disk to partition, and the <code>part</code> command to create the vmfs datastore. <code>installorupgrade</code> and <code>upgrade</code> are newly supported in ESXi 5.x.
serialnum	Deprecated in ESXi 5.0.x. Supported in ESXi 5.1.
vmserialnum	Deprecated in ESXi 5.0.x. Supported in ESXi 5.1.
timezone	Not supported in ESXi 5.x.
virtualdisk	Not supported in ESXi 5.x.
zerombr	Not supported in ESXi 5.x.

%firstboot	--level option not supported in ESXi 5.x.
%packages	Not supported in ESXi 5.x.

Disk Device Names

The `install`, `upgrade`, and `installorupgrade` installation script commands require the use of disk device names.

Table 7-3. Disk Device Names

Format	Examples	Description
VML	vml.00025261	The device name as reported by the vmkernel
MPX	mpx.vmhba0:C0:T0:L0	The device name

NOTE When you perform a scripted upgrade from ESX 4.x to ESXi 5.x, the MPX and VML disk names change, which might cause the upgrade to fail. To avoid this problem, use Network Address Authority Identifiers (NAA IDs) for the disk device instead of MPX and VML disk names.

After you obtain the NAA ID or VML number, typically from the BIOS of the storage adapter, you can modify the kickstart file (`ks.cfg`) for each host to identify the disk device by the NAA ID or VML number.

Some devices do not provide an NAA ID. In these circumstances, an MPX Identifier is generated by ESXi to represent the LUN or disk. The identifier takes a form similar to the canonical name of previous versions of ESXi with the `mpx.` prefix. This identifier can be used exactly as the NAA ID. See Knowledge Base article [1014953](#).

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` 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
kernel=FILEPATH
kernelopt=STRING
modules=FILEPATH1 --- FILEPATH2... --- FILEPATHn

# Any other line must remain unchanged.
```

The commands in `boot.cfg` configure the boot loader.

Table 7-4. Commands in `boot.cfg`.

Command	Description
<code>title=STRING</code>	Sets the boot loader title to <i>STRING</i> .
<code>kernel=FILEPATH</code>	Sets the kernel path to <i>FILEPATH</i> .

Table 7-4. Commands in `boot.cfg` . (Continued)

Command	Description
<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,”](#) on page 126, [“PXE Boot the ESXi Installer by Using PXELINUX and a PXE Configuration File,”](#) on page 130, [“PXE Boot the ESXi Installer by Using PXELINUX and an isolinux.cfg PXE Configuration File,”](#) on page 131, and [“PXE Booting the ESXi Installer,”](#) on page 127.

Install, Upgrade, or Migrate ESXi from a CD or DVD Using a Script

You can install, upgrade, or migrate ESXi from a CD/DVD drive 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,”](#) on page 126.


Prerequisites

Before you run the scripted installation, upgrade, or migration, verify that the following prerequisites are met:

- The system on which you are installing, upgrading, or migrating meets the hardware requirements. See [“ESXi Hardware Requirements,”](#) on page 13.
- You have the ESXi installer ISO on an installation CD/DVD. See [“Download and Burn the ESXi Installer ISO Image to a CD or DVD,”](#) on page 123.
- 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,”](#) on page 142.
- You have selected a boot command to run the scripted installation, upgrade or migration. See [“Enter Boot Options to Start an Installation or Upgrade Script,”](#) on page 140. For a complete list of boot commands, see [“Boot Options,”](#) on page 141.

Procedure

- 1 Boot the ESXi installer from the CD or DVD using the local CD/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, Upgrade, or Migrate ESXi from a USB Flash Drive Using a Script

You can install, upgrade, or migrate ESXi from a USB flash drive using a script that specifies the installation or upgrade options.

Supported boot options are listed in [“Boot Options,”](#) on page 141.


Prerequisites

Before running the scripted installation, upgrade, or migration, verify that the following prerequisites are met:

- The system that you are installing, upgrading, or migrating to ESXi meets the hardware requirements for the installation or upgrade. See [“ESXi Hardware Requirements,”](#) on page 13.
- 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,”](#) on page 123.
- 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,”](#) on page 142.
- 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,”](#) on page 140.

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 PXE Booting the Installer

ESXi 5.x provides many options for PXE booting the installer and using an installation or upgrade script.

- For information about setting up a PXE infrastructure, see [“PXE Booting the ESXi Installer,”](#) on page 127.
- For information about creating and locating an installation script, see [“About Installation and Upgrade Scripts,”](#) on page 142.
- For specific procedures to PXE boot the ESXi installer and use an installation script, see one of the following topics:
 - [“PXE Boot the ESXi Installer by Using PXELINUX and an isolinux.cfg PXE Configuration File,”](#) on page 131
 - [“PXE Boot the ESXi Installer by Using PXELINUX and a PXE Configuration File,”](#) on page 130
 - [“PXE Boot the ESXi Installer Using gPXE,”](#) on page 133
- For information about using Auto Deploy to perform a scripted installation by PXE booting, see [“Installing ESXi Using vSphere Auto Deploy,”](#) on page 155.

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. Auto Deploy can be used for stateless caching or stateful installs.

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.

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, and a vCenter Server location (folder or cluster) for each host.

Introduction to Auto Deploy

When you start a physical host that is set up for Auto Deploy, Auto Deploy uses a 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 Auto Deploy server manages state information for each host.

State Information for ESXi Hosts

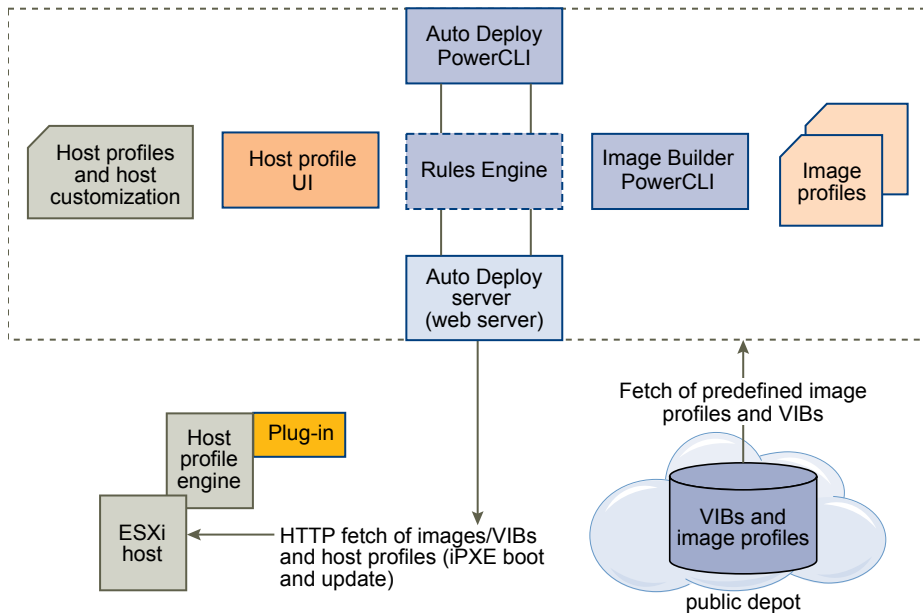
Auto Deploy stores the information for the ESXi hosts to be provisioned in different locations. Information about the location of image profiles and host profiles is initially specified in the rules that map machines to image profiles and host profiles.

Table 7-5. Auto Deploy Stores Information for Deployment

Information Type	Description	Source of State Information
Image state	Executable software to run on an ESXi host.	Image profile, created with Image Builder PowerCLI.
Configuration state	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	Runtime state that is generated by the running software, for example, generated private keys or runtime databases.	Stored in host memory and lost during reboot.
Virtual machine state	Virtual machines stored on a host and virtual machine autostart information (subsequent boots only).	vCenter Server must be available to supply virtual machine information to 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.	<p>You can create a host profile that requires user input for certain values.</p> <p>When Auto Deploy applies a host profile that requires an answer to a host, the host comes up in maintenance mode. Use the host profiles interface to check the host profile compliance, and respond to the prompt to customize the host.</p> <p>The host customization information is stored with the host.</p>

Auto Deploy Architecture

The Auto Deploy infrastructure consists of several components.

Figure 7-1. vSphere Auto Deploy Architecture

Auto Deploy server

Serves images and host profiles to ESXi hosts. The Auto Deploy server is at the heart of the Auto Deploy infrastructure.

Auto Deploy rules engine

Tells the Auto Deploy server which image profile and which host profile to serve to which host. Administrators use the Auto Deploy PowerCLI to define the rules that assign image profiles and host profiles to hosts.

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 the Image Builder PowerCLI to examine the depot and the 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.

Host profiles

Define machine-specific configuration such as networking or storage setup. Administrators create host profiles by using the host profile UI. 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.

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. See [“Host Customization in the vSphere Web Client,”](#) on page 195.

Host customization was called answer file in earlier releases of Auto Deploy.

Rules and Rule Sets

You specify the behavior of the Auto Deploy server by using a set of rules written in Power CLI. The Auto Deploy rules engine checks the rule set for matching host patterns to decide which items (image profile, host profile, or vCenter Server location) 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 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 Auto Deploy PowerCLI cmdlets 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.

NOTE You must test and repair rule compliance for any host managed by a vCenter Server system even if those hosts were not added to the vCenter Server system by Auto Deploy. See [“Test and Repair Rule Compliance,”](#) on page 174.

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 Auto Deploy PowerCLI cmdlets. 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. By default, Auto Deploy uses the name of the rule for the copy and hides the original rule.

Active Rule Set

When a newly started host contacts the Auto Deploy server with a request for an image profile, the Auto Deploy server checks the active rule set for matching rules. The image profile, host profile, and vCenter Server inventory location 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 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 Auto Deploy PowerCLI 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 Use cmdlets that execute 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 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 Auto Deploy with PowerCLI Cmdlets,”](#) on page 171.

Auto Deploy Boot Process

When you boot a host that you want to provision or reprovision with vSphere Auto Deploy, the Auto Deploy infrastructure supplies the image profile and, optionally, a host profile and a vCenter Server location for that host.

The boot process is different for hosts that have not yet been provisioned with Auto Deploy (first boot) and for hosts that have been provisioned with 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,”](#) on page 165.

- 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.
- Ensure that the Auto Deploy server has an IPv4 address. PXE booting is supported only with IPv4. Other components in your Auto Deploy infrastructure can communicate either with IPv4 or with 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 the Image Builder PowerCLI, and place the image profile in a depot that the Auto Deploy server can access. The image profile must include a base ESXi VIB.

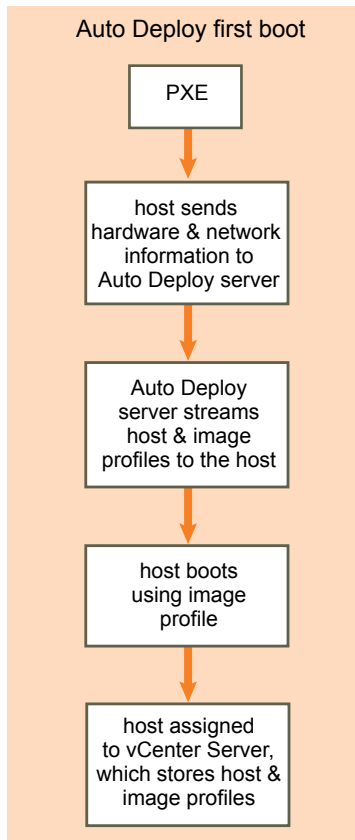
- (Optional) If you have a reference host in your environment, export the host profile of the reference host and define a rule that applies the host profile to one or more hosts. See [“Setting Up an Auto Deploy Reference Host,”](#) on page 185.
- Specify rules for the deployment of the host and add the 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 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 Auto Deploy server. The HTTP request includes hardware and network information.
- 4 In response, the 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 Auto Deploy server provided a host profile, the host profile is applied to the host.
- 6 Auto Deploy adds the host to the vCenter Server system that 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, 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 \(First Boot\),”](#) on page 175.

Figure 7-2. Auto Deploy Installation, First Boot**Subsequent Boots Without Updates**

For hosts that are provisioned with Auto Deploy and managed by a vCenter Server system, subsequent boots can become completely automatic.

- 1 The administrator reboots the host.
- 2 As the host boots up, 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 Auto Deploy and is provisioned with an image profile. The host continues to contact the Auto Deploy server until Auto Deploy reconnects to the vCenter Server system.

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,”](#) on page 176.

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,”](#) on page 178.

Subsequent Boots With Updates

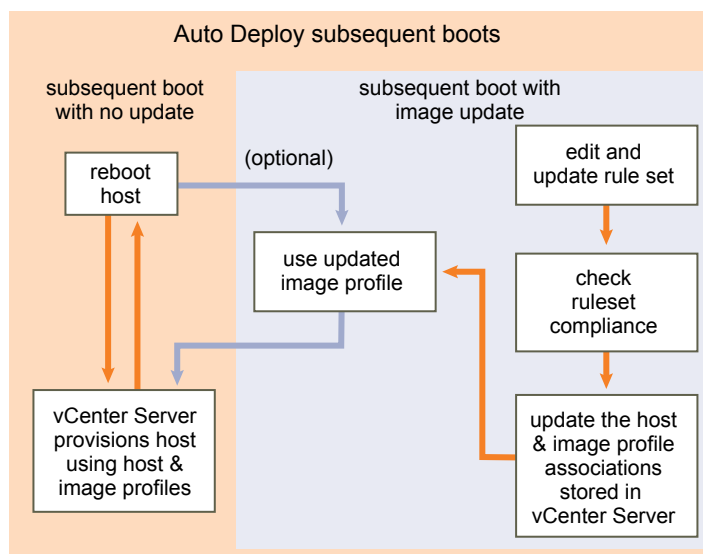
You can change the image profile, host profile, or vCenter Server location for hosts. The process includes changing rules and testing and repairing the host's rule compliance.

- 1 The administrator uses the `Copy-DeployRule` PowerCLI cmdlet to copy and edit one or more rules and updates the rule set. See [“Auto Deploy Roadmap,”](#) on page 162 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, or vCenter Server location 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,”](#) on page 178.

See [“Test and Repair Rule Compliance,”](#) on page 174.

Figure 7-3. Auto Deploy Installation, Subsequent Boots



Provisioning of Systems that Have Distributed Switches

You can configure the host profile of an 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 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 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.

Auto Deploy Roadmap and Cmdlet Overview

To be successful with Auto Deploy, you have to know the tasks involved in provisioning hosts, understand the Auto Deploy components and their interaction, and know the PowerCLI cmdlets.

Auto Deploy Roadmap

Getting started with Auto Deploy requires that you learn how Auto Deploy works, install the Auto Deploy server, install PowerCLI, write PowerCLI rules that provision hosts, and turn on your hosts to be booted with the image profile you specify. Customizations of the image profile, host profile, and vCenter Server location are supported.

See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for a step-by-step exercise that helps you set up your first Auto Deploy environment on a Windows 2008 system.

To successfully provision the hosts in your environment with Auto Deploy you can follow a few steps discussed in more detail in this document.

- 1 Install the Auto Deploy server.

Windows

The Auto Deploy server is included with the vCenter Server installation media. You can install the Auto Deploy server on the same system as vCenter Server or on a different system.

vCenter Server appliance

The vCenter Server appliance includes vCenter Server and the Auto Deploy server. The Auto Deploy server on the vCenter Server appliance is disabled by default.

You can use both servers on the appliance, use a standalone vCenter Server installation with Auto Deploy on the appliance, or use a standalone Auto Deploy installation with the vCenter Server appliance. See [“Using Auto Deploy with the VMware vCenter Server Appliance,”](#) on page 193 for configuration information.

NOTE You cannot use more than one Auto Deploy server with one vCenter Server system.

See [“Prepare Your System and Install the Auto Deploy Server,”](#) on page 165 for information on installing the software, setting up the DHCP server, and downloading the TFTP configuration file.

- 2 Install PowerCLI, which includes Auto Deploy and Image Builder cmdlets, and set up remote signing.
See [“Install PowerCLI and Prerequisite Software,”](#) on page 168 and [“Using Auto Deploy Cmdlets,”](#) on page 168.
- 3 Find the image profile that includes the VIBs that you want to deploy to your hosts.
 - In most cases, you add the depots that contain the software that you are interested in to your PowerCLI session, and then select an image profile from one of those depots.
 - To create a custom image profile, use 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.

Using Image Builder for customization is required 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.

- 4 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,”](#) on page 171.

NOTE 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. See [“Set Up Host Profiles for Static IP Addresses in the vSphere Web Client,”](#) on page 192.

- 5 Turn on the host to have Auto Deploy provision the host with the specified image profile.
- 6 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 an Auto Deploy Reference Host,”](#) on page 185.
- 7 Create and export a host profile for the reference host.
See the *Host Profiles* documentation.
- 8 To provision multiple hosts, you can use the `Copy-DeployRule` cmdlet.

You can revise the rule to assign not only an image profile but also a host profile and a cluster location .

