



VMware® ESXi™ 4.1 Operations Guide

TECHNICAL WHITE PAPER

Introduction

The hypervisor architecture of VMware® vSphere™ 4.1 (“vSphere”) plays a critical role in the management of the virtual infrastructure. The introduction of the bare-metal VMware ESX® architecture in 2001 significantly enhanced performance and reliability, which in turn enabled customers to extend the benefits of virtualization to their mission-critical applications. The introduction of the VMware ESXi™ architecture represents a similar leap forward in reliability and virtualization management. Less than 5 percent as large as VMware ESX, VMware ESXi runs independently of an operating system (OS) and improves hypervisor management in the areas of security, deployment and configuration, and ongoing administration. Yet none of this comes at the cost of functionality. All of the features offered by VMware vSphere 4.1, such as VMware vMotion™ (vMotion), VMware Storage vMotion (Storage vMotion), VMware High Availability (VMware HA), VMware Fault Tolerance (VMware FT) and VMware Distributed Resource Scheduler (VMware DRS), are fully supported on the VMware ESXi architecture.

This paper describes the architecture of VMware ESXi and then explains how various management tasks are performed in it. This information can be used to help plan a migration to the VMware ESXi architecture from the legacy VMware ESX framework and to improve or enhance day-to-day operations.

Architecture

In the original VMware ESX architecture, the virtualization kernel (VMkernel) is augmented by a management partition known as the console operating system (COS) or service console. The primary purpose of the COS is to provide a management interface with the host. Various VMware management agents are deployed in the COS, along with other infrastructure service agents (for example, name service, time service, logging, and so on). In this architecture, many customers deploy other agents from third parties to provide a particular functionality, such as hardware monitoring and system management. Furthermore, individual administrative users log in to the COS to run configuration and diagnostic commands and scripts.

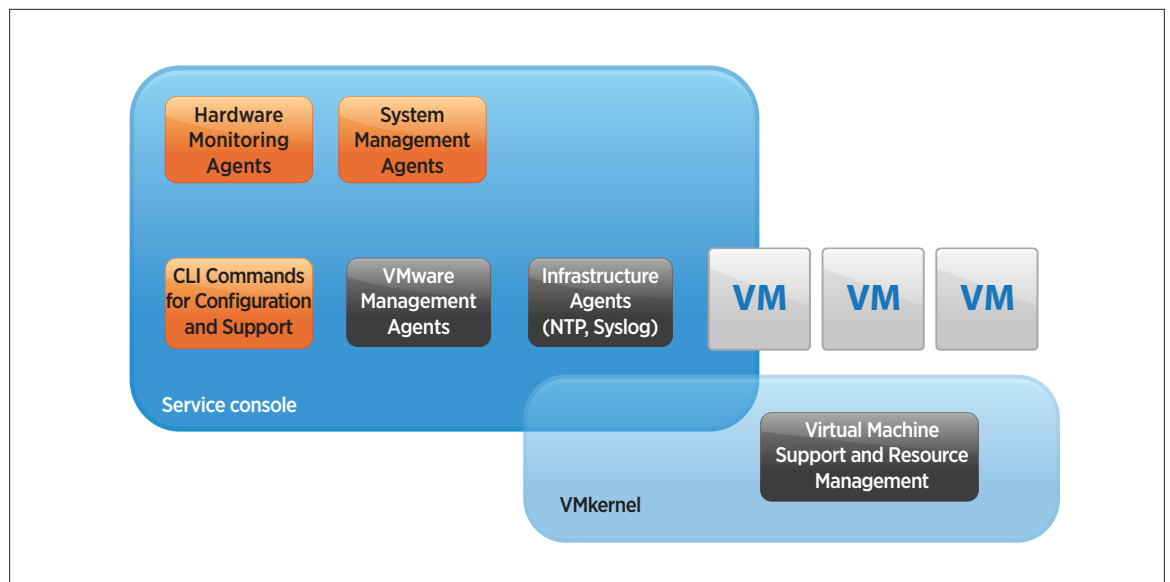


Figure 1. Architecture of VMware ESX

In the VMware ESXi architecture, the COS has been removed, and all of the VMware agents run directly on the VMkernel. Infrastructure services are provided natively through modules included in the VMkernel. Other authorized third-party modules, such as hardware drivers and hardware monitoring components, can run in the VMkernel as well. Only modules that have been digitally signed by VMware are allowed on the system, creating a tightly locked-down architecture. Preventing arbitrary code from running on the VMware ESXi host greatly improves the security and stability of the system.

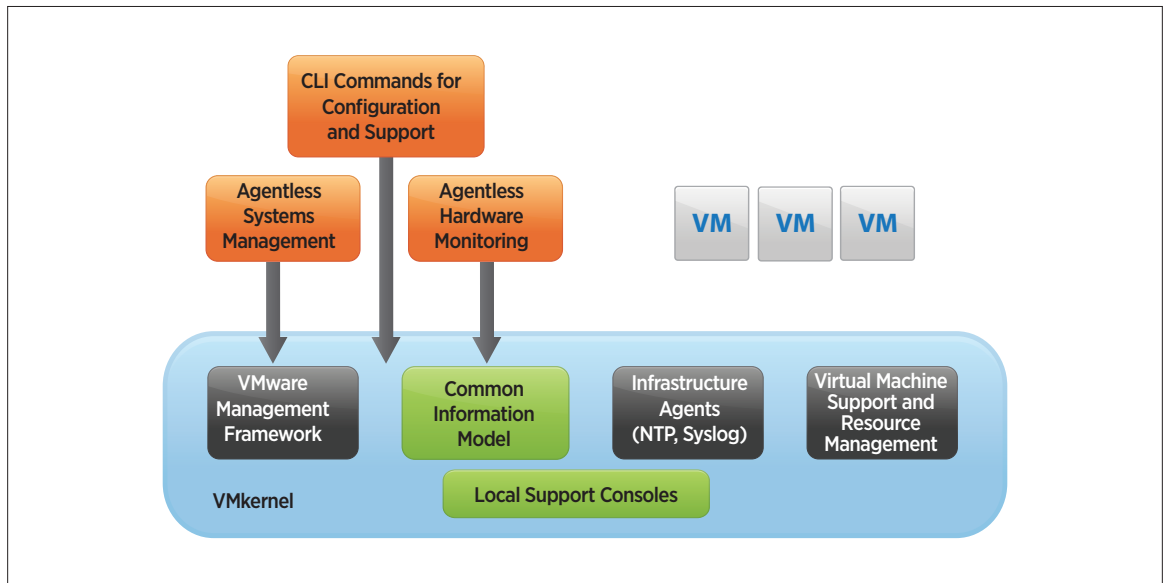


Figure 2. Architecture of VMware ESXi

Management

The management functionality that was provided by agents in the VMware ESX architecture is now exposed via APIs in the VMware ESXi architecture. This allows for an “agentless” approach to hardware monitoring and system management. VMware also created remote command lines, such as the VMware vSphere 4 Command-Line Interface (vCLI) and VMware vSphere 4 PowerCLI (PowerCLI), to provide command and scripting capabilities in a more controlled manner. These remote command-line sets include a variety of commands for configuration, diagnostics and troubleshooting. For low-level diagnostics and the initial configuration, menu-driven and command-line interfaces are available on the local console of the server. The following sections discuss individual management topics and describe how tasks are performed in the VMware ESXi architecture.