```
Copy-DeployRule -DeployRule "testrule" -ReplaceItem
my_host_profile_from_reference_host,my_targetcluster
-ReplacePattern "ipv4=192.XXX.1.10-192.XXX.1.20"
```

my_host_profile_from_reference_host is the name of the reference host profile. *my_targetcluster* is the name of the target cluster.

- 9 Turn on the hosts you want to provision.
If the hosts that are specified by the pattern are not currently managed by a vCenter Server system, Auto Deploy provisions them with the already stored image profile and the specified host profile and adds them to the target cluster.
- 10 Check 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.

Remedy host customization and compliance problems and reboot hosts until all hosts meet the requirements.

Read [“Understanding vSphere Auto Deploy,”](#) on page 155 for an introduction to the boot process, a discussion of differences between first and subsequent boots, and an overview of using host customization.

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 VMware PowerCLI.

If you are new to PowerCLI, read the PowerCLI documentation and review [“Using Auto Deploy Cmdlets,”](#) on page 168. 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 Auto Deploy cmdlets, provide all parameters on the command line when you invoke the cmdlet. Supplying parameters in interactive mode is not recommended.

Table 7-6. Rule Engine PowerCLI Cmdlets

Command	Description
<code>Get-DeployCommand</code>	Returns a list of Auto Deploy cmdlets.
<code>New-DeployRule</code>	Creates a new rule with the specified items and patterns.
<code>Set-DeployRule</code>	Updates an existing rule with the specified items and patterns. You cannot update a rule that is part of a rule set.
<code>Get-DeployRule</code>	Retrieves the rules with the specified names.
<code>Copy-DeployRule</code>	Clones and updates an existing rule.
<code>Add-DeployRule</code>	Adds one or more rules to the working rule set and, by default, also to the active rule set. Use the <code>NoActivate</code> parameter to add a rule only to the working rule set.
<code>Remove-DeployRule</code>	Removes one or more rules from the working rule set and from the active rule set. Run this command with the <code>-Delete</code> parameter to completely delete the rule.
<code>Set-DeployRuleset</code>	Explicitly sets the list of rules in the working rule set.
<code>Get-DeployRuleset</code>	Retrieves the current working rule set or the current active rule set.
<code>Switch-ActiveDeployRuleset</code>	Activates a rule set so that any new requests are evaluated through the rule set.
<code>Get-VMHostMatchingRules</code>	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.
<code>Test-DeployRulesetCompliance</code>	Checks whether the items associated with a specified host are in compliance with the active rule set.
<code>Repair-DeployRulesetCompliance</code>	Given the output of <code>Test-DeployRulesetCompliance</code> , 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.
<code>Apply-EsxImageProfile</code>	Associates the specified image profile with the specified host.
<code>Get-VMHostImageProfile</code>	Retrieves the image profile in use by a specified host. This cmdlet differs from the <code>Get-EsxImageProfile</code> cmdlet in the Image Builder PowerCLI.
<code>Repair-DeployImageCache</code>	Use this cmdlet only if the Auto Deploy image cache is accidentally deleted.
<code>Get-VMHostAttributes</code>	Retrieves the attributes for a host that are used when the Auto Deploy server evaluates the rules.
<code>Get-DeployMachineIdentity</code>	Returns a string value that Auto Deploy uses to logically link an ESXi host in vCenter to a physical machine.
<code>Set-DeployMachineIdentity</code>	Logically links a host object in the vCenter Server database to a physical machine. Use this cmdlet to add hosts without specifying rules.

Table 7-6. Rule Engine PowerCLI Cmdlets (Continued)

Command	Description
Get-DeployOption	Retrieves the 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 Auto Deploy. 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.

Preparing for vSphere Auto Deploy

Before you can start to use vSphere Auto Deploy, you must prepare your environment. You start with server setup and hardware preparation. You must register the Auto Deploy software with the vCenter Server system that you plan to use for managing the hosts you provision, and install the VMware PowerCLI.

- [Prepare Your System and Install the Auto Deploy Server](#) on page 165
Before you turn on a host for PXE boot with vSphere Auto Deploy, you must install prerequisite software and set up the DHCP and TFTP servers that Auto Deploy interacts with.
- [Install PowerCLI and Prerequisite Software](#) on page 168
Before you can run Auto Deploy cmdlets to create and modify the rules and rule sets that govern Auto Deploy behavior, you must install vSphere PowerCLI and all prerequisite software. The Auto Deploy cmdlets are included with the PowerCLI installation.
- [Using Auto Deploy Cmdlets](#) on page 168
Auto Deploy cmdlets are implemented as Microsoft PowerShell cmdlets and included in vSphere PowerCLI. Users of Auto Deploy cmdlets can take advantage of all PowerCLI features.
- [Set Up Bulk Licensing](#) on page 169
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 Auto Deploy.

Prepare Your System and Install the Auto Deploy Server

Before you turn on a host for PXE boot with vSphere Auto Deploy, you must install prerequisite software and set up the DHCP and TFTP servers that Auto Deploy interacts with.

Prerequisites

- Ensure that the hosts that you will provision with Auto Deploy meet the hardware requirements for ESXi.

See [“ESXi Hardware Requirements,”](#) on page 13.

NOTE You cannot provision EFI hosts with Auto Deploy unless you switch the EFI system to BIOS compatibility mode.

- Ensure that the ESXi hosts have network connectivity to vCenter Server and that all port requirements are met.

See [“Required Ports for vCenter Server,”](#) on page 23.

- If you want to use VLANs in your Auto Deploy environment, you must set up the end to end networking properly. When the host is PXE booting, the UNDI 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 BIOS. 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.
- Ensure that you have enough storage for the Auto Deploy repository. The 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 2GB to have enough room for four image profiles and some extra space. Each image profile requires approximately 350MB. Determine how much space to reserve for the Auto Deploy repository by considering how many image profiles you expect to use.

- Obtain the vCenter Server installation media, which include the Auto Deploy installer, or deploy the vCenter Server Appliance.

See [Chapter 4, “Installing vCenter Server,”](#) on page 69.

See [“Using Auto Deploy with the VMware vCenter Server Appliance,”](#) on page 193.

- Ensure that a TFTP server is available in your environment. If you require a supported solution, purchase a supported TFTP server from your vendor of choice.
- 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 Auto Deploy setup, replace the `gpxelinux.0` file name with `undionly.kpxe.vmw-hardwired`.
- Secure your network as you would for any other PXE-based deployment method. Auto Deploy transfers data over SSL to prevent casual interference and snooping. However, the authenticity of the client or the Auto Deploy server is not checked during a PXE boot. .

NOTE Auto Deploy is not supported with NPIV (N_Port ID Virtualization).

- 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 and set up your first host so 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,”](#) on page 187.

See also [“Install or Upgrade vSphere ESXi Dump Collector,”](#) on page 104.

- Auto Deploy does not support a pure IPv6 environment because the PXE boot specifications do not support IPv6. However, after the initial PXE boot state, the rest of the communication can happen over IPv6. You can register Auto Deploy to the vCenter Server system with IPv6, and you can set up the host profiles to bring up hosts with IPv6 addresses. Only the initial boot process requires an IPv4 address.

Procedure

- 1 Install the vSphere Auto Deploy server as part of a vCenter Server installation or standalone on a Windows system, or deploy the vCenter Server Appliance to an ESXi system of your choice.

Location	Description
vCenter Server system	Use the vCenter Server installation media to install Auto Deploy on the same host as the vCenter Server system itself. That vCenter Server system manages all hosts that you provision with this Auto Deploy installation. See “Install or Upgrade vSphere Auto Deploy,” on page 107.
Windows system	Use the vCenter Server installation media to install Auto Deploy on a Microsoft Windows system that does not have a vCenter Server system installed. The installer prompts you for a vCenter Server system to register Auto Deploy with. That vCenter Server system manages all hosts that you provision with this Auto Deploy installation. See “Install or Upgrade vSphere Auto Deploy,” on page 107.
vCenter Server Appliance	Deploy the vCenter Server Appliance to the ESXi host of your choice. The appliance includes an Auto Deploy server, which is disabled by default. By default, the vCenter Server system on the appliance manages all hosts you provision with the appliance Auto Deploy installation. Other configurations are supported. See “Using Auto Deploy with the VMware vCenter Server Appliance,” on page 193.

- 2 Configure the TFTP server.
 - a In a vSphere Web Client connected to the vCenter Server system that Auto Deploy is registered with, 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 Log** to download the TFTP configuration file and unzip the file to the directory in which your TFTP server stores files.
- 3 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 `undionly.kpxe.vmw-hardwired` in the DHCP option 67 (frequently called boot-filename).
- 4 Set each host you want to provision with Auto Deploy to network boot or PXE boot, following the manufacturer's instructions.
- 5 Locate the image profile that you want to use and the depot in which it is located.

In most cases, you point to an image profile that VMware makes available in a public depot. If you want to include custom VIBs with the base image, you can use the Image Builder PowerCLI to create an image profile and use that image profile. See the *Image Builder PowerCLI* documentation.
- 6 Write a rule that assigns an image profile to hosts.
- 7 (Optional) 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 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/`.

When you start a host that is set up for Auto Deploy, the host contacts the DHCP server and is directed to the Auto Deploy server, which provisions the host with the image profile specified in the active rule set.

What to do next

- Install PowerCLI. See [“Install PowerCLI and Prerequisite Software,”](#) on page 168.
- Use the PowerCLI cmdlets to define a rule that assigns an image profile and optional host profile to the host. See [“Prepare Your System and Install the Auto Deploy Server,”](#) on page 165.
- (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.
- If you want to have 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 [“Consider and Implement Your Partitioning Strategy,”](#) on page 191.
- If you have to configure host-specific information, set up the host profile of the reference host to prompt for user input. See [“Host Customization in the vSphere Web Client,”](#) on page 195.

Install PowerCLI and Prerequisite Software

Before you can run Auto Deploy cmdlets to create and modify the rules and rule sets that govern Auto Deploy behavior, you must install vSphere PowerCLI and all prerequisite software. The Auto Deploy cmdlets are included with the PowerCLI installation.

You install vSphere PowerCLI and prerequisite software on a Microsoft Windows system. See the Microsoft Web site for information about installing the Microsoft software. See the *vSphere PowerCLI Installation Guide* for detailed instructions for PowerCLI installation.

Procedure

- 1 Verify that Microsoft .NET 2.0 is installed, or install it from the Microsoft Web site following the instructions on that Web site.
- 2 Verify that Microsoft Powershell 2.0 is installed, or install it from the Microsoft Web site following the instructions on that Web site.
- 3 Install vSphere vSphere PowerCLI, which includes the Auto Deploy cmdlets.

What to do next

Review [“Using Auto Deploy Cmdlets,”](#) on page 168. If you are new to PowerCLI, read the PowerCLI documentation.

Use Auto Deploy cmdlets and other PowerCLI cmdlets and PowerShell cmdlets to manage Auto Deploy rules and rule sets. Use `Get-Help <cmdlet_name>` at any time for command-line help.

Using Auto Deploy Cmdlets

Auto Deploy cmdlets are implemented as Microsoft PowerShell cmdlets and included in vSphere PowerCLI. Users of Auto Deploy cmdlets can take advantage of all PowerCLI features.

Experienced PowerShell users can use 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-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`.

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 documentation pass in parameters by name.

Passing Parameters as Objects

You can pass parameters as objects if you want to do 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.

```
$str = Test-DeployRuleSetCompliance MyEsxi42
```

- 2 Display the itemlist property of the object to see the difference between what is in the rule set and what the host is currently using.

```
$str.itemlist
```

- 3 Remediate the host to use the revised rule set by passing the object to a call to Repair-DeployRuleSetCompliance.

```
Repair-DeployRuleSetCompliance $str
```

The example remediates the host the next time you boot the host.

Setting Properties to Support Remote Signing

For security reasons, Windows PowerShell supports an execution policy feature. It determines whether scripts are allowed to run and whether they must be digitally signed. By default, the execution policy is set to Restricted, which is the most secure policy. If you want to run scripts or load configuration files, you can change the execution policy by using the Set-ExecutionPolicy cmdlet. To do this, type the following in the vSphere PowerCLI console window.

```
Set-ExecutionPolicy RemoteSigned
```

If the command is successful, you can run scripts and load configuration files. For more information about the execution policy and digital signing in Windows PowerShell, use the following command.

```
Get-Help About_Signing
```

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 Auto Deploy.

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

Install PowerCLI. See [“Install PowerCLI and Prerequisite Software,”](#) on page 168.

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 the vCenter Server system, the host is assigned a license key. A license key that is assigned through the 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.

Procedure

- 1 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 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 Auto Deploy with PowerCLI Cmdlets

You can use Auto Deploy PowerCLI cmdlets to create rules that associate hosts with image profiles, host profiles, and a location on the vCenter Server target. You can also update hosts by testing rule compliance and repairing compliance issues.

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 Auto Deploy.

Prerequisites

- Install VMware PowerCLI and all prerequisite software.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Auto Deploy Cmdlets,”](#) on page 168.

Procedure

- 1 Run the `Connect-VIServer` PowerCLI cmdlet to connect to the vCenter Server system that Auto Deploy is registered with.

Connect-VIServer 192.XXX.X.XX

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 using the Image Builder PowerCLI.
- 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:\file_path\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:Xa:Xb:Xc:Xx:XX
* . model=MyVendorModel
* . oemstring=Product ID: XXXXXX-XXX
* . serial=XX XX XX XX XX XX...
* . uuid=XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX
* . vendor=MyVendor
*****
```

What to do next

- For hosts already provisioned with Auto Deploy, perform the compliance testing and repair operations to provision them with the new image profile. See [“Test and Repair Rule Compliance,”](#) on page 174.
- Turn on unprovisioned hosts to provision them with the new image profile.

Assign a Host Profile to Hosts

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.

The following procedure explains how to write a rule that assigns a host profile to hosts. To assign the host profiles to hosts already provisioned with Auto Deploy, you must also perform a test and repair cycle. See [“Test and Repair Rule Compliance,”](#) on page 174.

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

- Install vSphere PowerCLI and all prerequisite software.
- Export the host profile that you want to use.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Auto Deploy Cmdlets,”](#) on page 168.

Procedure

- 1 Run the `Connect-VIServer` PowerCLI cmdlet to connect to the vCenter Server system that Auto Deploy is registered with.

Connect-VIServer 192.XXX.X.XX

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 hosts with certain attributes, for example a range of IP addresses, are assigned to the host profile.

**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

- Upgrade existing hosts to use the new host profile by performing compliance test and repair operations on those hosts. See [“Test and Repair Rule Compliance,”](#) on page 174.
- Turn on unprovisioned hosts to provision them with the host profile.

Assign a Host to a Folder or Cluster

Auto Deploy can assign a host to a folder or cluster. When the host boots, Auto Deploy adds it to the specified location on the vCenter Server. Hosts assigned to a cluster inherit the cluster's host profile.

The following procedure explains how to write a rule that assigns a host to a folder or cluster. To assign a host already provisioned with Auto Deploy to a new folder or cluster, you must also perform a test and repair cycle. See [“Test and Repair Rule Compliance,”](#) on page 174.

NOTE The folder you select must be in a datacenter or in a cluster. You cannot assign the host to a standalone top-level folder.

Prerequisites

- Install VMware PowerCLI and all prerequisite software.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Auto Deploy Cmdlets,”](#) on page 168.

Procedure

- 1 Run the Connect-VIServer PowerCLI cmdlet to connect to the vCenter Server system that Auto Deploy is registered with.

Connect-VIServer 192.XXX.X.XX

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 datacenter 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

- Upgrade existing hosts to be added to the specified vCenter Server location by performing test and repair compliance operations on those hosts. See [“Test and Repair Rule Compliance,”](#) on page 174.
- Turn on unprovisioned hosts to add them to the specified vCenter Server location.

Test and Repair Rule Compliance

When you add a rule to the Auto Deploy rule set or make changes to one or more rules, hosts are not updated automatically. Auto Deploy applies the new rules only when you test their rule compliance and perform remediation.

This task assumes that your infrastructure includes one or more ESXi hosts provisioned with Auto Deploy, and that the host on which you installed vSphere PowerCLI can access those ESXi hosts.

Prerequisites

- Install vSphere PowerCLI and all prerequisite software.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Auto Deploy Cmdlets,”](#) on page 168.

Procedure

- 1 Use PowerCLI to check which Auto Deploy rules are currently available.

Get-DeployRule

The system returns the rules and the associated items and patterns.

- 2 Make a change to one of the available rules, for example, you might change the image profile and the name of the rule.

Copy-DeployRule -DeployRule testrule -ReplaceItem MyNewProfile

You cannot edit a rule already added to a rule set. Instead, you copy the rule and replace the item or pattern you want to change. By default, PowerCLI uses the old name for the copy and hides the old rule.

- 3 Verify that the host that you want to test rule set compliance for is accessible.

```
Get-VMHost -Name MyEsxi42
```

- 4 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
```

- 5 Examine the differences between what is in the rule set and what the host is currently using.

```
$tr.itemlist
```

The system returns a table of current and expected items.

CurrentItem	ExpectedItem
-----	-----
My Profile 25	MyProfileUpdate

- 6 Remediate the host to use the revised rule set the next time you boot the host.

```
Repair-DeployRuleSetCompliance $tr
```

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, boot your host to have Auto Deploy apply the new rule and to achieve compliance between the rule set and the host.

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, or folder or cluster location.

The 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 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 Image Builder PowerCLI cmdlets.

Prerequisites

- Make sure your host meets the hardware requirements for ESXi hosts.
See [“ESXi Hardware Requirements,”](#) on page 13.
- Prepare the system for vSphere Auto Deploy (see [“Preparing for vSphere Auto Deploy,”](#) on page 165).
- 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 Auto Deploy with PowerCLI Cmdlets,”](#) on page 171.

When setup is complete, the Auto Deploy server and PowerCLI are installed, 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 Auto Deploy server provisions the host with the image specified by the rule engine. The Auto Deploy server might also apply a host profile to the host if one is specified in the rule set. Finally, Auto Deploy adds the host to the vCenter Server system that is specified in the rule set.

- 2 (Optional) If 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,”](#) on page 176.

If you want to change the image profile, host profile, or location of the host, update the rules and perform a test and repair compliance operation. See [“Test and Repair Rule Compliance,”](#) on page 174.

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 Auto Deploy requires that you set up your environment and add rules to the rule set. See [“Preparing for vSphere Auto Deploy,”](#) on page 165.

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 Auto Deploy requires only that all prerequisites are still met. The process uses the previously assigned image profile, host profile, and vCenter Server location.

Setup includes DHCP server setup, writing rules, and making an image profile available to the Auto Deploy infrastructure.

Prerequisites

Make sure the setup you performed during the first boot operation is in place.

Procedure

- 1 Check that the image profile and host profile for the host are still available, and that the host has the identifying information (asset tag, IP address) it had during previous boot operations.

- 2 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.

- 3 Reboot the host.

The host shuts down. When the host reboots, it uses the image profile that the Auto Deploy server provides. The Auto Deploy server also applies the host profile stored on the vCenter Server system.

Reprovision a Host with a New Image Profile

You can reprovision the host with a new image profile, host profile, or vCenter Server location 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

- Create the image profile you want to boot the host with. Use the Image Builder PowerCLI, discussed in [“Using vSphere ESXi Image Builder CLI,”](#) on page 223.
- Make sure that the setup that 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 Auto Deploy is registered with.

Connect-VIServer myVCServer

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 the Image Builder PowerCLI.
- 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:\<i>file_path</i>\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 and repair rule compliance for each host that you want to deploy the image to.

See [“Test and Repair Rule Compliance,”](#) on page 174.

When you reboot hosts after compliance repair, Auto Deploy provisions the hosts 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 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 In the vSphere Web Client, remediate the host.
 - a Right click the host and click **All vCenter Actions > Host Profiles > Remediate**.
- 3 When prompted, provide the user input.

You can now direct the host to exit maintenance mode.

The host customization is saved. The next time you boot, the host customization is applied to the host.

Using Auto Deploy for Stateless Caching and Stateful Installs

The Auto Deploy stateless caching feature allows you to cache the host's image locally on the host or on a network drive and continue to provision the host with Auto Deploy. The Auto Deploy stateful installs feature allows you to install hosts over the network without setting up a complete PXE boot infrastructure. After the initial network boot, these hosts boot like other hosts on which ESXi has been installed.

- [Introduction](#) on page 179

The System Cache Configuration host profile supports stateless caching and stateful installs.

- [Understanding Stateless Caching and Stateful Installs](#) on page 180

When you want to use Auto Deploy with stateless caching or stateful installs, you must set up a host profile, apply the host profile, and set the boot order.

- [Set Up Stateless Hosts to Use Auto Deploy with Caching](#) on page 182

You can set up your system to provision hosts with Auto Deploy, and configure the hosts to use stateless caching. If the Auto Deploy server is not available when a host reboots, the host uses the cached image.

- [Enable Stateful Installs for Hosts Provisioned with Auto Deploy](#) on page 184

You can set up hosts provisioned with Auto Deploy to cache the image to disk and to use the cached image on subsequent boots. After the image is cached, the hosts act like hosts on which an image is installed.

Introduction

The System Cache Configuration host profile supports stateless caching and stateful installs.

Stateless caching is a good solution when you use the Auto Deploy infrastructure, but you require a safeguard in case the Auto Deploy server is unavailable. Hosts provisioned with stateless caching host profile settings continue to be provisioned with Auto Deploy. Stateful installs support network installation through Auto Deploy. After the initial installation, hosts that are provisioned with stateful install host profile settings will boot from disk.

Use Cases

The System Cache Configuration host profile supports the following use cases.

Hosts provisioned with 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 Auto Deploy. If the Auto Deploy server becomes unavailable, the host boots from the cache.

Hosts provisioned with Auto Deploy become stateful hosts

Set up and apply a host profile for stateful installs. When you provision a host with Auto Deploy, the image is installed on the local disk, a remote disk, or a USB drive. For subsequent boots, you boot from disk. The host no longer uses Auto Deploy.

Preparation

To successfully use stateless caching or stateful installs, decide how to set up the system and set the boot order.

Table 7-7. Preparation for Stateless Caching or Stateful Installs

Requirement or Decision	Description
Decide on VMFS partition overwrite	<p>When you install ESXi with the interactive installer, you are prompted whether you want to overwrite an existing VMFS datastore. The System Cache Configuration host profile allows you to overwrite existing VMFS partitions by selecting a check box.</p> <p>The check box 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 Auto Deploy with stateless caching, you can set up a highly available 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 becomes temporarily available.</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 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 Auto Deploy is not available, the host boots using the cache. ■ To set up 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 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 Auto Deploy server, the vCenter Server system, or both, some limitations apply when you next reboot.

- If vCenter Server is available but the Auto Deploy server is unavailable, hosts do not connect to the vCenter Server automatically. You can manually connect the hosts to the vCenter Server, or wait until the Auto Deploy server is available again.
- If both vCenter Server and vSphere Auto Deploy do not work, you can connect to each ESXi host by using the vSphere Web Client, and assign virtual machines to each host.
- If vCenter Server is not available, vSphere DRS does not work. The Auto Deploy server cannot add hosts to the vCenter Server system. You can connect to each ESXi host by using the vSphere Web Client, and assign virtual machines to each host.
- If you make changes to your setup while connectivity is lost, you lose these changes when the Auto Deploy server is restored after the outage.

Understanding Stateless Caching and Stateful Installs

When you want to use 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, 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.

- With the **Enable stateless caching on the host** host profile, Auto Deploy caches the image when you apply the host profile. No reboot is required. When you later reboot, the host continues to use the Auto Deploy infrastructure to retrieve its image. If the Auto Deploy server is not available, the host uses the cached image.

- With the **Enable stateful installs on the host** host profile, Auto Deploy installs the image. When you reboot the host, the host boots from disk, just like a host that was provisioned with the installer. Auto Deploy no longer provisions the host.

You can apply the host profile from a vSphere Web Client, or write an Auto Deploy PowerCLI rule that applies the host profile.

Each workflow supports stateless caching and stateful installs.

Table 7-8. Workflows that set up hosts for stateless caching or stateful installs

Workflow	Stateless caching	Stateful install
Apply host profile from vSphere Web Client	Apply the host profile either to individual hosts or to all hosts in a folder or cluster. No reboot required.	Apply the host profile either to individual hosts or to all hosts in a folder or cluster. Reboot is required.
Write and apply PowerCLI rule	Set up a reference host with a host profile that has the caching setup you want. Write a PowerCLI rule that provisions the host by using Auto Deploy and that applies a host profile that is set up for stateless caching. Reboot is required.	Set up a reference host with a host profile that has the caching setup you want. Write a PowerCLI rule that provisions the host by using Auto Deploy and applies a host profile that is set up for stateful installs. Reboot is required.

Applying the System Cache Configuration Host Profile from the vSphere Web Client

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 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.
- 4 A reboot is required so the changes can take effect.

Applying the System Cache Configuration with PowerCLI

You can create a host profile for a reference host and write an Auto Deploy PowerCLI rule that applies that host profile to other target hosts. The following workflow results.

- 1 You provision a reference with Auto Deploy and create a host profile to enable a form of caching.
- 2 You write a rule that provisions additional hosts with Auto Deploy and that applies the host profile of the reference host to those hosts.

- 3 Auto Deploy provisions each host with the new image profile. The exact effect of applying the host profile depends on the host profile you selected and on whether the host was previously provisioned with Auto Deploy.

Table 7-9. First Boot and Subsequent Boots Comparison

First Boot	Subsequent Boots
For stateful installs, Auto Deploy installs the image.	For stateful installs, the host boots from disk.
For stateless caching, Auto Deploy provisions the host and caches the image.	For stateless caching, Auto Deploy provisions the host. <ul style="list-style-type: none"> ■ If Auto Deploy provisioned the host before but stateless caching was not set up before, Auto Deploy caches the image. ■ If Auto Deploy provided the host before and cached the image, Auto Deploy provisions the host using the information in the rules. ■ If Auto Deploy is unavailable, the host boots from the cached image.

Set Up Stateless Hosts to Use Auto Deploy with Caching

You can set up your system to provision hosts with Auto Deploy, and configure the hosts to use stateless caching. If the Auto Deploy server is not available when a host reboots, the host uses the cached image.

A host that is set up for stateless caching uses the cached image only if the Auto Deploy server is not available when the host reboots. In all other situations, the host is provisioned with Auto Deploy. If you change the rule that applies an image profile to the host, and you perform a test and repair compliance operation, Auto Deploy provisions the host with the new image and the new image is cached.

Set up a highly available Auto Deploy infrastructure to guarantee that virtual machines are migrated to the host if the host reboots. Because vCenter Server assigns virtual machines to the host, vCenter Server must be available. See [“Set up a Highly Available Auto Deploy Infrastructure,”](#) on page 202.

You can set up your environment for stateless caching by applying host profiles directly or by using PowerCLI rules.

Table 7-10. Setting up hosts for stateless caching or stateful installs

Workflow	Stateless caching	Stateful install
Apply host profile directly	Apply the host profile either to individual hosts or to all hosts in a folder or cluster. See “Configure a Host Profile to Use Stateless Caching,” on page 183.	Apply the host profile either to individual hosts or to all hosts in a folder or cluster. See “Configure a Host Profile to Enable Stateful Installs,” on page 184.
Write and apply PowerCLI rules	Set up a reference host with a host profile that has the caching setup you want. Write an Auto Deploy PowerCLI rule that provisions the host and that applies a host profile that is set up for stateless caching. See “Assign a Host Profile to Hosts,” on page 172.	Set up a reference host with a host profile that has the caching setup you want. Write an Auto Deploy PowerCLI rule that provisions the host and that applies a host profile that is set up for stateful installs. See “Assign a Host Profile to Hosts,” on page 172.

Prepare for Auto Deploy with Stateless Caching

Before you can start provisioning a host that uses stateless caching with Auto Deploy, you must verify that your environment is set up for Auto Deploy, prepare Auto Deploy PowerCLI rules, and set the host boot order.

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 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, protect other elements of your infrastructure. See [“Set up a Highly Available Auto Deploy Infrastructure,”](#) on page 202.

Procedure

- 1 Set up your environment for Auto Deploy and install PowerCLI.
See [“Preparing for vSphere Auto Deploy,”](#) on page 165.
- 2 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.
- 3 Set up the host to first attempt a network boot and to boot from disk if network boot fails.
See your hardware vendor's documentation.

What to do next

Set up a host profile for stateless caching. In most cases, you set up the host profile on a reference host and apply that host profile to other hosts.

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 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.

You can configure the host profile on a single host that you want to set up to use caching. You can also create a host profile that uses caching on a reference host and apply that host profile to other hosts.

Prerequisites

Prepare your host for stateless caching. See [“Prepare for Auto Deploy with Stateless Caching,”](#) on page 182.

Procedure

- 1 In the vSphere Web Client, create a host profile.
See the *Host Profiles* documentation.
- 2 Select the host profile and click **Edit Host Profile**.
- 3 Leave the name and description and click **Next**.
- 4 Click **Advanced Configuration Settings** and click the **System Image Cache Configuration** folder.
- 5 Click the **System Image Cache Configuration** icon.
- 6 In the System Image Cache Profile Settings drop-down menu, make your selection.

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 If you selected **Enable stateless caching on the host**, specify information about the disk to use.

Option	Description
Arguments for first disk	By default, the system attempts to replace an existing ESXi installation, and then attempts to write to the local disk. 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.
Check to overwrite any VMFS volumes on the selected disk	If you click this check box, the system overwrites existing VMFS volumes if not enough space is available to store the image, image profile, and host profile.

- 8 Click **Finish** to complete the host profile configuration.
- 9 Apply the host profile with the vSphere Web Client or the vSphere PowerCLI.

Option	Description
vSphere Web Client	Use the host profiles interface of the vSphere Web Client. See the <i>Host Profiles</i> documentation.
vSphere PowerCLI	See “ Assign a Host Profile to Hosts ,” on page 172.

Enable Stateful Installs for Hosts Provisioned with Auto Deploy

You can set up hosts provisioned with Auto Deploy to cache the image to disk and to use the cached image on subsequent boots. After the image is cached, the hosts act like hosts on which an image is installed.

Prepare Hosts Provisioned with Auto Deploy for Stateful Installs

In some situations, it is useful to provision hosts with Auto Deploy and to perform all subsequent boots from disk. This approach is called Stateful Installs.

Prerequisites

Decide which disk to use for storing the image, and determine whether the new image will overwrite an existing VMFS partition.

Procedure

- Set up your environment for Auto Deploy and install PowerCLI.
See “[Preparing for vSphere Auto Deploy](#),” on page 165.
- Verify that a disk with at least 1GB of free space is available.
If the disk is not partitioned, partitioning happens when you apply the host profile.
- Set up the host to boot from disk.
See your hardware vendor's documentation.

Configure a Host Profile to Enable Stateful Installs

To set up a host provisioned with 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

Make sure that your host is configured for Auto Deploy and that you meet other prerequisites for stateful installs. See [“Prepare Hosts Provisioned with Auto Deploy for Stateful Installs,”](#) on page 184.

Procedure

- 1 In the vSphere Web Client, create a host profile.
See the *Host Profiles* documentation.
- 2 With the host profile object displayed, click the Edit host profile settings icon.
- 3 Leave the name and description and click **Next**.
- 4 Click **Advanced Configuration Settings** and click the **System Image Cache Configuration** folder.
- 5 Click the **System Image Cache Configuration** icon.
- 6 In the System Image Cache Profile Settings drop-down menu, make your selection.

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 If you select **Enable stateful installs on the host**, specify information about the disk to use.

Option	Description
Arguments for first disk	By default, the system attempts to replace an existing ESXi installation, and then attempts to write to the local disk. 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.
Check to overwrite any VMFS volumes on the selected disk	If you click this check box, the system overwrites existing VMFS volumes if not enough space is available to store the image, image profile, and host profile.

- 8 Click **Finish** to complete the host profile configuration.
- 9 Apply the host profile with the vSphere Web Client or the vSphere PowerCLI.

Option	Description
vSphere Web Client	To apply the host profile to individual hosts, use the host profiles interface of the vSphere Web Client. See the <i>Host Profiles</i> documentation.
vSphere PowerCLI	To apply the host profile to one or more hosts by using PowerCLI, see “Assign a Host Profile to Hosts,” on page 172.

Setting Up an 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 might also include setup of security, storage, networking, and ESXi Dump Collector. You can then apply the host setup to other hosts with 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 with 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 with the `esxcli system syslog vCLI` command or with the vSphere Web Client. Setting up logging on a remote host is especially important for hosts provisioned with 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 set up the reference host to use ESXi Dump Collector. See [“Configure ESXi Dump Collector with ESXCLI,”](#) on page 187. 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 Auto Deploy must have the same security settings. Make any customization in your reference host. You can, for example, set up the firewall to allow certain services access to the ESXi system. See the *vSphere Security* documentation. Security setup includes shared user access settings for all hosts. You can achieve unified user access by setting up your reference host for Microsoft Active Directory.

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 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 in the environment.

Figure 7-4. Auto Deploy Reference Host Setup

[“Configuring an Auto Deploy Reference Host,”](#) on page 187 explains how to perform this setup.

Configuring an Auto Deploy Reference Host

vSphere allows you to configure a reference host by using the vSphere Web Client, by using vCLI, or by using host profiles.

To set up a host profile, 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. You can completely set up your environment in the reference and create a host profile from the reference host for use by Auto Deploy.

vSphere Command-Line Interface

You can use vCLI commands for setup of many aspects of your host. vCLI is especially suitable for configuring some of the services in the vSphere environment. Commands include `vicfg-ntp` (set up an NTP server), `esxcli system syslog` (set up a syslog server), and `esxcli network route` (add routes and set up the default route). See [“Configure ESXi Dump Collector with ESXCLI,”](#) on page 187.

Host Profiles Interface

Best practice is to set up a host with vSphere Web Client or vCLI and create a host profile from that host. You can also configure the host profiles directly with the Host Profiles interface in the vSphere Web Client. See [“Configure Host Profiles for an Auto Deploy Reference Host with the vSphere Web Client,”](#) on page 188.

Configure ESXi Dump Collector with ESXCLI

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. You can use ESXi Dump Collector to keep core dumps on a network server for use during debugging. ESXi Dump Collector is especially useful for Auto Deploy, but is supported for any ESXi host. ESXi Dump Collector supports other customization, including sending core dumps to the local disk.

Prerequisites

Install ESXi Dump Collector, a support tool that is included with the vCenter Server `autorun.exe` application and that is also included in the vCenter Server Appliance.

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-ipv4 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. If you configure an ESXi system that is running inside a virtual machine that is using a vSphere standard switch, you must choose a VMkernel port that is in promiscuous mode.

- 2 Enable ESXi Dump Collector.

```
esxcli system coredump network set --enable true
```

- 3 (Optional) Check that ESXi Dump Collector is configured correctly.

```
esxcli system coredump network check
```

The host on which you have set up ESXi Dump Collector is set up to send core dumps to the specified server by using the specified VMkernel NIC and optional port.

What to do next

- Write a rule that applies the host profile to all hosts that you want to provision with the settings that you specified in the reference host (see [“Assign a Host Profile to Hosts,”](#) on page 172).
- For hosts that are already provisioned with Auto Deploy, perform the test and repair compliance operations to provision them with the new host profile. See [“Test and Repair Rule Compliance,”](#) on page 174.
- Turn on unprovisioned hosts to provision them with the new host profile.

Configure Host Profiles for an Auto Deploy Reference Host with the vSphere Web Client

You can set up host profiles in a reference host and apply the host profile settings to all other hosts that you provision with vSphere Auto Deploy. You can either configure the reference host and export the host profile or, for small changes, edit the host profiles directly.

Prerequisites

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 **Rules and Profiles** and click **Host Profiles**.
- 2 For a new profile, click the **Create Profile from a host** icon, or right-click a profile that you want to modify and select **Edit Host Profile**.
- 3 Customize your reference host by using vCLI, by using the client UI, or by using the Host Profiles interface.

Policy	Description
ESXi Dump Collector	Set up ESXi Dump Collector with the <code>esxcli system coredump</code> command and save the host profile (best practice), or configure the host profile directly. See “Set Up Syslog from the Host Profiles Interface in the vSphere Web Client,” on page 190.
Syslog	Set up syslog for the host with the <code>esxcli system syslog</code> command. Save the host profile (best practice) or configure the host profile directly. See “Set Up Syslog from the Host Profiles Interface in the vSphere Web Client,” on page 190.

Policy	Description
NTP	Use the <code>vicfg-ntp</code> vCLI command or the vSphere Web Client to set up a host. If you use the vSphere Web Client to start the NTP Server, make sure the startup policy for the NTP Daemon is set appropriately. <ol style="list-style-type: none"> In the vSphere Web Client, select the host. Select the Manage tab and click Time Configuration. Click Edit and click Use Network Time Protocol (Enable NTP client). Select Start and stop with host as the NTP Service Startup Policy.
Security	Set up the firewall configuration, security configuration, user configuration, and user group configuration for the reference host with the vSphere Web Client or with vCLI commands. See the <i>vSphere Security</i> documentation.
Networking and Storage	Set up the networking and storage policies for the reference host with the vSphere Web Client or vCLI command.

- Click **OK** to save the host profile settings.

What to do next

Write a rule that applies the host profile to all hosts that you want to provision with the settings that you specified in the reference host (see [“Assign a Host Profile to Hosts,”](#) on page 172). Perform a test-and-repair compliance operation.

Set Up ESXi Dump Collector from the Host Profiles Interface in the vSphere Web Client

You can set up ESXi Dump Collector for a reference host with `esxcli` or directly in the Host Profiles panels of the vSphere Web Client. You can export the host profile and write a rule that applies the profile to all hosts provisioned with Auto Deploy.

Best practice is to set up hosts to use ESXi Dump Collector with the `esxcli system coredump` command and save the host profile (see [“Configure ESXi Dump Collector with ESXCLI,”](#) on page 187). If you prefer to use a GUI, set up ESXi Dump Collector from the Host Profiles interface.

Prerequisites

Verify that at least one partition has sufficient storage capability for core dumps from multiple hosts provisioned with vSphere Auto Deploy.

Procedure

- In the vSphere Web Client, click **Rules and Profiles** and click **Host Profiles**.
- For a new profile, click the **Create Profile from a host** icon, or right-click a profile that you want to modify and select **Edit Host Profile**.
- Leave the name and description and click **Next**.
- Select **Network Configuration**.
- Select **Network Coredump Settings**.
- Click the **Enabled** check box.
- Specify the host NIC to use, the Network Coredump Server IP, and the Network Coredump Server Port.
- Click **Finish** to save the host profile settings.

What to do next

- Write a rule that applies the host profile to all hosts that you want to provision with the settings that you specified in the reference host. See [“Assign a Host Profile to Hosts,”](#) on page 172.

- For hosts already provisioned with Auto Deploy, perform the compliance testing and repair operations to provision them with the new host profile. See [“Test and Repair Rule Compliance,”](#) on page 174.
- Turn on unprovisioned hosts to provision them with the new host profile.

Set Up Syslog from the Host Profiles Interface in the vSphere Web Client

Hosts provisioned with 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 interface.

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 **Rules and Profiles** and click **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 Host Profile**.
- 4 Leave the name and description and click **Next**.
- 5 Click **Advanced Configuration Settings** click the **Advanced Options** folder, and click **Advanced configuration options**.

You can specify syslog settings from here.

- 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 plus sign.
 - b Click the new Advanced configuration option at the top of the option list and select **Configure a fixed option** from the drop-down menu.
 - c Specify Syslog.global.loghost as the option, and your host as the value.

If you are configuring an ESXi version 5.1 or later host or an ESXi 5.0 host that has syslog configured, Syslog.global.loghost is already in the list of advanced options.

- 7 Click **OK** to save the host profile settings.

What to do next

- Write a rule that applies the host profile to all hosts that you want to provision with the settings that you specified in the reference host (see [“Assign a Host Profile to Hosts,”](#) on page 172).
- For hosts already provisioned with Auto Deploy, perform the compliance testing and repair operations to provision them with the new image profile. See [“Test and Repair Rule Compliance,”](#) on page 174.
- Turn on unprovisioned hosts to provision them with the new image profile.

Set Up Networking for Your Auto Deploy Host in the vSphere Web Client

You can set up networking for your 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 Auto Deploy.

Procedure

- 1 In the vSphere Web Client, select the host and click the **Networking** tab.
- 2 Perform networking setup.
If you are using virtual switches and not vSphere Distributed Switch, do not add other VMkernel NICs to vSwitch0.
- 3 After the reference host is configured, reboot the system to verify that vmk0 is connected to the Management Network.
- 4 Create a host profile from the host.

What to do next

- Write a rule that applies the host profile to all hosts that you want to provision with the settings that you specified in the reference host. See [“Assign a Host Profile to Hosts,”](#) on page 172.
- For hosts already provisioned with Auto Deploy, perform the compliance testing and repair operations to provision them with the new host profile. See [“Test and Repair Rule Compliance,”](#) on page 174.
- Turn on unprovisioned hosts to provision them with the new host profile.

Consider and Implement Your Partitioning Strategy

By default, Auto Deploy provisions hosts only if a partition is available on the host. You can set up a reference host to auto-partition all hosts that you provision with Auto Deploy.



CAUTION If you change the default auto-partitioning behavior, 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 you wish to use as your reference host with an ESXi image by using Auto Deploy.
- 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, select the host you want to use as a reference host and click **Manage**.
- 2 Click **Settings**.
- 3 Click **System** to open the system options and click **Advanced System Settings**.
- 4 Scroll to VMkernel.Boot.autoPartition and set the value to true.
- 5 (Optional) If you want local SSDs to remain unpartitioned, scroll to VMkernel.Boot.skipPartitioningSsds and set the value to true.
- 6 If no host profile exists for your reference host, create it now.

- 7 Use the Auto Deploy PowerCLI to write a rule that applies the host profile of your reference host to all hosts immediately when they boot.

Auto-partitioning is performed when the hosts boot.

Advanced Management Tasks

In most cases, you manage your Auto Deploy environment by preparing system setup, writing rules, and provisioning hosts. In some cases, you might perform advanced management tasks such as reregistering the Auto Deploy server or assigning a static IP address to each host.

Reregister Auto Deploy

If the IP address of the vCenter Server changes, you must reregister Auto Deploy.

Regardless of whether you are using the vCenter Server Appliance or a vCenter Server system installed on Windows, you must stop the Auto Deploy process and reregister Auto Deploy if the vCenter Server IP address changes.

Procedure

- ◆ Perform the reregistration task depending on the operating system.

Platform	Task
Windows	a Stop the Auto Deploy process.
	<code>net stop vmware-autodeploy-waiter</code>
	b Edit the setup file located at <code>c:\ProgramData\VMware\VMware vCenter Auto Deploy\vmconfig-autodeploy.xml</code> and specify the new IP address.
	You do not have to edit the file if you reregister for other reasons than a new IP address.
VMware vCenter Server Appliance	c Run the <code>autodeploy-register.exe</code> command, specifying all required options.
	<code>autodeploy-register.exe -R -a vCenter-IP -p vCenter-Port -u user_name -w password -s setup-file-path</code>
	d Restart the Auto Deploy process.
	<code>net start vmware-autodeploy-waiter</code>
	a Stop the Auto Deploy process.
	<code>/etc/init.d/vmware-rbd-watchdog stop</code>
	b Run the <code>autodeploy-register</code> command, specifying all required options.
	<code>autodeploy-register -R -a vCenter-IP -p vCenter-Port -u user_name -w password -s setup-file-path</code>
	c Restart the Auto Deploy process.
	<code>/etc/init.d/vmware-rbd-watchdog start</code>
You can also use the Start ESXi Services and Stop ESXi Services buttons in the vCenter Server Appliance.	

Set Up Host Profiles for Static IP Addresses in the vSphere Web Client

By default, hosts provisioned with Auto Deploy are assigned DHCP addresses by a DHCP server. You can use the Auto Deploy host customization mechanism to assign static IP addresses to hosts.

Prerequisites

- Set up your Auto Deploy environment.

- Boot the host using Auto Deploy.
- Extract a host profile from the host.

Procedure

- 1 With vSphere Web Client, connect to the vCenter Server that manages the Auto Deploy host, select **Rules and Profiles**, and select **Host Profiles**.
- 2 Right-click the host profile you just extracted and click **Edit Host Profile**.
- 3 Use the default name and description and click **Next**.
- 4 Follow these steps to change the default IP address settings.
 - a Click **Networking configuration**.
 - b Click **Host port group**.
 - c Click **Management Network**.
 - d Click **IP address settings**
- 5 From the IPv4 address drop-down menu, select **User specified IP address to be used while applying the configuration**.
- 6 If the host is in a different subnet than the vCenter Server system, select **Network Configuration > IP route configuration > IP route config** and supply the default route in the Default IPv4 gateway field.
- 7 Select **Networking configuration** and click **DNS configuration**. In the DNS settings field, make sure the **Flag indicating if DHCP should be used** check box is not selected.
- 8 Remediate the host to update the host customization information.
 - a Right-click the host, select **All vCenter Actions > Host Profiles > Attach Host Profile**, select the profile to attach, and click **OK**.
 - b Right-click the host and select **Enter Maintenance Mode**.
 - c Right-click the host and select **All vCenter Actions > Host Profiles > Remediate**.
 - d When prompted, provide the IP address.
 - e Right-click the host and select **Exit Maintenance Mode**.

The IP address is saved as a host customization. The next time you boot, the host customization information is applied to the host.

Using Auto Deploy with the VMware vCenter Server Appliance

The VMware vCenter Server Appliance is a preconfigured Linux-based virtual machine optimized for running vCenter Server and associated services. The appliance includes an Auto Deploy server that is disabled by default.

You can use Auto Deploy with the vCenter Server Appliance in different ways.

- Use the vCenter Server system on the appliance in conjunction with the Auto Deploy server on the appliance.
- Use the vCenter Server system on the appliance in conjunction with an Auto Deploy server that you install separately on a Windows system.
- Use the Auto Deploy server on the appliance in conjunction with a vCenter Server system that you install on a different vCenter Server Appliance.

- Use the Auto Deploy server on the appliance in conjunction with a vCenter Server system that you install separately on a Windows system.

NOTE If you want to use the Auto Deploy server on the vCenter Server Appliance, you must start the service first. You can start the service from the Summary tab of the appliance.

In all cases, you are responsible for setting up your DHCP server. See [“Prepare Your System and Install the Auto Deploy Server,”](#) on page 165.

NOTE You can register only one Auto Deploy instance with a vCenter Server system, and only one vCenter Server system with an Auto Deploy server.

You can set up a mixed-mode environment that includes an IPv6 vCenter Server. The Auto Deploy server must have an IPv4 address because PXE booting is supported only with IPv4.

Set Up the vCenter Server Appliance to Use a Standalone Auto Deploy Server

The vCenter Server Appliance includes an Auto Deploy server. You can also use the appliance with a standalone Auto Deploy server.

Prerequisites

- Deploy the vCenter Server Appliance.
- Obtain the vCenter Server installation media, which include the Auto Deploy installer.
- Verify that the Windows system that you want to use for Auto Deploy meets the requirements for vCenter Server installation. Auto Deploy has the same requirements.

Procedure

- 1 If Auto Deploy is running on the vCenter Server Appliance, stop the built-in Auto Deploy server and unregister Auto Deploy from the Linux command prompt.

```
service vmware-rdb-watchdog stop
autodeploy-register --unregister -a localhost -l
```

If you never started Auto Deploy on the appliance, this step is not necessary.

- 2 Using the vCenter Server installation media, install Auto Deploy on a Windows system and provide the vCenter Server Appliance information when prompted for the vCenter Server.

Your standalone Auto Deploy installation is now associated with the vCenter Server on the vCenter Server Appliance.

Set Up Auto Deploy on the vCenter Server Appliance

By default, Auto Deploy on the vCenter Server Appliance is preconfigured to use the vCenter Server system running on the appliance. You can set up Auto Deploy to use a different vCenter Server system.

Prerequisites

- Deploy the vCenter Server Appliance.
- Install a vCenter Server system to use with Auto Deploy.

Procedure

- 1 If Auto Deploy is running on the vCenter Server Appliance, stop the built-in Auto Deploy server and unregister Auto Deploy from the Linux command prompt.

```
service vmware-rdb-watchdog stop
autodeploy-register --unregister -a localhost -l
```

If you never started Auto Deploy on the appliance, this step is not necessary.

By default, the Auto Deploy daemon is not running.

- 2 The next step depends on whether you want to confirm the vCenter Server thumbprint while registering.

- Register and pass in a thumbprint that you created earlier as part of the registration.

```
autodeploy-register --register -a vcenter-ip -u username -w password -R -T thumbprint
```

- Register without using a thumbprint.

```
autodeploy-register --register -a vcenter-ip -u username -w password -R
```

- 3 Restart the Auto Deploy daemon

```
service vmware-rbd-watchdog restart
```

Host Customization in the vSphere Web Client

To customize hosts with shared attributes, you can create a host profile in a reference host. To customize individual hosts, you can set up some fields in the host profile to prompt the user for input for each host.

Host profiles allow you to prespecify information, for example, the storage setup or Syslog setup in a reference host to and apply the host profile to a set of target hosts that share the same settings. You can also use host profiles to specify that certain settings are host dependent. If you do so, the host comes up in maintenance mode when you provision it with Auto Deploy. Remediate the host or reset the host customization to be prompted for input. The system stores your input and uses it the next time the host boots.

NOTE The host customization is not stored in a location or format that administrators can access. Use the Host Profiles UI in the vSphere Web Client to modify customization.

When the host profile is set to prompt for user input, you must specify a value in the dialog that appears when you reset the host customization. An error results if you do not specify a value.

Table 7-11. Host Profile Options that Prompt for iSCSI User Input

Information to Request User Input For		Setting the Host Profile Option
When you apply a host profile on a system that includes a profile for iSCSI, you are prompted for several properties. For many of the properties, a system default is available. For some properties, you must specify a value or an error results.		<ol style="list-style-type: none"> 1 Select Edit Host Profile, click Storage configuration, and click iSCSI Initiator Configuration. 2 Select the folder for an already enabled initiator and set up the initiator. 3 Set up the initiator. For many fields, the user is prompted as part of host customization.
IQN name	If the iSCSI setup uses an IQN name, you are prompted when you apply the host profile. You cannot continue until you provide the name.	
CHAP information	If you set up iSCSI to require CHAP authentication, you are prompted for CHAP information including the user name and the secret when you apply the host profile. You cannot continue until you provide the name.	

Table 7-12. Host Profile Options that Prompt for Storage User Input

Information to Request User Input For	Setting the Host Profile Option
You are setting up the Fixed PSP configuration and want to prompt for the adapter and target IDs for the storage arrays that should use the Fixed PSP.	<p>You can set the option only if the adapter is set up to use the Fixed PSP.</p> <ol style="list-style-type: none"> 1 Select Edit Host Profile, click Storage configuration. 2 Click Native Multipathing (NMP). 3 Click Path Selection Policy (PSP) configuration. 4 In the Preferred Path window, select Prompt the user for adapter and target IDs on the host.
Configure FCoE adapter activation based on a user-specified MAC address.	<p>You can set the option only if an activation profile exists.</p> <ol style="list-style-type: none"> 1 Select Edit Host Profile, click Storage configuration. 2 Click Software FCoE configuration. 3 Click Adapter Configuration. 4 Click the activation profile and click Policy Profile. 5 Select Activation policy based on adapter MAC address from the drop-down menu.

Table 7-13. Host Profile Options that Prompt for Security User Input

Information to Request User Input For	Setting the Host Profile Option
Administrator password for ESXi host when the host boots for the first time.	<ol style="list-style-type: none"> 1 Select Edit Host Profile, click Security Settings and click Security configuration. 2 In the right panel, select User Input Password to be Used to Configure Administrator Password from the Administrator password drop-down menu.
Preconfigures a user for the ESXi host but prompts for the password for that user on each host when the host boots for the first time.	<p>You can perform this task only if a user configuration already exists. Configure the user by selecting one of the options.</p> <ul style="list-style-type: none"> ■ Assigned fixed user configurations is available for compatibility with ESX/ESXi 4.1 system, this option displays the password in the clear. ■ Assign advanced fixed user configurations is for users of ESXi 5.0 and later systems. ■ Specify the user configuration in the profile but prompt for password during host configuration allows you to specify the information about the user but prompt for a password on each host.
Prompt the user for credentials when the host joins the Active Directory domain.	<ol style="list-style-type: none"> 1 Set the Authentication configuration profile to use a fixed domain. <ol style="list-style-type: none"> a Select Edit Host Profile, click Security and Services and click Authentication configuration. b Click Active Directory configuration. c In the Domain Name drop-down menu, select Configure a fixed domain name. 2 Set the method for joining the domain to prompt the user. <ol style="list-style-type: none"> a Select Edit Host Profile, click Security and Services and click Authentication configuration. b Click Active Directory configuration. c In the Join Domain Method drop-down menu, select Use user specified AD credentials to join the host to domain.

Table 7-14. Host Profile Options that Prompt for Networking User Input

Information to Request User Input For	Setting the Host Profile Option
Prompt the user for the MAC address for a port group. You can have the system prompt the user in all cases (User specified MAC address...) or prompt the user only if no default is available.	<ol style="list-style-type: none"> 1 Select Edit Host Profile, click Networking configuration, and click Host port group. 2 Click Management Network. 3 In the Determine how MAC address for vmknics should be decided field, and select how the system manages the MAC address. <ul style="list-style-type: none"> ■ User specified MAC Address to be used while applying the configuration ■ Prompt the user for the MAC Address if no default is available
Prompt the user for the IPv4 address for each ESXi host to which the profile is applied. You can have the system prompt the user in all cases (User specified IPv4 address...) or prompt the user only if no default is available.	<ol style="list-style-type: none"> 1 Select Edit Host Profile, click Networking configuration, and click Host port group. 2 Click Management Network and click IP address settings. 3 In the IPv4 address field, select how the system manages the IPv4 address. <ul style="list-style-type: none"> ■ User specified IPv4 Address to be used while applying the configuration ■ Prompt the user for the IPv4 Address if no default is available
Prompt the user for the IPv6 address for each ESXi host to which the profile is applied. You can have the system prompt the user in all cases (User specified IPv6 address...) or prompt the user only if no default is available.	<ol style="list-style-type: none"> 1 Select Edit Host Profile, click Networking configuration, and click Host port group. 2 Click Management Network and click IP address settings. 3 In the Static IPv6 address field, select how the system manages the IPv6 address. <ul style="list-style-type: none"> ■ User specified IPv6 Address to be used while applying the configuration ■ Prompt the user for the IPv6 Address if no default is available
Prompt the user for the DNS name of the host. You can have the system prompt the user in all cases (User specified host name...) or prompt the user only if no default is available.	<ol style="list-style-type: none"> 1 Select Edit Host Profile, click Networking configuration, and click DNS configuration. 2 In the Host name field, select how the system manages the DNS configuration. <ul style="list-style-type: none"> ■ Prompt the user for host name if default is not available ■ User specified host name to be used while applying the configuration
Prompt the user for the MAC address for a distributed switch, its port group, or one of its services. Right-click the Host virtual NIC folder icon and click the Add sub-profile icon to determine the component to which the setting is applied. You can decide to prompt the user in all cases or only if no default is available.	<ol style="list-style-type: none"> 1 Open Networking configuration. 2 Click Host virtual NIC. 3 In the Determine how MAC address for vmknics should be decided field, select how the system manages the MAC address for the distributed switch. <ul style="list-style-type: none"> ■ User specified MAC address to be used while applying the configuration ■ Prompt the user for the MAC address if no default is available

Table 7-14. Host Profile Options that Prompt for Networking User Input (Continued)

Information to Request User Input For	Setting the Host Profile Option
<p>Prompt the user for the IPv4 address for a distributed switch, its port group, or one of its services. Right-click the Host virtual NIC folder icon and click the Add sub-profile icon to determine the component to which the setting is applied.</p> <p>You can decide to prompt the user only if no default is available or in all cases.</p>	<ol style="list-style-type: none"> 1 Open Networking configuration. 2 Click Host virtual NIC. 3 Click IP address settings. 4 In the IPv4 address field, select how the system handles the IPv4 address for the distributed switch. <ul style="list-style-type: none"> ■ User specified IPv4 address to be used while applying the configuration ■ Prompt the user for IPv4 address if no default is available
<p>Prompt the user for the IPv6 address for a distributed switch, its port group, or one of its services. Right-click the Host virtual NIC folder icon and click the Add sub-profile icon to determine the component to which the setting is applied.</p> <p>You can decide to prompt the user only if no default is available or in all cases.</p>	<ol style="list-style-type: none"> 1 Open Networking configuration. 2 Open Host virtual NIC. 3 Open IP address settings. 4 In the Static IPv6 address field, select how the system manages the IPv6 address for the distributed switch. <ul style="list-style-type: none"> ■ User specified IPv6 address to be used while applying the configuration ■ Prompt the user for IPv6 address if no default is available

Auto Deploy Best Practices and Security Consideration

Follow best practices when installing vSphere Auto Deploy and when using Auto Deploy with other vSphere components. Set up a highly available 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.

Auto Deploy Best Practices

This section discusses several Auto Deploy best practices. See the VMware Knowledge Base for additional best practice information.

Auto Deploy and vSphere HA Best Practices

You can improve the availability of the virtual machines running on hosts provisioned with Auto Deploy by following best practices.

- Some environments configure the hosts provisioned with Auto Deploy with a distributed switch or configure virtual machines running on the hosts with Auto Start Manager. In those environments, deploy the vCenter Server system so that its availability matches the availability of the Auto Deploy server. Several approaches are possible.
 - In a proof of concept environment, deploy the vCenter Server system and the Auto Deploy server on the same system. In all other situations, install the two servers on separate systems.
 - Deploy vCenter Server Heartbeat.

VMware vCenter Server Heartbeat delivers high availability for VMware vCenter Server, protecting the virtual and cloud infrastructure from application, configuration, operating system, or hardware related outages.

- Deploy the vCenter Server system in 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 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 wish to use DRS in the cluster. The greater the number of hosts that are not managed by Auto Deploy the greater your resilience to host failures.

NOTE This approach is not suitable if you use Auto Start Manager because Auto Start Manager is not supported in a cluster enabled for vSphere HA.

Auto Deploy Networking Best Practices

Prevent networking problems by following Auto Deploy networking best practices.

IP Address Allocation Using DHCP reservations is recommended for address allocation. Fixed IP addresses are supported by the host customization mechanism, but providing input for each host is not recommended.

VLAN Considerations Using Auto Deploy in environments that do not use VLANs is recommended.

If you intend to use Auto Deploy in an environment that uses VLANs, you must make sure that the hosts 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 by the router, or you might be able to set the VLAN ID 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.

Auto Deploy and VMware Tools Best Practices

See the VMware Knowledge Base article 2004018 for Auto Deploy and VMware Tools best practices.

Auto Deploy Load Management Best Practice

Simultaneously booting large numbers of hosts places a significant load on the Auto Deploy server. Because 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 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 Auto Deploy server. See the VMware Technical Publications Video *Using Reverse Web Proxy Servers for Auto Deploy*.

Configure the hosts to boot off the reverse proxy by using multiple TFTP servers, one for each reverse proxy server. Finally, set up the DHCP server to send different hosts to different TFTP servers.

When you boot the hosts, the DHCP server sends them to different TFTP servers. Each TFTP server sends hosts to a different server, either the Auto Deploy server or a reverse proxy server, significantly reducing the load on the Auto Deploy server.

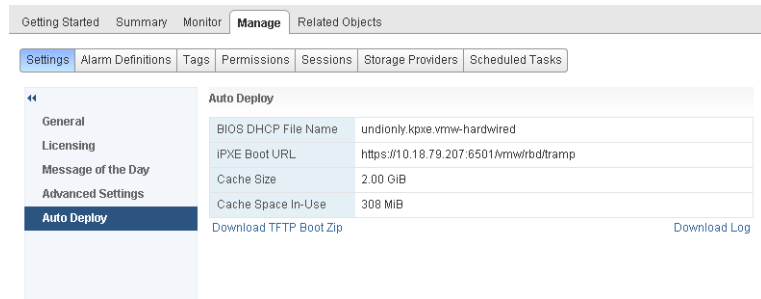
After a massive power outage, VMware recommends that you bring up the hosts on a per-cluster basis. If you bring up multiple clusters simultaneously, the Auto Deploy server might experience CPU bottlenecks. All hosts come up after a potential delay. The bottleneck is less severe if you set up the reverse proxy.

vSphere Auto Deploy Logging and Troubleshooting Best Practices

To resolve problems you encounter with vSphere Auto Deploy, use the Auto Deploy logging information from the vSphere Web Client and set up your environment to send logging information and core dumps to remote hosts.

Auto Deploy Logs

- 1 In a vSphere Web Client connected to the vCenter Server system that Auto Deploy is registered with, go to the inventory list and select the vCenter Server system.
- 2 Click the Manage tab, select Settings, and click Auto Deploy.
- 3 Click **Download TFTP Boot Log** to download the TFTP configuration file and unzip the file to the directory in which your TFTP server stores files.



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 and enables network logging and combining of logs from multiple hosts.

Setting Up ESXi Dump Collector

Hosts provisioned with Auto Deploy do not have a local disk to store core dumps on. Install ESXi Dump Collector and set up your first host so 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,”](#) on page 187.

Using 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 Auto Deploy server. [“Auto Deploy and vSphere HA Best Practices,”](#) on page 198 gives an overview of the options you have.
- Protect all other servers in your environment including the DHCP server and the TFTP server.
- Follow VMware security guidelines, including those outlined in [“Auto Deploy Security Considerations,”](#) on page 203.

Auto Deploy Device Alias Best Practices

A compliance check for a device alias profile might result in a compliance failure that cannot be remediated without modifying the hardware itself. These errors occur when the hosts are not fully homogenous (for example, the hosts contain different PCI cards or have different BIOS levels).

The following are remediation and reboot best practices for different situations:

- If the compliance check shows devices whose aliases (for example, names such as vmhba3) are different from the reference host, then remediation is possible. The remediation will cause the device to names to align exactly with the reference host.
- The remediation step (for hosts not provisioned with Auto Deploy) always requires reboot.
- The remediation step for hosts provisioned with Auto Deploy is unnecessary. Reboot the host without remediation.
- The compliance check step ignores extra devices on the host that were not on the reference host. You should choose the host with the fewest devices as the reference host.

Auto Deploy Device Alias Naming

In order to create uniform, persistent, and stable device naming across all hosts, the device alias profile should be used with homogeneous hosts only. These are hosts that are all identically configured with the same network and storage cards in the PCI bus. The device alias profile is mandatory for all hosts provisioned with Auto Deploy.

In previous versions of vSphere, devices were named based on the order in which they enumerated in the BIOS. VMware did not guarantee that devices in the same slot on the same bus would have the same name when the system rebooted if any PCI bus configuration changes occurred in the interim. Some PCI bus configuration changes might be adding or removing cards, or changes resulting from card malfunction. In vSphere 5.5, a mapping now exists between the physical location of the device (bus address) of devices and their names (aliases). This means that when ESXi hosts reboot, devices that have not physically moved are guaranteed to have the same name.

NOTE For systems with older versions of the BIOS, the older 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. Always bring the BIOS up to the latest level.

The heuristic is designed to keep built-in devices that lack proper BIOS information stable even when cards in slots are added or removed. VMware cannot guarantee this will function under all conditions (for example, if changes are made in the BIOS setting or if devices fail).

For normally installed hosts, the alias mapping is stored on the system's hard disk or USB boot device. For hosts provisioned with Auto Deploy, the mapping is stored in the Host Profile (device alias sub-profile). Because the mapping is critical to keep the device naming stable the device alias sub-profile should never be disabled for hosts provisioned with Auto Deploy.

For installed hosts, the device alias sub-profile can be used to align the device names (aliases) between multiple hosts. However, those hosts must be homogeneous for this to work. If the hosts are not homogeneous, you should disable the device alias sub-profile.

To avoid errors, VMware advises that you do not edit the device alias configuration policy in the Host Profile.

Upgrading Systems for Device Alias Profiles

In previous versions of ESXi, the device alias profile did not exist. Consider the following issues when you upgrade from previous versions of ESXi to ESXi 5.5:

- For installed hosts (host not provisioned with Auto Deploy), upgrading the ESXi host preserves aliases. After they are upgraded, aliases remain stable as long as the BIOS provides the information.
- When upgrading the ESXi host provisioned with Auto Deploy image for a stateless cluster, the aliases will not change because ESXi 5.5 uses the same algorithm to generate aliases (even though the ESXi 5.x host profile does not contain aliases). A new Host Profile should then be generated from the reference host. This Host Profile does contain aliases.

- 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 fresh install.

Set up a Highly Available Auto Deploy Infrastructure

In many production situations, a highly available Auto Deploy infrastructure is required to prevent data loss. Such an infrastructure is also a prerequisite for using Auto Deploy with stateless caching.

Prerequisites

For the management cluster, install ESXi on three hosts. Do not provision the management cluster hosts with Auto Deploy.

Procedure

- 1 Enable vSphere HA and vSphere DRS on the management cluster.
- 2 Set up the following virtual machines on the management cluster.

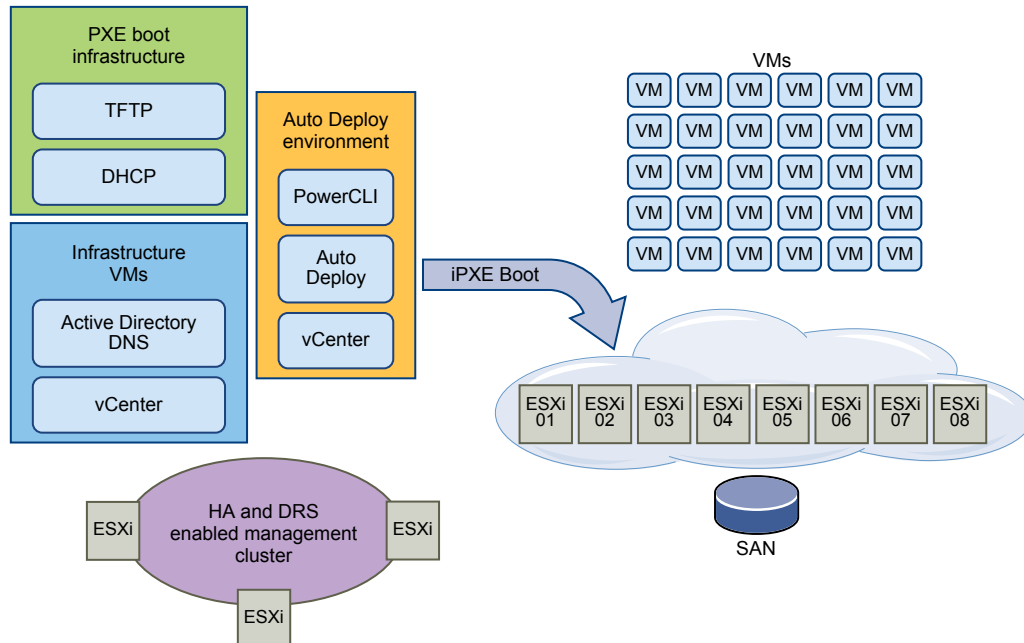
Infrastructure component	Description
PXE boot infrastructure	TFTP and DHCP servers.
Infrastructure VM	Active Directory, DNS, vCenter Server.
Auto Deploy environment	PowerCLI, Auto Deploy server, vCenter Server. Set up this environment on a single virtual machine or on three separate virtual machines in production systems.

The vCenter Server on the infrastructure virtual machine differs from the vCenter Server in the Auto Deploy environment.

- 3 Set up Auto Deploy to provision other hosts as needed.

Because the components on the management cluster are protected with vSphere HA, high availability is supported.

Figure 7-5. Highly Available Auto Deploy Infrastructure



Auto Deploy Security Considerations

Understanding potential security risks helps you set up your environment in a secure manner.

Secure your network as you would for any other PXE-based deployment method. 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.

The boot image that the 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 a host customization setting.
 - 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 Service for setting up Active Directory to avoid exposing the Active Directory passwords.

- The host's public and private SSL key and certificate are included in the boot image.

You can greatly reduce the security risk of Auto Deploy by completely isolating the network where Auto Deploy is used.

Troubleshooting Auto Deploy

The Auto Deploy troubleshooting topics offer solutions for situations when provisioning hosts with Auto Deploy does not work as expected.

Auto Deploy TFTP Timeout Error at Boot Time

A TFTP Timeout error message appears when a host provisioned by Auto Deploy boots. The text of the message depends on the BIOS.

Problem

A TFTP Timeout error message appears when a host provisioned by Auto Deploy boots. The text of the message depends on the BIOS.

Cause

The TFTP server is down or unreachable.

Solution

- ◆ Ensure that your TFTP service is running and reachable by the host that you are trying to boot.

Auto Deploy Host Boots with Wrong Configuration

A host is booting with a different ESXi image, host profile, or folder location than the one specified in the rules.

Problem

A host is booting with a different ESXi image profile or configuration than the image profile or configuration that the rules specify. For example, you change the rules to assign a different image profile, but the host still uses the old image profile.

Cause

After the host has been added to a vCenter Server system, the boot configuration is determined by the vCenter Server system. The vCenter Server system associates an image profile, host profile, or folder location with the host.

Solution

- ◆ Use the `Test-DeployRuleSetCompliance` and `Repair-DeployRuleSetCompliance` PowerCLI cmdlets to reevaluate the rules and to associate the correct image profile, host profile, or folder location with the host.

Host Is Not Redirected to Auto Deploy Server

During boot, a host that you want to provision with Auto Deploy loads iPXE. The host is not redirected to the Auto Deploy server.

Problem

During boot, a host that you want to provision with Auto Deploy loads iPXE. The host is not redirected to the AutoDeploy server.

Cause

The `tramp` file that is included in the TFTP ZIP file has the wrong IP address for the Auto Deploy server.

Solution

- ◆ Correct the IP address of the Auto Deploy server in the `tramp` file, as explained in the *vSphere Installation and Setup* documentation.

Auto Deploy Host with a Built-In USB Flash Drive Does Not Send Coredumps to Local Disk

If your Auto Deploy host has a built-in USB flash drive, and an error results in a coredump, the coredump is lost. Set up your system to use ESXi Dump Collector to store coredumps on a networked host.

Problem

If your Auto Deploy host has a built-in USB Flash, and if it encounters an error that results in a coredump, the coredump is not sent to the local disk.

Solution

- 1 Install ESXi Dump Collector on a system of your choice.
ESXi Dump Collector is included with the vCenter Server installer.
- 2 Use ESXCLI to configure the host to use ESXi Dump Collector.

```
esxcli conn_options system coredump network set IP-addr,port
esxcli system coredump network set -e true
```
- 3 Use ESXCLI to disable local coredump partitions.

```
esxcli conn_options system coredump partition set -e false
```

Package Warning Message When You Assign an Image Profile to Auto Deploy Host

When you run a PowerCLI cmdlet that assigns an image profile that is not Auto Deploy ready, a warning message appears.

Problem

When you write or modify rules to assign an image profile to one or more hosts, the following error results:

Warning: Image Profile <name-here> contains one or more software packages that are not stateless-ready. You may experience problems when using this profile with Auto Deploy.

Cause

Each VIB in an image profile has a `stateless-ready` flag that indicates that the VIB is meant for use with Auto Deploy. You get the error if you attempt to write an Auto Deploy rule that uses an image profile in which one or more VIBs have that flag set to `FALSE`.

NOTE You can use hosts provisioned with Auto Deploy that include VIBs that are not stateless ready without problems. However booting with an image profile that includes VIBs that are not stateless ready is treated like a fresh install. Each time you boot the host, you lose any configuration data that would otherwise be available across reboots for hosts provisioned with Auto Deploy.

Solution

- 1 Use Image Builder PowerCLI cmdlets to view the VIBs in the image profile.
- 2 Remove any VIBs that are not stateless-ready.
- 3 Rerun the Auto Deploy PowerCLI cmdlet.

Auto Deploy Host Reboots After Five Minutes

An Auto Deploy host boots and displays iPXE information, but reboots after five minutes.

Problem

A host to be provisioned with Auto Deploy boots from iPXE and displays iPXE information on the console. However, after five minutes, the host displays the following message to the console and reboots.

```
This host is attempting to network-boot using VMware
AutoDeploy. However, there is no ESXi image associated with this host.
Details: No rules containing an Image Profile match this
host. You can create a rule with the New-DeployRule PowerCLI cmdlet
and add it to the rule set with Add-DeployRule or Set-DeployRuleSet.
The rule should have a pattern that matches one or more of the attributes
listed below.
```

The host might also display the following details:

```
Details: This host has been added to VC, but no Image Profile
is associated with it. You can use Apply-ESXImageProfile in the
PowerCLI to associate an Image Profile with this host.
Alternatively, you can reevaluate the rules for this host with the
Test-DeployRuleSetCompliance and Repair-DeployRuleSetCompliance cmdlets.
```

The console then displays the host's machine attributes including vendor, serial number, IP address, and so on.

Cause

No image profile is currently associated with this host.

Solution

You can temporarily assign an image profile to the host by running the `Apply-ESXImageProfile` cmdlet.

You can permanently assign an image profile to the host as follows.

- 1 Run the `New-DeployRule` cmdlet to create a rule that includes a pattern that matches the host with an image profile.
- 2 Run the `Add-DeployRule` cmdlet to add the rule to a ruleset.
- 3 Run the `Test-DeployRuleSetCompliance` cmdlet and use the output of that cmdlet as the input to the `Repair-DeployRuleSetCompliance` cmdlet.

Auto Deploy Host Does Not Network Boot

The host you provision with Auto Deploy comes up but does not network boot.

Problem

When you attempt to boot a host provisioned with Auto Deploy, the host does not start the network boot process.

Cause

You did not enable your host for network boot.

Solution

- 1 Reboot the host and follow the on-screen instructions to access the BIOS configuration.
If you have an EFI host, you must switch the EFI system to BIOS compatibility mode.
- 2 In the BIOS configuration, enable Network Boot in the Boot Device configuration.

Auto Deploy Host Does Not Get a DHCP Assigned Address

The host you provision with Auto Deploy fails to get a DHCP Address.

Problem

When you attempt to boot a host provisioned with Auto Deploy, the host performs a network boot but is not assigned a DHCP address. The Auto Deploy server cannot provision the host with the image profile.

Cause

You might have a problem with the DHCP service or with the firewall setup.

Solution

- 1 Check that the DHCP server service is running on the Windows system on which the DHCP server is set up to provision hosts.
 - a Click **Start > Settings > Control Panel > Administrative Tools**.
 - b Double-click **Services** to open the Services Management panel.
 - c In the Services field, look for the DHCP server service and restart the service if it is not running.
- 2 If the DHCP server is running, recheck the DHCP scope and the DHCP reservations that you configured for your target hosts.

If the DHCP scope and reservations are configured correctly, the problem most likely involves the firewall.

- 3 As a temporary workaround, turn off the firewall to see whether that resolves the problem.
 - a Open the command prompt by clicking **Start > Program > Accessories > Command prompt**.
 - b Type the following command to temporarily turn off the firewall. Do not turn off the firewall in a production environment.


```
netsh firewall set opmode disable
```
 - c Attempt to provision the host with Auto Deploy.
 - d Type the following command to turn the firewall back on.


```
netsh firewall set opmode enable
```
- 4 Set up rules to allow DHCP network traffic to the target hosts.

See the firewall documentation for DHCP and for the Windows system on which the DHCP server is running for details.

Auto Deploy Host Cannot Contact TFTP Server

The host that you provision with Auto Deploy cannot contact the TFTP server.

Problem

When you attempt to boot a host provisioned with Auto Deploy, the host performs a network boot and is assigned a DHCP address by the DHCP server, but the host cannot contact the TFTP server.

Cause

The TFTP server might have stopped running, or a firewall might block the TFTP port.

Solution

- If you installed the WinAgents TFTP server, open the WinAgents TFTP management console and verify that the service is running. If the service is running, check the Windows firewall's inbound rules to make sure the TFTP port is not blocked. Turn off the firewall temporarily to see whether the firewall is the problem.
- For all other TFTP servers, see the server documentation for debugging procedures.

Auto Deploy Host Cannot Retrieve ESXi Image from Auto Deploy Server

The host that you provision with Auto Deploy stops at the iPXE boot screen.

Problem

When you attempt to boot a host provisioned with Auto Deploy, the boot process stops at the iPXE boot screen and the status message indicates that the host is attempting to get the ESXi image from the Auto Deploy server.

Cause

The Auto Deploy service might be stopped or the Auto Deploy server might be inaccessible.

Solution

- 1 Log in to the system on which you installed the Auto Deploy server.

- 2 Check that the Auto Deploy server is running.
 - a Click **Start > Settings > Control Panel > Administrative Tools**.
 - b Double-click **Services** to open the Services Management panel.
 - c In the Services field, look for the VMware vSphere Auto Deploy Waiter service and restart the service if it is not running.
- 3 Open a Web browser, enter the following URL, and check whether the Auto Deploy server is accessible.
`https://Auto_Deploy_Server_IP_Address:Auto_Deploy_Server_Port/vmw/rdb`

NOTE Use this address only to check whether the server is accessible.

- 4 If the server is not accessible, a firewall problem is likely.
 - a Try setting up permissive TCP Inbound rules for the Auto Deploy server port.
 The port is 6501 unless you specified a different port during installation.
 - b As a last resort, disable the firewall temporarily and enable it again after you verified whether it blocked the traffic. Do not disable the firewall on production environments.
 To disable the firewall, run **netsh firewall set opmode disable**. To enable the firewall, run **netsh firewall set opmode enable**.

Recovering from Database Corruption on the Auto Deploy Server

In some situations, you might have a problem with the Auto Deploy database. The most efficient recovery option is to replace the existing database file with the most recent backup.

Problem

When you use Auto Deploy to provision the ESXi hosts in your environment, you might encounter a problem with the Auto Deploy database.

IMPORTANT This is a rare problem. Follow all other Auto Deploy troubleshooting strategies before you replace the current database file. Rules or associations that you created since the backup you choose are lost.

Cause

This problem happens only with hosts that are provisioned with Auto Deploy.

Solution

- 1 Stop the Auto Deploy server service.
- 2 Find the Auto Deploy log by going to the Auto Deploy page in the vSphere Web Client.
- 3 Check the logs for the following message:
`DatabaseError: database disk image is malformed.`
 If you see the message, replace the existing database with the most recent backup.

- 4 Go to the Auto Deploy data directory.

Operating System	File Location
vCenter Server appliance	<code>/var/lib/rbd</code>
Microsoft Windows	The data directory you selected during installation. To find it, type the following command into a command prompt. <code>reg.exe QUERY "HKLM\SOFTWARE\WOW6432Node\VMware, Inc.\VMware vSphere Auto Deploy" /v DataPath</code>

The directory contains a file named `db`, and backup files named `db-yyy-mm-dd`.

- 5 Rename the current `db` file.
VMware Support might ask for that file if you call for assistance.
- 6 Rename the most recent backup to `db`.
- 7 Restart the Auto Deploy server service.
- 8 If the message still appears in the log, repeat the steps to use the next recent backup until Auto Deploy works without database errors.

Problems if You Upgrade vCenter Server But Do Not Upgrade Auto Deploy Server

When you upgrade vCenter Server, you can upgrade the Auto Deploy Server at the same time. If you postpone the update, problems with the vSphere HA agent might result.

Problem

When you upgrade vCenter Server, vCenter Server replaces the vSphere HA agent (`vmware-fdm`) version 5.0 with vSphere HA agent version 5.1 or later on each ESXi host. On hosts provisioned with Auto Deploy, the replacement is not permanent because no state is on the host. If vCenter Server is not available, the ESXi hosts do not have the correct vSphere HA agent and cannot join a cluster.

Cause

The Auto Deploy 5.0 server does not automatically upgrade the FDM VIB to version 5.1 or later. Unless you create a new image that includes the VIB, Auto Deploy reverts to the FDM VIB version 5.0 after reboot.

Solution

Upgrade the Auto Deploy server.

If you cannot upgrade the Auto Deploy server, you can use Image Builder PowerCLI cmdlets included with vSphere PowerCLI to create an ESXi 5.0 image profile that includes the new `vmware-fdm` VIB. You can supply your hosts with that image profile.

- 1 At the PowerCLI prompt, add the ESXi 5.0 software depot and add the software depot that contains the new `vmware-fdm` VIB.

```
Add-EsxSoftwareDepot
C:\Path\VMware-Esxi-5.0.0-buildnumber-depot.zip
```

```
Add-EsxSoftwareDepot http://vcenter_server/vSphere-HA-depot
```

- 2 Create a rule that assigns the new image profile to your hosts, and add the rule to the ruleset.

```
New-DeployRule -Name "Rule Name"
-Item "ImageName"
-Pattern "my host pattern"
Add-DeployRule -DeployRule "Rule Name"
```

- 3 Perform a test-and-repair compliance operation for the hosts to permanently include the vSphere HA agent on the hosts.

```
$result = Test-DeployRuleSetCompliance Host_list
Repair-DeployRuleSetCompliance -TestResult $result
```

Auto Deploy Proof of Concept Setup

A proof of concept setup of an Auto Deploy environment helps administrators to evaluate the product and demonstrate its capabilities to management. When you complete the proof of concept setup workflow, you have a working Auto Deploy environment that includes a reference host and one or more other target hosts.

The proof of concept setup is intended for a test or development environment, but your completed setup can be the basis for a production environment. The set of tasks starts in an environment in which no Auto Deploy components are installed. 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, you should have the following background knowledge and privileges.

- Experience with vSphere (vCenter Server, ESX, and ESXi).
- Basic knowledge of Microsoft PowerShell and vSphere PowerCLI.
- Administrator rights to the target Windows and vCenter Server systems.

Follow the tasks in the order presented in this document. Some steps can be performed in a different order, but the order used here limits repeated manipulation of some components.

You can set up a mixed-mode environment that includes an IPv6 vCenter Server. The Auto Deploy server must have an IPv4 address because PXE booting is supported only with IPv4.

Proof of Concept Preinstallation Checklist

Before you can start the proof of concept setup, make sure that your environment meets the hardware and software requirements and that you have the necessary permissions for the components that are included in the setup.

You need the following hardware and software for your proof of concept setup.

- vCenter Server 5.1 or later is installed on a Windows system. In this proof of concept setup, you install the Auto Deploy server and the vSphere PowerCLI on the host on which the vCenter Server is running. You perform many of the setup tasks by logging in to that host, either directly into the console or by using Remote Desktop (RDP).
 - Datacenter, clusters, and folders configured on the vCenter Server system.
 - At least 4GB of free space on the vCenter Server system. 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 see the LUNs.
 - List of target IP addresses for NFS or iSCSI.
 - List of target volume information for NFS or iSCSI.
- Two or more hosts to be provisioned with Auto Deploy, and the following information for each host.
 - List of MAC addresses for each physical NIC.
 - List of IP addresses and fully qualified host names preassigned for target ESXi installs.
 - Default route, netmask, and primary and secondary DNS server IP addresses.
 - IP address and netmask for the VMkernel primary (management) network.

- IP address and netmask for other VMkernel networks such as storage, vSphere FT, or VMware vMotion.

Auto Deploy does not overwrite existing partitions by default.

- vSphere installer (DVD or ISO).
- Window 7 or Windows Server 2008 system with Microsoft PowerShell preinstalled.
- vSphere PowerCLI installer binaries downloaded from the Downloads page on the VMware Web site.
- Location of the ESXi software depot on the Downloads page of the VMware Web site. You will use a URL to point to the image profile stored at that location or download a ZIP file to work with a local depot. Do not download the ESXi image.
- TFTP installer software such as WinAgents TFTP server. The TFTP server included in Windows 2008 is closely tied to Windows network deployment and is not suitable.
- DHCP server. The DHCP server included with Windows 2008 is suitable for this proof of concept setup.

You also need information about and administrator privileges to the environment's core servers 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

Auto Deploy relies on a TFTP server for sending the boot image to the hosts that it provisions. You must install a TFTP server in your environment.

This task only installs the TFTP server. You later download a configuration file to the server. See [“Configure the Auto Deploy and TFTP Environment in the vSphere Web Client,”](#) on page 216.

Prerequisites

Make sure your system meets the requirements in the Preinstallation Checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.

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 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 and set up vSphere PowerCLI. You use PowerCLI cmdlets to write the rules that assign image profiles and host profiles to hosts. See [“Install and Set Up vSphere PowerCLI,”](#) on page 211.

Install and Set Up vSphere PowerCLI

You manage Auto Deploy with rules that you create with vSphere PowerCLI cmdlets.

This proof of concept setup installs vSphere 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 2.0 is installed, or install it from the Microsoft Web site following the instructions on that Web site.
- Verify that Microsoft Powershell 2.0 is installed, or install it from the Microsoft website following the instructions on that Web site.

PowerShell 2.0 is preinstalled on Windows 2008 and Windows 7 systems.

Procedure

- 1 Log in to the console of the Windows system on which vCenter Server is installed with administrator privileges, either directly by using RDP.
- 2 Open a command prompt and type the following commands in sequence, pressing Enter after each line.

```
powershell
Set-ExecutionPolicy RemoteSigned
Exit
```

- 3 Download vSphere PowerCLI from the Download page of the VMware Web site and install the vSphere PowerCLI software.
- 4 Confirm that PowerCLI is working.
 - a Double-click the PowerCLI icon on the desktop to open a PowerCLI window.
 - b Ignore the SSL error, type **Get-DeployCommand**, and press Enter.

PowerCLI displays a list of cmdlets and their definitions in the PowerCLI window.

What to do next

- If you do not see a list of cmdlets when you run **Get-DeployCommand**, check your PowerCLI version and uninstall and reinstall if necessary.
- For some background information on PowerCLI, see [“Using Auto Deploy Cmdlets,”](#) on page 168. See the vSphere PowerCLI documentation set for details.
- Prepare the hosts you want to provision with Auto Deploy. See [“Prepare Auto Deploy Target Hosts,”](#) on page 212.

Prepare Auto Deploy Target Hosts

You must prepare all target hosts for Auto Deploy.

Prerequisites

Hosts that you want to provision with Auto Deploy must meet the requirements for ESXi.

See [“ESXi Hardware Requirements,”](#) on page 13.

NOTE You cannot provision EFI hosts with Auto Deploy unless you switch the EFI system to BIOS compatibility mode.

Procedure

- 1 Change each physical host's BIOS settings to force the host to boot from the primary network device.
- 2 Reconfirm the MAC address of the primary network device.

What to do next

Prepare the DHCP Server. See [“Prepare the DHCP Server,”](#) on page 213.

Prepare the DHCP Server

The DHCP Server in your proof of concept environment must be set up to serve each target host with an iPXE binary.

The proof of concept environment uses Active Directory with DNS and DHCP.

The proof of concept illustrates how to use DHCP reservations. Setting up fixed IP addresses for each host is time consuming and not recommended.

Prerequisites

- Make sure your system meets the requirements in the preinstallation checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.
- Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.

Procedure

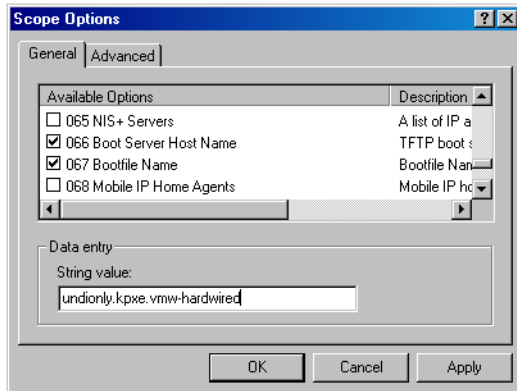
- 1 Log in to your DHCP Server as an administrator user.
- 2 Create a DHCP scope for your IP address range.
 - a Click **Start > Settings > Control Panel > Administrative Tools** and click **DHCP**.
 - b Drill down to **DHCP > *hostname* > IPv4**, right click, and click **New Scope**.
 - c Click **Next** on the Welcome screen and specify a name and description for the scope.
 - d Specify an IP address range and click **Next**.
 - e Click **Next** until you reach the Configure DHCP Options screen and select **No, I will configure this option later**.
- 3 If you are planning on using DHCP reservations, create a DHCP reservation for each target ESXi host.
 - a In the DHCP window, drill down 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.

The precise process depends on the DHCP Server you are using. This example uses the DHCP server included with Windows 2008.

- a In the DHCP window, drill down 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

Prepare the DNS Server. See [“Prepare the DNS Server,”](#) on page 214.

Prepare the DNS Server

Preparing the DNS server involves adding the DHCP information to the DNS server and verifying that the DNS entries are working. This task is optional.

The example environment uses Active Directory with DNS and DHCP.

Prerequisites

Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.

Procedure

- 1 Log in to the DNS server.
- 2 Add the DHCP reservation IP addresses and the associated host names as static DNS entries.
Be sure to add entries in both Forward (ARecord) and Reverse (PTR Record) Zones.

- 3 Log in to the console of the Windows system on which vCenter Server is installed with administrator privileges, either directly or by using RDP.
- 4 Open a command prompt and perform an nslookup of ESXi host names to validate that the DNS entries are working.
Use both forward (Short and FQDN) and reverse lookups.
- 5 Log out of your DNS server.

What to do next

Install the Auto Deploy Server software. See [“Install Auto Deploy Server Software,”](#) on page 215.

Install Auto Deploy Server Software

Auto Deploy server software is included with the vCenter Server installation media. This proof of concept setup installs the Auto Deploy server on the system on which vCenter Server is installed.

Auto Deploy is supported only if you are using the same version of Auto Deploy and the vCenter Server system. For example, if you are using vCenter Server 5.5, install Auto Deploy server 5.5. If you have an Auto Deploy 5.0 server installed and you cannot upgrade the Auto Deploy server, see [“Problems if You Upgrade vCenter Server But Do Not Upgrade Auto Deploy Server,”](#) on page 209.

Prerequisites

- Make sure your system meets the requirements in the preinstallation checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.
- Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.

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 Secure the installation media for the Auto Deploy server software.

vCenter Server type	Action
Virtual	Download the vSphere installer CD ISO image and place the ISO on a data store the vCenter Server can access.
Physical	Download the vSphere installer CD ISO image and burn it to disk.

- 3 Make the ISO available to the vCenter Server.

vCenter Server type	Action
Virtual	Attach the CD-ROM drive to the ISO, either by using the ESXi or from the datastore.
Physical	Insert the DVD into the physical server's drive.

- 4 Run Autoplay on the DVD.
- 5 Select Auto Deploy Server and click **Install**.
- 6 When the installer wizard prompts you for the Repository Location, browse to the volume or hard drive that has enough storage for Auto Deploy use and select that location.
Network shares are not an option when you install Auto Deploy.
- 7 Leave the defaults for everything else.
- 8 When the installer prompts you for credentials, use your vCenter Server administrator credentials.

What to do next

- Request access to a vSphere Web Client that can access the vCenter Server system that will manage the Auto Deploy server.
- Configure the Auto Deploy and TFTP environment. See [“Configure the Auto Deploy and TFTP Environment in the vSphere Web Client,”](#) on page 216.

Configure the Auto Deploy and TFTP Environment in the vSphere Web Client

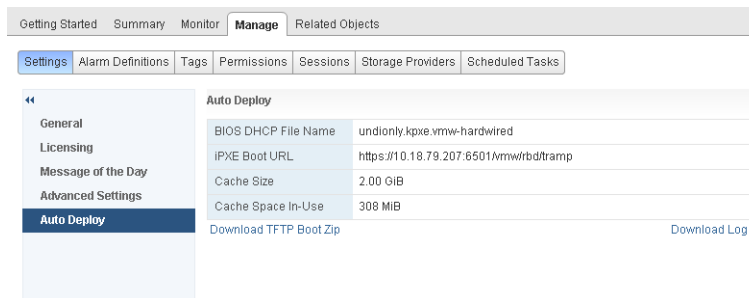
You must download a FTP Boot ZIP file from your Auto Deploy server. The customized FTP server serves the boot images that Auto Deploy provides. You can perform the task in the vSphere Web Client.

Prerequisites

- Make sure your system meets the requirements in the preinstallation checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.
- Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.

Procedure

- 1 From your Web browser, access the URL of the vSphere Web Client that connects to the vCenter Server system which manages the Auto Deploy Server.
- 2 When the Certificate warning appears continue to the vCenter Server system.
- 3 In the inventory, drill down to the vCenter Server system.
- 4 Click the Manage tab, select Settings, and click Auto Deploy.
- 5 Click **Download TFTP Boot Log** to download the TFTP configuration file.



- 6 Save the file (Deploy-tftp.zip) to the TFTP_Root directory that you created when you installed the TFTP Server and unzip the file into that directory.
- 7 Minimize the Web browser you are using with the vSphere Web Client.

What to do next

Prepare the depot from which Auto Deploy retrieves the ESXi software when it provisions the hosts. See [“Prepare the ESXi Software Depot,”](#) on page 217.

Prepare the ESXi Software Depot

Auto Deploy provisions hosts with images described by image profiles. Image profiles are stored in software depots. You must make sure the correct image profile is available before you start provisioning hosts.

The ESXi software depot contains the image profiles and software packages (VIBs) that are used to run ESXi. An image profile is a list of VIBs. This proof of concept setup uses a depot and image profile provided by VMware and does not create custom image profiles.

This proof of concept setup downloads the ZIP file that contains the image profile. You can instead point the Auto Deploy server to the HTTP URL of an image profile.

If you require custom VIBs such as custom drivers in your image profile, you can create a custom image profile by using the Image Builder PowerCLI.

The steps in this task instruct you to run PowerCLI cmdlets. For additional information on each cmdlet, type `Help cmdlet` at the PowerCLI prompt or search the vSphere Documentation Center.

Prerequisites

- Make sure your system meets the requirements in the preinstallation checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.
- Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.

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 the ESXi Depot ZIP file from the VMware Web site to a location the PowerCLI Windows system can access.

The file has a name that follows this pattern: `VMware-Esxi-N.N.N-xxxxx-depot.zip`.

- 3 Save the ZIP file to your local D:\ drive or any volume with enough space and note the location of the file.
- 4 Start a PowerCLI session and run the following cmdlets at the prompt.

```
Connect-VIServer -Server your_vc_hostname -User username -Password password <Enter>
Add-EsxSoftwareDepot path:\VMware-Esxi-version-xxxxx-depot.zip <Enter>
```

Include the complete path and file name of the ZIP file you downloaded.

- 5 Validate that you successfully added the ZIP file to the depot by checking the contents of the depot with the `Get-EsxImageProfile` cmdlet.

```
Get-EsxImageProfile <Enter>
```

The cmdlet returns information about all image profiles in the depot.

What to do next

Set up Auto Deploy to provision the first host and provision that host with the image profile in the depot. See [“Set Up the First Host and Provision with Auto Deploy,”](#) on page 218.

Set Up the First Host and Provision with Auto Deploy

Setting up the first host requires that you understand how to write Auto Deploy rules with vSphere PowerCLI. After you write the rules and add them to the ruleset, you can turn on the host to provision it.

You use the PowerCLI command-line interface to specify how Auto Deploy provisions the target hosts. You define rules and add each rule to the active ruleset. The Auto Deploy server checks the ruleset to determine which image profile to send to each ESXi host, which host profile to send to each ESXi host, and which location on the vCenter Server to place the host in.

A rule allows you to specify the following parameters.

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 to be used, a host profile to be used, or a vCenter Server inventory location (datacenter, folder, cluster) for the target host. You can specify multiple items separated by commas.																		
Pattern	The pattern specifies the host or group of hosts to which the rule applies. Choose one of the following. <table> <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>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> Specify <code>-AllHosts</code> to apply the item or items to all hosts.	vendor	Machine vendor name.	model	Machine model name.	serial	Machine serial number.	hostname	Machine hostname.	domain	Domain name.	ipv4	IPv4 address of the machine.	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.																		
mac	Boot NIC MAC address.																		
asset	Machine asset tag.																		
oemstring	OEM-specific strings in the SMBIOS.																		

This proof of concept setup first uses `-AllHosts` and later uses an IP address range to identify the hosts to provision.

Write Rules for the First Host

You specify the image profile to provision the host with by using PowerCLI to write a rule and adding the rule to the active ruleset.

This task assumes you have a basic knowledge of Microsoft PowerShell and vSphere PowerCLI.

Prerequisites

- Make sure your system meets the requirements in the preinstallation checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.
- Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.
- Make sure you can access the ESXi software from the system on which you run the PowerCLI cmdlets.

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.

This task assumes that you installed PowerCLI on the system on which the vCenter Server system is running.

- 2 Open the PowerCLI window and list the ESXi image profiles.

Get-EsxImageProfile

- 3 Create a new rule by running the following cmdlet, replacing ESXi-5.1.0-XXXXX-standard with the image profile that you want to use.

New-DeployRule -Name "InitialBootRule" -Item "Esxi-5.1.0-XXXXX-standard" -AllHosts

- 4 Add the new rule to the active rule set to make the rule available to the Auto Deploy server.

Add-DeployRule -DeployRule "InitialBootRule"

What to do next

Boot the host and check that Auto Deploy provisions the host and adds it to the vCenter Server inventory. See [“Provision the First Host,”](#) on page 219.

Provision the First Host

You can provision the first host and check its location on the vCenter Server to complete verification of the image provisioning of your setup.

Prerequisites

- Make sure your system meets the requirements in the preinstallation checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.
- Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.

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 a name you specified earlier in the DNS server. The host contacts the Auto Deploy server and downloads the ESXi binaries from the HTTP URL indicated in the iPXE tramp file that you downloaded into the TFTP_Root directory earlier. Each instance of Auto Deploy produces a custom set of files for the TFTP Server.

- 2 With a vSphere Web Client and connect to the vCenter Server system.

In this proof of concept setup, the vCenter Server system is localhost.

- 3 Click **Hosts and Clusters**.
 - 4 Check that the newly provisioned host is now in the vCenter Server inventory at the datacenter level.
- By default, Auto Deploy adds hosts at the datacenter level when the boot process completes.

What to do next

If you encounter problems, see [“Troubleshooting Auto Deploy,”](#) on page 203.

Configure the first host for use as a reference host and save its host profile for use with other hosts. See [“Configure the Proof of Concept Reference Host,”](#) on page 220.

Configure the Proof of Concept Reference Host

You can customize the first ESXi host that you boot for your environment and create a host profile. You can set up Auto Deploy to provision other target hosts with that host profile. The ESXi host you create the host profile from is considered your reference host or template host.

How you configure the reference host depends on what you want to do.

Shared settings	Specify settings that all hosts share and save a host profile for the host.
Host-specific settings	Customize hosts by setting up the host profile to prompt for user input for a limited number of options such as a static IP address. Host customizations are saved when you save the host profile. See “Host Customization in the vSphere Web Client,” on page 195.

Auto Deploy applies all common settings from the host profile to all target hosts. If you set up the host profile to ask 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.

NOTE Administrators cannot directly access or manipulate host customizations. Use the vSphere Web Client Host Profiles UI to work with host customizations.

Prerequisites

- Make sure your system meets the requirements in the preinstallation checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.
- Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.

Procedure

- 1 Use the vSphere Web Client to connect to the vCenter Server system.
In this proof of concept setup, the vCenter Server system is localhost.
- 2 Click Hosts and Clusters and select the host that Auto Deploy added to the first datacenter.
- 3 Configure the host.

The rest of the proof of concept setup assumes that you configure at least one setting that is different for different hosts.

Configuration	Description
Networking	Configure the following networking components. <ul style="list-style-type: none"> ■ Base virtual switch and management port group for VMkernel. ■ Storage network port group for VMkernel. ■ Virtual machine networking port group. ■ Any additional virtual switches and port groups. ■ Distributed switches, if necessary (transfer port groups to distributed switches if you use them).
Storage	Configure shared storage.
Time settings	Configure your time settings.
Security	Configure the security profile.
Authentication	Configure authentication.
DNS and routing	If necessary, configure DNS and route settings.
Other	Configure advanced settings or any other settings as required in the target environment.

What to do next

Create the host profile from the reference host for use with all other target hosts. See the *Host Profiles* documentation.

Create and Apply a Host Profile with the vSphere Web Client

Configuration that is shared by a group of hosts is stored in a host profile. You can create the host profile from your reference host. Configuration that differs for different hosts, such as a static IP address, can be managed through the host customization mechanism.

Auto Deploy can provision each host with the same host profile. In certain cases, Auto Deploy also uses host customizations that allow you to specify information that differs 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.

In this example, you extract a host profile from a reference host, attach the host profile to one other host, and check host profile compliance. In most cases, you do not perform these tasks manually but you write an Auto Deploy rule that applies a host profile to hosts that are provisioned with Auto Deploy. See [“Assign a Host Profile to Hosts,”](#) on page 172.

Prerequisites

- Make sure your system meets the requirements in the preinstallation checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.
- Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.

Procedure

- 1 Log in to a vSphere Web Client that is connected to the vCenter Server system with administrator privileges.
- 2 Click **Rules and Profiles** and select **Host Profiles**.
- 3 Click the Extract profile from host icon and respond to the wizard prompts.

Option	Description
Select Host	Select the reference host you configured earlier.
Name and Description	Name the profile ESXiGold and add a description.
Ready to Complete	Review the information and click Finish .

- 4 Right-click the ESXiGold host profile, select **Attach/Detach Hosts and Clusters**.
- 5 Select the ESXi host to which you want to attach the profile, click **Attach**, and click **Next**.
The wizard loads the host customization.
- 6 Provide any customization information and click **Finish**.

What to do next

Create a rule that assigns the image profile and the newly-created host profile to all hosts you want to provision with Auto Deploy. See [“Create a Rule for Other Target Hosts,”](#) on page 221.

Create a Rule for Other Target Hosts

You can create a rule that applies the previously verified image profile and the host profile that you just created to all target hosts.

This task assumes you have a basic knowledge of Microsoft PowerShell and vSphere PowerCLI.

Prerequisites

- Make sure your system meets the requirements in the preinstallation checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.
- Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.

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 Start a PowerCLI session and type the following commands, followed by Enter, at the prompt.

```
Connect-VIServer -Server your_vc_hostname -User username -Password password  
Add-EsxSoftwareDepot path:\VMware-Esxi-version-xxxxx-depot.zip
```

Include the complete path and file name of the ZIP file you downloaded earlier. Adding the software depot is required each time you start a new PowerCLI session.

- 3 (Optional) To display the rules in the active ruleset type the following cmdlet at the prompt and press Enter.

```
Get-DeployRuleset
```

- 4 To create a rule that instructs Auto Deploy to provision the set of hosts in the specified IP range with the image you selected and with the host profile you created from the reference host, type the following command and press Enter.

```
New-DeployRule -name "Production01Rule" -item "image_profile", ESXiGold,target_cluster -  
Pattern "ipv4=IP_range"
```

Option	Description
image_profile	The ESXi image profile you used in the first deploy rule.
target_cluster	Name of the cluster in vCenter Server to which you want to add all hosts.
IP_range	Either a single IP address or a range of IP addresses for the hosts you want to provision with the image profile and host profile.

When you specify a target cluster, the host profile is applied to all hosts in the cluster. Applying the host profile to each host is not required.

- 5 Add the new rule to the active ruleset.
- 6 (Optional) Remove the deploy rule you created for the initial boot operation.
- 7 Check the active rule set.

```
Get-DeployRuleset<Enter>
```

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 all hosts and set up host customizations for each host. See [“Provision All Hosts and Set Up Host Customizations,”](#) on page 223.

Provision All Hosts and Set Up Host Customizations

With the rule in place that provisions hosts using an image profile, and with the host profile created from the reference host available, you can provision all 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.

Prerequisites

- Make sure your system meets the requirements in the preinstallation checklist. See [“Proof of Concept Preinstallation Checklist,”](#) on page 210.
- Perform all preceding proof of concept setup tasks. See [“Auto Deploy Proof of Concept Setup,”](#) on page 210 for the complete list.
- Open a console to each host you want to provision to monitor boot progress.

Procedure

- 1 Boot the remaining hosts.

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 With a vSphere Web Client, connect to the vCenter Server system.
- 3 Click **Home** and select **Host Profiles**.
- 4 In the panel on the left, select the ESXiGold profile and add the newly booted hosts to that profile.
- 5 Apply the host profile to each of the hosts, provide the user input information, and 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 retrieve that information and boot completely.

What to do next

With your proof of concept implementation completed successfully, you can start planning your production setup.

Using vSphere ESXi Image Builder CLI

The ESXi Image Builder CLI is a set of PowerCLI cmdlets that you can use to manage vSphere image profiles and VIB packages, such as driver VIBs and update VIBs. You can also use Image Builder cmdlets to export an image profile to an ISO or offline depot ZIP file that you can use to install ESXi with a customized set of updates, patches, and drivers.

Understanding Image Builder

You can use the 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.

Image Builder Overview

The Image Builder PowerCLI supports management of vSphere image profiles and VIBs.

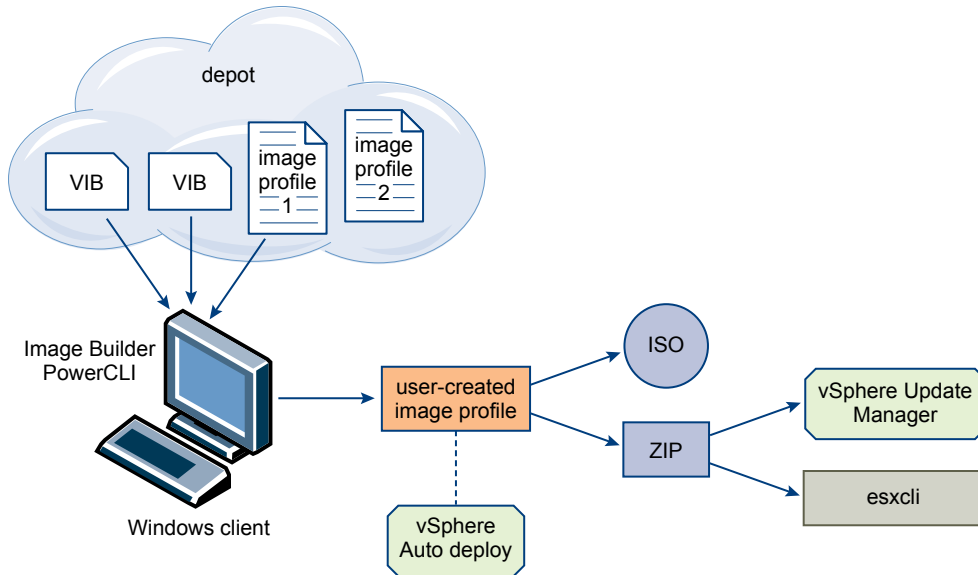
You can manage VIBs and image profiles with vSphere Image Builder. VIBs are software packages, and image profiles specify a set of software packages. See [“Software Depots and Their Components,”](#) on page 225.

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

Create image profiles for use by Auto Deploy.	Use Image Builder to create an image profile that defines the VIBs that Auto Deploy uses to provision hosts.
Add custom third-party drivers to existing image profile and export to ISO or bundle.	To add third-party driver or extension custom VIBs to your ESXi hosts, use 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 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.	Some customers require a minimal footprint image. These customers can clone the ESXi base image profile and remove VIBs using Image Builder.

The Image Builder PowerCLI cmdlets take image profiles and VIBs as input and produce different outputs.

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 ISO or an offline depot ZIP file. ISO images are used for installation. The ZIP depot can be used by vSphere Update Manager or by <code>esxcli software</code> commands to update or install images. Image profiles are also used in vSphere Auto Deploy rules to specify the software to provision ESXi hosts with.

Figure 7-6. Image Builder Architecture

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,”](#) on page 229.

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 using the Image Builder PowerCLI.

See [“ImageProfile Object Properties,”](#) on page 228.

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.

Image Builder PowerCLI Overview

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

The Image Builder PowerCLI includes the following cmdlets.

NOTE When you run 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 7-15. 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 “Create an Image Profile,” on page 233.
<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,” on page 227.
<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 Auto Deploy and to other ESXi 5.x hosts. You define and manipulate image profiles with the Image Builder PowerCLI.

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 Image Builder PowerCLI 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, Image Builder checks that the change does not invalidate the profile.

Dependency Validation

When you add or remove a VIB, 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,”](#) on page 228

Acceptance Level Validation

Image Builder performs acceptance level validation each time an image profile is created or changed. 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, 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 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 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 IOVP 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 plugins 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 Community Supported 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

In some situations, 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,” on page 227.
Author	System.String	Person who created the profile. 60 characters or fewer.
CreationTime	System.DateTime	Timestamp of creation time.
Description	System.String	Full text description of profile. No length limit.
GUID	System.String	Globally unique ID of the image profile.
ModifiedTime	System.DateTime	Timestamp of last modification time.
Name	System.String	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[]	Displays any OEM hardware requirements that the image profile might have. Auto Deploy checks this property when deploying an image profile and deploys the profile if matching hardware is available.
Vendor	System.String	Organization that publishes the profile. 40 characters or fewer.
VibList	SoftwarePackage[]	List of VIB IDs the image consists of.

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	Acceptance level of this VIB.
Conflicts	SoftwareConstraint[]	List of VIBs that cannot be installed at the same time as this VIB. Each constraint uses the following format: <code>package-name[<< <= >= >>version]</code>
Depends	SoftwareConstraint[]	List of VIBs that must be installed at the same time as this VIB. Same constraint format as <code>Conflicts</code> property.
Description	System.String	Long description of the VIB.
Guid	System.String	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	Name of the VIB. Usually uniquely describes the package on a running ESXi system.
Provides	SoftwareProvides	List of virtual packages or interfaces this VIB provides. See “SoftwareProvide Object Properties,” on page 231.
ReferenceURLs	SupportReference[]	List of <code>SupportReference</code> objects with in-depth support information. The <code>SupportReference</code> object has two properties, <code>Title</code> and <code>URL</code> , both of type <code>System.String</code> .
Replaces	SoftwareConstraint[]	List of <code>SoftwareConstraint</code> 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[]	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	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.

Name	Type	Description
Vendor	System.String	VIB vendor or publisher.
Version	System.String	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,” on page 230

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

SoftwareVersion Object Properties

The `SoftwareVersion` object allows you to 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 higher than the second version string. `Compare` returns 0 if two versions strings are equal. `Compare` returns -1 if the second version string is higher than the first string. The object has the following properties.

Name	Type	Description
Version	System.String	The part of the version before the hyphen. This part indicates the primary version.
Release	System.String	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 False otherwise.

The SoftwareConstraint object also includes the following properties.

Name	Type	Description
Name	System.String	Name of the constraint. This name should match a corresponding SoftwareProvides Name 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 Relation and Version property.
Version	System.String	The version to match the constraint against. This property can be \$null if the constraint does not have a Relation and Version property.
VersionObject	SoftwareVersion	The version represented by a SoftwareVersion object.

SoftwareProvide Object Properties

The SoftwareProvide object includes the following properties.

Name	Type	Description
Name	System.String	Name of the provide
Version	System.String	Version of the provide. Can be \$null if the provide does not specify a version.
Release	System.String	Version of the provide as represented by a SoftwareVersion object. See “SoftwareVersion Object Properties,” on page 230.

Image Builder Installation and Usage

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

Install Image Builder PowerCLI and Prerequisite Software

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

You install Image Builder and prerequisite software on a Microsoft Windows system.

Procedure

- 1 Install Microsoft .NET 2.0 from the Microsoft website following the instructions on that website.
- 2 Install Microsoft PowerShell 2.0. from the Microsoft website following the instructions on that website.
- 3 Install vSphere PowerCLI, which includes the Image Builder cmdlets.

See the *vSphere PowerCLI Installation Guide* for detailed instructions.

What to do next

Review [“Using Image Builder Cmdlets,”](#) on page 232. If you are new to PowerCLI, read the PowerCLI documentation.

Use 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.

Using Image Builder Cmdlets

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

Experienced PowerShell users can use Image Builder cmdlets just like other PowerShell cmdlets. If you are new to PowerShell and PowerCLI, the following tips help you come up to speed.

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`.

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-ExsSoftwarePackage -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-ExsImageProfile
```

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

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

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

Most examples in the documentation pass in parameters by name. [“Image Builder Workflows,”](#) on page 240 includes examples that pass parameters as objects.

Setting Properties to Support Remote Signing

For security reasons, Windows PowerShell supports an execution policy feature. It determines whether scripts are allowed to run and whether they must be digitally signed. By default, the execution policy is set to Restricted, which is the most secure policy. If you want to run scripts or load configuration files, you can change the execution policy by using the Set-ExecutionPolicy cmdlet. To do this, type the following in the vSphere PowerCLI console window.

```
Set-ExecutionPolicy RemoteSigned
```

If the command is successful, you can run scripts and load configuration files. For more information about the execution policy and digital signing in Windows PowerShell, use the following cmdlet.

```
Get-Help About_Signing
```

Image Builder Common Tasks

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

Create 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 profile from scratch.

Administrators performing this task must have some experience with PowerCLI or Microsoft PowerShell.

Prerequisites

- Install the VMware PowerCLI and all prerequisite software. See [“Image Builder Installation and Usage,”](#) on page 231.
- Verify that you have access to the software depot that contains the image profile you want to clone.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Image Builder Cmdlets,”](#) on page 232.

Procedure

- 1 At the PowerShell prompt, add the depot that contains the profile you want to clone to the current session.

Depot Type	Cmdlet
Remote depot	Run Add-EsxSoftwareDepot -DepotUrl <i>depot_url</i> .
ZIP file	a Download the ZIP file to a local file path. b Run Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip

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.

Get-EsxImageProfile

You can use filtering options with Get-EsxImageProfile.

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

```
New-EsxImageProfile -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 [“Examining Depot Contents,”](#) on page 240 for some examples of filtering.

Customize the image profile by adding or removing VIBs. See [“Add VIBs to an Image Profile,”](#) on page 234.

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 `ReadOnly`. 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, the Image Builder PowerCLI performs validation. If you add VIBs from two or more OEM partners, no errors are reported but the resulting image profile might not work. Install VIBs from only one OEM vendor at a time.

You might be able to add VIBs even if the resulting image profile is invalid.

NOTE VMware can support only environments and configurations that have been proven to be stable and fully functional through rigorous and extensive testing. Use only those supported configurations. If you have to go outside of that clearly defined space, use custom VIBs. Custom VIBs lower your host compliance and, as a result, supportability. But you know the changes you made, and can maintain those changes with the custom VIB mechanism.

Prerequisites

- Install the VMware PowerCLI and all prerequisite software. See [“Install Image Builder PowerCLI and Prerequisite Software,”](#) on page 231.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Image Builder Cmdlets,”](#) on page 232.
- Administrators performing this task must have some experience with PowerCLI or Microsoft PowerShell.

Procedure

- 1 Run `Add-EsxSoftwareDepot` for each depot you want to work with.

Depot Type	Cmdlet
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>

The cmdlet returns one or more `SoftwareDepot` objects.

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

Get-EsxImageProfile

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 clone the image profile. The vendor parameter is required.

- 4 Run `Add-ESXSoftwarePackage` to add a new package to one of the image profile.

```
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 results. Change the acceptance level of the image profile to add the VIB.

- 5 (Optional) Change the acceptance level of the image profile if an error about acceptance level problems displays.

VIB acceptance levels are set during VIB creation and cannot be changed.

The image profile includes the new VIB.

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.

Administrators performing this task must have some experience with PowerCLI or Microsoft PowerShell.

Prerequisites

- Install the VMware PowerCLI and all prerequisite software. See [“Install Image Builder PowerCLI and Prerequisite Software,”](#) on page 231.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Image Builder Cmdlets,”](#) on page 232.

Procedure

- 1 Run `Add-ESXSoftwareDepot` to connect to the depot that contains the image profile to export.

Depot Type	Cmdlet
Remote depot	Run <code>Add-ESXSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	a Download the ZIP file to a local file path. 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.

To build...	Run...
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, Image Builder validates VIB signatures, adds VIB binaries to the image, and downloads the image to the specified location. For the ZIP file, Image Builder validates VIB signatures and downloads the VIB binaries to the specified location.

Example: Exporting an Image Profile to ISO or Offline Bundle

Export an image profile to ISO from the PowerCLI prompt.

- 1 Add the software depot.

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

- 2 Display 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
```

Export an image profile to an offline bundle ZIP from the PowerCLI prompt.

- 1 Add the software depot.

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

- 2 Display 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 vSphere Update 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 VMware PowerCLI and all prerequisite software. See [“Install Image Builder PowerCLI and Prerequisite Software,”](#) on page 231.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Image Builder Cmdlets,”](#) on page 232.

Procedure

- 1 In a PowerCLI session, create an image profile, for example by cloning an existing image profile and adding a VIB.

- 2 Before you exit the session, 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, you can add the depot to access your image profile.

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

Working with Acceptance Levels

Hosts, image profiles, and individual VIBs have acceptance levels. VIB acceptance levels show tested the VIB. 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, VMware refers you to the 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,”](#) on page 239.

The vSphere Update Manager does not display the actual acceptance level. Use Image Builder PowerCLI 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 results. Change the acceptance level of the host to install the image profile or VIB. See [“Change the Host Acceptance Level,”](#) on page 238. Changing the acceptance level of the host changes the system support for that host.

The acceptance level of a host, image profile, or VIB allows you to 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 IOVP 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 plugins 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 Community Supported 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 host acceptance level. 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.

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

- 2 Get the host acceptance level.

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

- 3 If the acceptance level of the VIB is lower than the acceptance level of the host, 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 your setup is no longer consistent. The warning is repeated when you install VIBs, remove VIBs, and 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 VIB acceptance level is lower than the image profile acceptance level, you can clone the image profile with a lower acceptance level or change the image profile acceptance level.

Prerequisites

- Install the VMware PowerCLI and all prerequisite software. See [“Install Image Builder PowerCLI and Prerequisite Software,”](#) on page 231.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Image Builder Cmdlets,”](#) on page 232.

Procedure

- 1 Run `Add-EsxSoftwareDepot` for each depot you want to work with.

Depot Type	Cmdlet
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>.

- 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
```

You can specify `VMwareCertified`, `VMwareAccepted`, `PartnerSupported`, or `CommunitySupported` as an acceptance level. If you lower the acceptance level, the level of support for the image profile and hosts that you provision with it changes. See [“Acceptance Levels,”](#) on page 227.

Image Builder Workflows

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.

Examining Depot Contents

You can examine software depots and VIBs with Image Builder PowerCLI cmdlets. This workflow illustrates examining depot contents and includes examples of wildcard usage.

The workflow itself passes parameters by name, the example below passes parameters as objects by accessing variables.

Before you use the cmdlets in this workflow, make sure your environment meets the following requirements.

- VMware PowerCLI and prerequisite is software installed. See [“Install Image Builder PowerCLI and Prerequisite Software,”](#) on page 231.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Image Builder Cmdlets,”](#) on page 232.

Examining depot contents is facilitated by the use of filtering options and wildcard characters.

- 1 At the PowerShell prompt, add the depot that you want to examine to the current session.

For remote depots, run **Add-EsxSoftwareDepot -DepotUrl *depot_url***.

For an offline depot ZIP file, you must download the ZIP file first.

- a Download the ZIP file to a local file path, but do not unzip it.
- b Run **Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip**

- 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 create by a vendor with a name that starts with a C.

- 3 Retrieve software packages by using Get-EsxSoftwarePackage.

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 a V.
- **Get-EsxSoftwarePackage -Vendor "V" -Name "*scsi"**
Returns all software packages with a name that has the string scsi in it from a vendor with a name that starts with a V.
- **Get-EsxSoftwarePackage -Version "2.0"**
Returns all software packages with a version string that starts with a 2.0.

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

- **Get-EsxSoftwarePackage -Vendor "V*" -Newest**

Returns the newest package for the vendors starting with V and displays the information as a table.

- **Get-EsxSoftwarePackage -Vendor "V*" -Newest | format-list**

Pipes the output of the request for software packages to the PowerShell `format-list` cmdlet and displays detailed information about each software package.

- 5 Display the list of VIBs in the image profile.

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

`VibList` is a property of the `ImageProfile` object. See [“Structure of ImageProfile, SoftwarePackage, and ImageProfileDiff Objects,”](#) on page 228.

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

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

Example: Depot Content Examination Using Variables

This example workflow 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-EsxSoftwarePackage -Vendor "V*"
Get-EsxSoftwarePackage -Vendor "V*" -Name "r*"
Get-EsxSoftwarePackage -Version "2.0*"
$ip1 = Get-EsxImageProfile -name ESX-5.0.0-123456-full
$ip1.VibList
Get-EsxSoftwarePackage -CreatedAfter 7/1/2010
```

Creating Image Profiles by Cloning Workflow

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

Before you use the cmdlets in this workflow, make sure your environment meets the following requirements.

- VMware PowerCLI and prerequisite is software installed. See [“Install Image Builder PowerCLI and Prerequisite Software,”](#) on page 231.
- If you encounter problems running PowerCLI cmdlets, consider changing the execution policy. See [“Using Image Builder Cmdlets,”](#) on page 232.

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.

- 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 necessary, add the depot that includes the profile you want to clone to the current session.

For remote depots, run **Add-EsxSoftwareDepot -DepotUrl *depot_url***.

For an offline depot ZIP file, you must download the ZIP file first.

- a Download the ZIP file to a local file path.
- b Run **Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip**

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

- 3 Check the `$DefaultSoftwareDepots` variable, which now returns the newly-added depot.

\$DefaultSoftwareDepots

- 4 Display all currently available image profiles.

Get-EsxImageProfile

The list helps you pick a likely candidate for cloning.

- 5 Clone one of the image profiles by specifying the name displayed in the Name column, a name for the new profile, and the vendor.

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

- 6 Display the newly created image profile.

\$ip

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 sample cmdlet sequence repeats the steps of this 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.

```
$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
```

Creating Image Profiles from Scratch Workflow

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

Before you use the cmdlets in this workflow, make sure your environment meets the following requirements.

- VMware PowerCLI and prerequisite is software installed. See [“Install Image Builder PowerCLI and Prerequisite Software,”](#) on page 231.
- 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 unzip to your local environment and access by using a file path.

See [“Create an Image Profile,”](#) on page 233 for an example of cloning and modifying an image profile.

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. Pass in the `-AcceptanceLevel` parameter to change the acceptance level of the image profile if you have to add a VIB with a lower 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.

The following workflow illustrates creating image profiles from scratch.

- 1 At the PowerShell prompt, add the depot that contains the packages you want to use to the current session.

For remote depots, run **Add-EsxSoftwareDepot -DepotUrl *depot_url***.

For an offline depot ZIP file, you must download the ZIP file first.

- a Download the ZIP file to a local file path.
- b Run **Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip**

- 2 List the available packages that you are interested in and bind them to a variable.

Get-EsxSoftwarePackage -CreatedAfter 7/1/2010

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

New-EsxImageProfile -NewProfile -Name "Test #2" -vendor "Vendor42" -SoftwarePackage esx-base[0]

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

- 4 Pipe the information about the new image profile to `format-list` for detailed information about the new package.

(Get-EsxImageProfile -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 from the PowerCLI prompt.

```
Add-EsxSoftwareDepot depoturl
$pkgs = Get-EsxSoftwarePackage -CreatedAfter 7/1/2010
$ip2 = New-EsxImageProfile -Name "Test #2" -vendor "Vendor42" -SoftwarePackage $pkgs[0]
$ip2 | format-list
```

Editing Image Profiles Workflow

You can create a custom image by cloning and editing an image profile. You can add or replace one or more VIBs in the existing profile. If adding or replacing VIBs would make the image profile inconsistent, an error results.

Before you use the cmdlets in this workflow, make sure your environment meets the following requirements.

- VMware PowerCLI and prerequisite is software installed. See [“Install Image Builder PowerCLI and Prerequisite Software,”](#) on page 231.
 - 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.
- 1 At the PowerShell prompt, add the depot that contains the image profile that you want to edit to the current session.

For remote depots, run **Add-EsxSoftwareDepot -DepotUrl *depot_url***.

For an offline depot ZIP file, you must download the ZIP file first.

- a Download the ZIP file to a local file path.
- b Run **Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip**
- 2 Pipe the image profile you intend to edit to **format-list** to see detailed information. In this example, the image profile created in [“Creating Image Profiles from Scratch Workflow,”](#) on page 242 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 formatted 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 the image profile's level, 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

- 5 Display 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 with 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
```

Setting Up ESXi

These topics provide information about using the direct console user interface and configuring defaults for ESXi.

This chapter includes the following topics:

- [“ESXi Autoconfiguration,”](#) on page 246
- [“About the Direct Console ESXi Interface,”](#) on page 246
- [“Set the Password for the Administrator Account,”](#) on page 249
- [“Configuring the BIOS Boot Settings,”](#) on page 249
- [“Host Fails to Boot After You Install ESXi in UEFI Mode,”](#) on page 250
- [“Network Access to Your ESXi Host,”](#) on page 251
- [“Configure the Network Settings on a Host That Is Not Attached to the Network,”](#) on page 251
- [“Managing ESXi Remotely,”](#) on page 252
- [“Configuring Network Settings,”](#) on page 252
- [“Storage Behavior,”](#) on page 256
- [“View System Logs,”](#) on page 259
- [“Configure Syslog on ESXi Hosts,”](#) on page 260
- [“Enable Lockdown Mode Using the Direct Console,”](#) on page 261
- [“Enable Lockdown Mode Using the vSphere Web Client,”](#) on page 261
- [“Enable ESXi Shell and SSH Access with the Direct Console User Interface,”](#) on page 262
- [“Set the Host Image Profile Acceptance Level,”](#) on page 262
- [“Reset the System Configuration,”](#) on page 263
- [“Remove All Custom Packages on ESXi,”](#) on page 264
- [“Disable Support for Non-ASCII Characters in Virtual Machine File and Directory Names,”](#) on page 264
- [“Disable ESXi,”](#) on page 264

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 through vCenter Server.

Table 8-1. 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
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](#) on page 247
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](#) on page 248
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](#) on page 248
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.

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,”](#) on page 248

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.

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.

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.

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 8-2. 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. 	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. You can override the link local IP address by configuring a static IP address using the direct console.
The ESXi host is connected to a functioning DHCP server, but you do not want to use the DHCP-configured IP address.	During the autoconfiguration phase, the software assigns a DHCP-configured IP address. You can make the initial connection by using the DHCP-configured IP address. Then you can configure a static IP address. 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.
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,” on page 251.

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 to configure the password for the administrator account (root).
- 3 Use the direct console 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.

Managing ESXi Remotely

You can use the vSphere Web Client and vCenter Server to manage the host.

For instructions about downloading and installing vCenter Server and the vSphere Web Client, see the following topics.

- [“Download the vCenter Server Installer,”](#) on page 73
- [Install vCenter Server as Part of a Custom Install](#)
- [“Install or Upgrade the vSphere Web Client,”](#) on page 101

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 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 could happen if the autoconfiguration phase did not succeed in configuring DHCP.
- The wrong network adapter was selected during the autoconfiguration phase.

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,”](#) on page 134

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 **Manage** tab, select **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.

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 through 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`, `slpd`, `wsman`, and `vobd`. The software also restarts Fault Domain Manager (FDM) if it is installed.

Users accessing this host through 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.

Disable the Management Network

The management network synchronizes VMware components and lets you access the ESXi host through the vSphere Web Client and vCenter Server. It is installed with the vSphere software. You might need to disable the management network to isolate a host from the vCenter Server inventory.

Users who access this host through the vSphere Web Client and vCenter Server lose connectivity when you disable the management network.

One reason to disable the management network is to isolate an ESXi host from an HA and DRS cluster, without losing your static IP and DNS configurations or rebooting the host.

This operation does not require downtime for virtual machines. The virtual machines continue to run while the host is disconnected from vCenter Server and the vSphere Web Client.

Procedure

- 1 From the direct console, select **Disable Management Network** and press Enter.

- 2 Press F11 to confirm.

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.

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.

Storage Behavior

When you start ESXi, the host enters an autoconfiguration phase during which system storage devices are configured with defaults.

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 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 8-3](#).

Table 8-3. 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,” on page 258. The VFAT scratch partition is created only on the disk from which the ESXi host is booting.</p> <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> <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 4GB 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 5.2GB 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. In some cases, 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 field **ScratchConfig.ConfiguredScratchLocation**, enter a directory path that is unique for this host.
- 7 Reboot the host for the changes to take effect.

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 field **ScratchConfig.CurrentScratchLocation** shows the current location of the scratch partition.

- 8 In the field **ScratchConfig.ConfiguredScratchLocation**, enter a directory path that is unique for this host.
- 9 Reboot the host for the changes to take effect.

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,”](#) on page 260.

Configure Syslog on ESXi Hosts

All ESXi hosts run a syslog service (`vm syslogd`), which logs messages from the VMkernel and other system components to log files.

You can use the vSphere Web Client or the `esxcli system syslog` vCLI command to configure the syslog service.

For more information about using vCLI commands, see *Getting Started with vSphere Command-Line Interfaces*.

Procedure

- 1 In the vSphere Web Client inventory, select the host.
- 2 Click the **Manage** tab.
- 3 In the System panel, click **Advanced System Settings**.
- 4 Locate the **Syslog** section of the Advanced System Settings list.
- 5 To set up logging globally, select the setting to change and click the Edit icon.

Option	Description
Syslog.global.defaultRotate	Sets the maximum number of archives to keep. You can set this number globally and for individual subloggers.
Syslog.global.defaultSize	Sets the 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. The directory should be specified 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://hostname1:514</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 you that want to customize.
 - b Click the Edit icon and enter the number of rotations and log size you want.
- 7 Click **OK**.

Changes to the syslog options take effect immediately.

Enable Lockdown Mode Using the Direct Console

To increase the security of your ESXi hosts, you can put them in lockdown mode.

When you enable lockdown mode, no users other than `vpxuser` have authentication permissions, nor can they perform operations against the host directly. Lockdown mode forces all operations to be performed through vCenter Server.

When a host is in lockdown mode, you cannot run vSphere CLI commands from an administration server, from a script, or from vMA against the host. External software or management tools might not be able to retrieve or modify information from the ESXi host.

NOTE Users with the DCUI Access privilege are authorized to log in to the Direct Console User Interface (DCUI) when lockdown mode is enabled. When you disable lockdown mode using the DCUI, all users with the DCUI Access privilege are granted the Administrator role on the host. You grant the DCUI Access privilege in Advanced Settings.

Enabling or disabling lockdown mode affects which types of users are authorized to access host services, but it does not affect the availability of those services. In other words, if the ESXi Shell, SSH, or Direct Console User Interface (DCUI) services are enabled, they will continue to run whether or not the host is in lockdown mode.

You can enable lockdown mode using the Add Host wizard to add a host to vCenter Server, using the vSphere Web Client to manage a host, or using the direct console user interface.

NOTE If you enable or disable lockdown mode using the Direct Console User Interface (DCUI), permissions for users and groups on the host are discarded. To preserve these permissions, you must enable and disable lockdown mode using the vSphere Web Client connected to vCenter Server.

Lockdown mode is available only on ESXi hosts that you add to vCenter Server.

See the *vSphere Security* documentation for more information about lockdown mode.

Procedure

- 1 In the direct console, select **Configure Lockdown Mode** and press Enter.
- 2 Press the spacebar to select **Enable Lockdown Mode** and press Enter.
- 3 Press Enter.

The host is in lockdown mode.

Enable Lockdown Mode Using the vSphere Web Client

Enable lockdown mode to require that all configuration changes go through vCenter Server. You can also enable or disable lockdown mode through the Direct Console User Interface (DCUI).

Procedure

- 1 Browse to the host in the vSphere Web Client inventory.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under System, select **Security Profile**.
- 4 In the Lockdown Mode panel, click **Edit**.
- 5 Select **Enable Lockdown Mode**.
- 6 Click **OK**.

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.

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,”](#) on page 227.

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 8-4. 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

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.

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 will display the virtual machine names in the non-ASCII characters, but ESXi will convert 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.

Disable ESXi

If you do not want your server to be an ESXi host, you can deactivate the ESXi setup.

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.

After You Install and Set Up ESXi

After ESXi is installed and set up, you can manage the host through the vSphere Web Client and vCenter Server, license the host, and back up your ESXi configuration.

This chapter includes the following topics:

- [“Managing the ESXi Host with the vSphere Web Client,”](#) on page 265
- [“Licensing ESXi Hosts,”](#) on page 265

Managing the ESXi Host with the vSphere Web Client

The vSphere Web Client provides the simplest way to manage your ESXi host and operate its virtual machines.

You can use the vSphere Web Client to connect to and operate vCenter Server by Web browser. To install the vSphere Web Client, see [“Install or Upgrade the vSphere Web Client,”](#) on page 101.

Licensing ESXi Hosts

You can use the vSphere Web Client and vCenter Server to license an individual host.

For information about managing host licenses, see *vCenter Server and Host Management*.

You can set up bulk licensing 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,”](#) on page 169.

You can also operate ESXi without licensing for a 60-day evaluation period, during which you can access the full ESXi feature set. See [“About ESXi Evaluation and Licensed Modes,”](#) on page 122

About ESXi Evaluation and Licensed Modes

After you purchase vSphere licenses, VMware provides a serial number that you use to license ESXi hosts. You can use evaluation mode to explore the entire set of features that are available for ESXi hosts, including features that are not included in the license that you have.

For example, in evaluation mode, you can use vMotion, HA, DRS, and other features, even if you have not licensed those features.

The installable version of ESXi is always installed in evaluation mode. ESXi Embedded is preinstalled on an internal USB 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, even if you start in licensed mode rather than evaluation mode. Any time during the 60-day evaluation period, you can convert from licensed mode to evaluation mode. To take full advantage of the 60-day evaluation period, you should convert to evaluation mode as soon as possible after you first power on the host.

For information about managing licensing and setting an ESXi host to evaluation mode, see the *vCenter Server and Host Management* documentation.

Recording the ESXi License Key

All ESXi editions have license keys associated with them. VMware recommends that you write down the license key and tape it to the server, or put the license key in a secure, easily accessible location.

You can access the license key from the direct console or the vSphere Web Client. If the host becomes inaccessible or unbootable, it is important that you have a record of the 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.

Access the ESXi License Key from the vSphere Web Client

You can use the vSphere Web Client to access the ESXi license key.

Procedure

- 1 From the vSphere Web Client, connect to the vCenter Server.
- 2 Select the host in the inventory.
- 3 Select the **Manage** tab.
- 4 Select **Settings**.
- 5 Select **System**.
- 6 Select **Licensing**.

The license key appears in the form XXXXX-XXXXX-XXXXX-XXXXX-XXXXX.

Index

Symbols

%include command 143
%post command 143
%pre command 143

Numerics

3rd-party modules, removing 264
64-bit DSN requirement 31

A

acceptance levels
 image profiles 239
 VIBs 227
acceptance levels, host 238
accepteula command 143
access, restricting 261
Active Directory identity source 86
Active Directory LDAP Server identity source 87
active rule set 157
Add-DeployRule 221
administrative password 249
administrator user, setting for vCenter Server 51
advanced management (Auto Deploy) 192
Apply-EsxImageProfile cmdlet 177
authenticating to vCenter Server 57
Authentication Proxy, *See also* vSphere
 Authentication Proxy
Auto Deploy
 best practices 198
 boot file 165
 boot operation 155
 boot process 158
 cached 178
 caching 182
 caching usage scenarios 179
 change target vCenter Server 194
 coredump 204
 DHCP address 206
 DHCP reservations 213
 DHCP server 165, 213
 DNS Server 214
 EFI 165
 enable caching 180
 failing to complete boot 205
 failure to boot 207
 highly available 202
 host profiles 221
 image profile warning 205
 installing 215
 iPXE boot 219
 network boot problem 206
 networking 191
 PowerCLI cmdlets 163
 PowerCLI installation 211
 PowerCLI setup 211
 preparing 165
 proof of concept 210
 proof of concept checklist 210
 provisioning hosts 175
 rebooting 176
 redirection problem 204
 reference host 185, 220
 reprovisioning hosts with 176
 reregister 192
 rule set compliance 174
 rules 218
 static IP address 192
 tasks 162
 TFTP server 165, 207
 timeout error 203
 troubleshooting 203
 tutorial 210
 usage scenarios for caching 179
 user input 176
 VLANs 165
 wrong image 203
 See also vSphere Auto Deploy
Auto Deploy daemon 192
Auto Deploy image 217
Auto Deploy on vCenter Server Appliance 194
Auto Deploy PowerCLI 168, 171
Auto Deploy PowerCLI cmdlets 157
Auto Deploy roadmap 162
Auto Deploy rules 172, 173
Auto Deploy server 155
Auto Deploy upgrade 209
Auto Deploy with caching 182
auto-deploy register command 194
auto-partitioning 191

B

- banner, security **247**
- best practices, Auto Deploy **198**
- BIOS **249**
- BIOS UUID **155**
- boot command line options **141**
- boot commands, entering **140**
- boot disk, shared **259**
- boot failure in UEFI mode. **250**
- boot file (Auto Deploy) **165**
- boot operations **155**
- boot process, Auto Deploy **158**
- boot prompt **141**
- boot setting **249**
- boot.cfg file **152**
- bootloader kernel options **141**
- bulk licensing **169**

C

- CD-ROM, booting from virtual **250**
- CD/DVD, burning the ESXi ISO image **123**
- clearpart command **143**
- Client Integration Plug-in, installing **100**
- clients, firewall **23, 26**
- cluster location, assign with Auto Deploy **173**
- components included with the vCenter Server installer **72**
- computer name
 - Oracle **31**
 - SQL Server **31**
- configuration defaults, resetting **263**
- configuring ports **23, 26**
- configuring the keyboard **246**
- Connect-VIServer cmdlet **171–173, 177**
- connecting
 - Oracle database **46, 47**
 - SQL Server database **40**
- Copy-DeployRule cmdlet **177**
- creating a SQL Server database **33**
- creating an Oracle database **44**
- custom packages, removing **264**

D

- data source name **31**
- database corruption, Auto Deploy **208**
- Database Monitoring,enabling for Microsoft SQL Server user **42**
- Database Monitoring,enabling for Oracle user **48**
- database roles, setting vCenter user rights **34**
- databases
 - maintaining **32**

- Oracle **46**

- preparing **111**

- SQL Server **40, 41**

- DBO privileges **33**

- deactivating ESXi **264**

- default installation scripts **142**

- default root password **142**

- default storage behavior **256**

- defaults, restoring **263**

- deployment scenarios, vCenter Single Sign-On **57**

- deployment modes, vCenter Single Sign-On **52**

- depots **240**

- DHCP

- direct console **253, 254**

- for PXE booting the ESXi installer **128**

- vSphere Web Client **253**

- DHCP reservations, Auto Deploy **213**

- DHCP Scope **213**

- DHCP Server, Auto Deploy **213**

- DHCP server for Auto Deploy **165**

- direct console

- boot setting **249**

- configuring the keyboard **246**

- DHCP **253, 254**

- DNS **254**

- IP addressing **253, 254**

- management network **251, 252**

- navigating **246**

- network adapters **252**

- network settings **251, 252**

- password configuration **249**

- redirecting by setting the boot options **247**

- redirecting to a serial port **247, 248**

- security banner **247**

- static addressing **253, 254**

- testing management network **254, 256**

- VLAN ID **253**

- direct console, redirecting to a serial port in an Auto Deploy host **248**

- Directory Services **111, 113**

- disabling the management network **255**

- disk device names **152**

- Distributed Switch, See vSphere Distributed Switch

- distributed switches, permission **88**

- DNS **114, 254**

- DNS Requirements **27**

- DNS Server, Auto Deploy **214**

- DNS suffixes, direct console **254**

- domain controller **114**

- Download TFTP ZIP **165**

download the ESXi installer **135**
 download the vCenter Server installer **73**
 DRAC **28**
 dryrun command **143**
 DSN, 64-bit requirement **31**
 Dump Collector, *See* vSphere ESXi Dump Collector

E

EFI, Auto Deploy **165**
 enable caching **180**
 esxcli system coredump **187**
 ESXi
 about **246**
 deactivating **264**
 downloading the installer **135**
 installation options **121**
 installing **137**
 installing interactively **137**
 managing remotely **252**
 syslog service **260**
 system requirements **13**
 ESXi Dump Collector
 host profiles **189**
 reference host **189**
 ESXi hosts, licensing **265**
 ESXi Image Builder CLI, customized ESXi installation images **122**
 ESXi installation, required information **134**
 ESXi installation script, about **142**
 ESXi installation, Auto Deploy options **122**
 ESXi ISO image, burning on a CD/DVD **123**
 ESXi setup, post-setup **265**
 ESXi Shell access to host **262**
 ESXi, before you install **121**
 ESXi, installing **137**
 evaluation mode **122, 265**

F

factory defaults, restoring **263**
 FCoE, installing and booting ESXi from **134**
 firewall, network-based **114**
 floppy, booting from virtual **250**
 flowchart for installation and sign-in **69**
 folder location, assign with Auto Deploy **173**
 FTP **127**
 FTP Boot ZIP **216**

G

Get-Help PowerShell cmdlet **168**
 global data **111–113**
 gpupdate /force command **114**
 gPXE **127**

group policy update **114**
 groups, requirements **111**
 GUID **114**

H

hardware requirements
 ESXi **13**
 vCenter Server **17**
 vCenter Server Appliance **17**
 hardware requirements, ESXi **15**
 high availability, and vCenter Single Sign-On **54**
 highly available Auto Deploy environment **182**
 host customization **155, 178, 192, 195**
 host customizations **223**
 host image profile acceptance level **262**
 host profile from reference host **220**
 host profiles
 assign with Auto Deploy **172**
 Auto Deploy **188**
 Auto Deploy rule **221**
 caching **183**
 Network Coredump **188**
 reference host for Auto Deploy **220**
 Host profiles, stateful installs **184**
 host provisioning **155**
 hosts, reprovisioning with Auto Deploy **176**
 hosts firewall **23, 26**

I

IDE disks **13, 15**
 identity source, adding to vCenter Single Sign-On **85**
 identity sources for vCenter Single Sign-On **58**
 IIS, conflict with vCenter Server over port 80 **27**
 ILO **28**
 Image Builder
 and Auto Deploy **224**
 overview **224**
 See also ESXi Image Builder CLI
 Image Builder CLI, *See* vSphere ESXi Image Builder CLI
 Image Builder sessions **236**
 Image Builder, installing **231**
 Image Builder, workflows **240**
 image profile **217**
 image profiles
 acceptance level **239**
 editing **243**
 Image profiles, cloning **233**
 Image profiles, creating **233, 241**
 image profiles, export **235**
 image profiles, validation **226**

- ImageProfile structure **228**
- include command **143**
- install command **143**
- install vCenter Server as part of Simple Install **75**
- install vCenter Server in a Custom Install **83**
- install vCenter Server using Simple Install **73**
- install vCenter Server using Custom Install **76**
- install vCenter Single Sign-On using Simple Install **74**
- installation and sign-in **69**
- Installation overview **11**
- installation script
 - customized in ISO image **126**
 - path to **143**
 - supported locations **143**
- installation script, creating **140**
- installation scripts, default **142**
- installing
 - Client Integration Plug-in **100**
 - ESXi **137**
 - Update Manager server **103**
 - VirtualCenter Server **111**
 - VMware vSphere Web Client **80, 101**
- installing ESXi, scripted **140**
- installing ESXi interactively **137**
- installing ESXi with software FCoE **134**
- installorupgrade command **143**
- interactive installation **121**
- Inventory Service
 - install using Simple Install **73**
 - prerequisites for installing **49**
 - required information for installation or upgrade **62**
 - See also* vCenter Inventory Service
- Inventory Service database
 - back up on Linux **117**
 - back up on Windows **116**
 - restore on Linux **117**
 - restore on Windows **116**
- Inventory Service, install separately **81**
- Inventory Service,enabling IPv6 support **119**
- IP, on a detached host **251**
- IP address, vCenter Server with multiple network interfaces **113**
- IP addressing
 - direct console **253, 254**
 - vSphere Web Client **253**
- IPv6 **60**
- IPv6 address, format **60**
- IPv6 support, enabling for Inventory Service **119**
- iSCSI software disk, installing ESXi on **139**

- ISO
 - create **235**
 - export **235**
- ISO image, with custom installation script **126**

J

- Java Components (JRE), installing or upgrading separately **90**
- JDBC **41**
- JDBC URL formats **61**
- JRE,installing or upgrading separately **90**
- JVM heap settings, recommended for vCenter Virtual Appliance **17**

K

- keyboard command **143**
- keyboard, localizing **246**
- kickstart file, creating **140**
- ks.cfg **142**

L

- LDAP **112**
- license key **266**
- license key,accessing from the vSphere Web Client **266**
- licensed mode **122, 265**
- LicenseDataManager **169**
- licensing, bulk licensing **169**
- licensing ESXi hosts **265**
- Linked Mode
 - and databases **110, 111**
 - and permissions **110, 111**
 - reachability **114**
 - requirements **111**
 - troubleshooting **114**
- local Oracle database **46, 47**
- local SQL Server database **33**
- localizing, keyboard **246**
- lockdown mode
 - enabling **261**
 - vSphere Web Client **261**
- log in to vCenter Server **57**
- logging, providing space for **23**
- logging in to vCenter Server **57**
- logical volume management **256**
- Lookup Service, *See* vCenter Lookup Service
- LVM **256**

M

- MAC address **130, 155**
- maintaining the database **32**
- managed entities, permissions **88**
- management agents, restarting **255**

- management network
 - direct console **251, 252**
 - disabling **255**
 - restarting **255**
 - testing **254, 256**
- media options, ESXi installer, supported **123**
- memory, ESXi requirements **13, 15**
- message, security **247**
- Microsoft .NET **168, 231**
- Microsoft .NET Framework **72**
- Microsoft PowerShell **168, 231**
- Microsoft SQL Server, requirements **30**
- Microsoft SQL database permissions, setting by
 - using the dbo schema and db_owner database role **35**
- Microsoft SQL Server 2008 R2 Express **72**
- Microsoft SQL Server database schema,
 - creating with a script (recommended method) **36**
- Microsoft Windows
 - authentication for SQL Server **60**
 - system account **60**
- Microsoft Windows Installer **72**

N

- navigating, direct console **246**
- network adapters, direct console **252**
- network boot **212**
- network command **130, 143**
- network core dump **187**
- network drive, installing from **62**
- network settings, direct console **251, 252**
- New-DeployRule **218, 221**
- New-DeployRule cmdlet **171–173**
- New-EsxImageProfile cmdlet **233**
- NewEsxImageProfile cmdlet **241**
- non-ASCII characters, disable support for **264**

O

- ODBC databases **40**
- offline bundle
 - create **235**
 - export **235**
- online Help, deploying locally **102**
- OpenLDAP Server identity source **87**
- Oracle database
 - changing the computer name **31**
 - remote access **46**
 - requirements **30**
 - script for creating **44**
 - user **43**
- Oracle database schema **44**
- Oracle, preparing database **46**

P

- paranoid command **143**
- part command **143**
- partition command **143**
- partitions **256, 258**
- password, administrative **249**
- permissions
 - assigning **87**
 - distributed switches **88**
 - inheritance **88**
- plug-ins for vCenter Server **100**
- port 80 conflict between vCenter Server and IIS **27**
- ports
 - configuring **23, 26**
 - firewall **23, 26**
- ports used by vCenter Server **23**
- ports used by vCenter Server Appliance **26**
- PowerCLI **171**
- PowerCLI cmdlets, Auto Deploy **163**
- PowerCLI sessions **236**
- PowerCLI wildcard characters **240**
- Preface **7**
- preinstallation checklist **210**
- preparing database **47**
- prerequisites for installing vCenter Server **49**
- prerequisites for installing vCenter Single Sign-On **49**
- prerequisites for installing vCenter Inventory Service **49**
- PXE, configuration files **130**
- PXE boot ESXi installer using PXELINUX, setup procedure **130, 131, 133**
- PXELINUX
 - boot ESXi installer using **130, 133**
 - boot ESXi installer using **131**

R

- redirecting log files **259**
- reference host
 - Auto Deploy **188**
 - configuration options **187**
- reference host for Auto Deploy **185**
- registration of Auto Deploy **192**
- registry settings **114**
- reinstalling vCenter Server **109**
- remote access, restricting **261**
- remote management applications **134**
- remote management of ESXi **252**
- remote Oracle database **46**
- remote SQL Server database **33**
- removing 3rd-party modules **264**

- removing custom packages **264**
- removing vCenter Server **109**
- Repair-DeployRulesetCompliance cmdlet **174**
- requirements for vSphere Web Client **22**
- resetting configuration defaults **263**
- restarting the management agents **255**
- restarting the management network **255**
- restoring, factory defaults **263**
- restricting access **261**
- ROM image **127**
- root access, restricting **261**
- root password **249**
- rootpw command **143**
- RPCCfg.exe **114**
- RSA **28**
- rule rule enginesets **157**
- rule set **155**
- rule set compliance **174**
- rules **157, 221**
- rules engine **157**

S

- SAS disks **13, 15**
- SATA disks **13, 15**
- scratch partition, enabling **258**
- scratch storage **256, 258**
- script, for installing ESXi **142**
- script for Oracle database **44**
- script for SQL Server database **33**
- scripted installation, differences from ESXi 4.x **151**
- scripted installation of ESXi, by PXE Booting **155**
- scripted installation of ESXi, from a USB flash drive **154**
- scripted installation of ESXi, from a CD or DVD **153**
- scripted installation option **121**
- SCSI **13, 15**
- SDK **114**
- security **60**
- security banner **247**
- Security Token Service **56**
- serial port
 - redirecting the direct console from the vSphere Web Client **248**
 - redirecting the direct console to **247**
- services, syslogd **260**
- sessions, PowerCLI **236**
- Simple Install **73**
- Single Sign-On, *See* vCenter Single Sign-On
- SMBIOS information **155**

- snap-in, Auto Deploy **168**
- software depot **217**
- software depots, examining **240**
- SoftwarePackage structure **228**
- specifications
 - ESXi hardware requirements **13, 15**
 - performance recommendations **13, 15**
- SQL compatibility mode **68**
- SQL Server
 - changing the computer name **31**
 - Microsoft Windows authentication **60**
 - preparing the database **40, 41**
 - script for creating **33**
- SSH access to host **262**
- Standard switch, restoring **256**
- state **155**
- stateful installs **184**
- stateless caching **183, 202**
- static addressing, about **251, 252**
- static DNS **254**
- static DNS, direct console **254**
- static IP **253**
- static IP address for Auto Deploy **192**
- storage **256**
- STS (Security Token Service) **56**
- subnet mask **253**
- support information **266**
- synchronize ESX/ESXi clocks on vSphere network **59**
- synchronizing clocks on the vSphere network **59**
- syslog, host profile **190**
- Syslog Collector, *See* vSphere Syslog Collector
- syslog, Auto Deploy **190**
- system requirements, vCenter Server database **30**
- system swap **256, 258**

T

- target hosts **212**
- TCP/IP setup for SQL Server **41**
- template host for Auto Deploy **185**
- Test-DeployRuleSetCompliance cmdlet **174**
- testing management network, direct console **256**
- TFTP **127**
- TFTP server
 - Auto Deploy **207**
 - installing **211**
- TFTP server for Auto Deploy **165**
- tftp-hpa **127**
- tfptd32 **127**
- timeout error, Auto Deploy **203**

Tomcat **114**
troubleshooting, Linked Mode **114**

U

UEFI mode, ESXi fails to boot **250**
uninstalling vCenter Server **109**
unregister Auto Deploy **194**
updated information **9**
upgrade command **143**
upgrading ESXi, scripted **140**
upgrading vSphere Web Client **80, 101**
USB, bootable ESXi installation **123**
USB, ESXi installation script **125**
user and group for Oracle database **43**
user input **223**
user input for Auto Deploy **178**
user input for Auto Deploy hosts **176**
user repositories for vCenter Single Sign-On **58**

V

vCenter Inventory Service
 hardware requirements **17**
 updating URL **118**
vCenter Lookup Service **56**
vCenter Orchestrator **72**
vCenter Server
 additional components **100**
 before you install **29**
 components **72**
 downloading the installer **73**
 hardware requirements **17**
 install using Simple Install **73**
 install as part of Simple Install **75**
 install in a Custom Install **83**
 installing **69**
 installing from a network drive **62**
 installing on IPv6 machine **60**
 joining a group **111–113**
 Linked Mode **109**
 logging in **57**
 plug-ins **100**
 ports **23**
 prerequisites for installing **49**
 required information for installation or upgrade **62**
 required information for vCenter Server installation **62**
 requirements for joining a group **111**
 setting the administrator user **56**
 setting user rights through database roles **34**
 software requirements **22**
 system requirements **13**

vCenter Server Appliance, redirect to the vCenter Single Sign-On of a different virtual machine **95**
vCenter Server databases, preparing **30**
vCenter Server installation, post-installation **99**
vCenter Server administrator user, setting **51**
vCenter Server Appliance
 Auto Deploy **193, 194**
 configuration file format **96**
 ports **26**
 synchronize clock with NTP server **60**
 See also VMware vCenter Server Appliance
vCenter Server Appliance, creating a custom password **94**
vCenter Server Appliance: Auto Deploy **193, 194**
vCenter Server database
 Microsoft SQL Server **32**
 Oracle **42**
vCenter Server MSSQL database objects, creating manually with a script **37**
vCenter Server to Server, installing or upgrading separately **91**
vCenter Server upgrade **209**
vCenter Server VMware vCenter Server - tc Server Settings **114**
vCenter Simple Install **73**
vCenter Single Sign-On
 Active Directory **85**
 deployment modes **52**
 deployment scenarios **57**
 effect on vCenter Server installation and upgrades **51**
 high availability **54**
 identity sources **58, 85**
 install using Simple Install **73**
 installation fails **91**
 LDAP **85**
 OpenLDAP **85**
 Prerequisites for Installing vCenter Single Sign-On, Inventory Service, and vCenter Server **49**
 redirect vCenter Server Appliance to **95**
 required information for installation or upgrade **62**
 User repositories **58**
vCenter Single Sign-On , custom install first or only instance **77**
vCenter Single Sign-On, install additional node at existing site **78**
vCenter Single Sign-On, install additional node at new site **79**
vCenter Single Sign-On, install using Simple Install **74**
vCenter Virtual Appliance, JVM heap settings **17**

- VIB, third party **224**
- VIB structure **228**
- VIBs **224, 225**
- VIBs, acceptance levels **227**
- VIBs, validation **226**
- viewing, log files **259**
- virtual CD **134**
- virtual machine console, installing **100**
- virtual machines, RAM requirements **13, 15**
- virtual media **250**
- VirtualCenter Management Webservices **99**
- VLAN ID, direct console **253**
- VLANs, Auto Deploy **165**
- vmaccepteula command **143**
- VMFS **256**
- vmk0 **191**
- vmkernel module, removing **264**
- VMware vCenter Server - tc Server settings in vCenter Server **114**
- VMware Directory Service **56**
- VMware vCenter Management Webservices **116**
- VMware vCenter Server Appliance
 - downloading and deploying **92**
 - hardware requirements **17**
 - software requirements **22**
- VMware vSphere Web Client, installing or upgrading **80, 101**
- vmware-fdm **209**
- vmware-rbd-watchdog **194**
- vSphere Authentication Proxy
 - IIS installation causes port 80 conflict **27**
 - install or upgrade **108**
- vSphere Auto Deploy
 - installing ESXi with **155**
 - installing or upgrading **107**
- vSphere CLI **258**
- vSphere Distributed Switch, restoring standard switch **256**
- vSphere ESXi Dump Collector, install or upgrade **104**
- vSphere ESXi Image Builder CLI, using **223**
- vSphere installation and setup, introduction **11**
- vSphere Syslog Collector, install or upgrade **106**
- vSphere Update Manager **72**
- vSphere Update Manager,databases **30**
- vSphere Web Client
 - DHCP **253**
 - hardware requirements **17**
 - managing ESXi host **265**
 - online Help **102**
 - requirements **22**
 - static addressing **253**

See also VMware vSphere Web Client

vSwitch0 **191**

VWS **114**

W

waiter.tgz file **203**

web client, See VMware vSphere Web Client

wildcard characters, PowerCLI **240**

working rule set **157**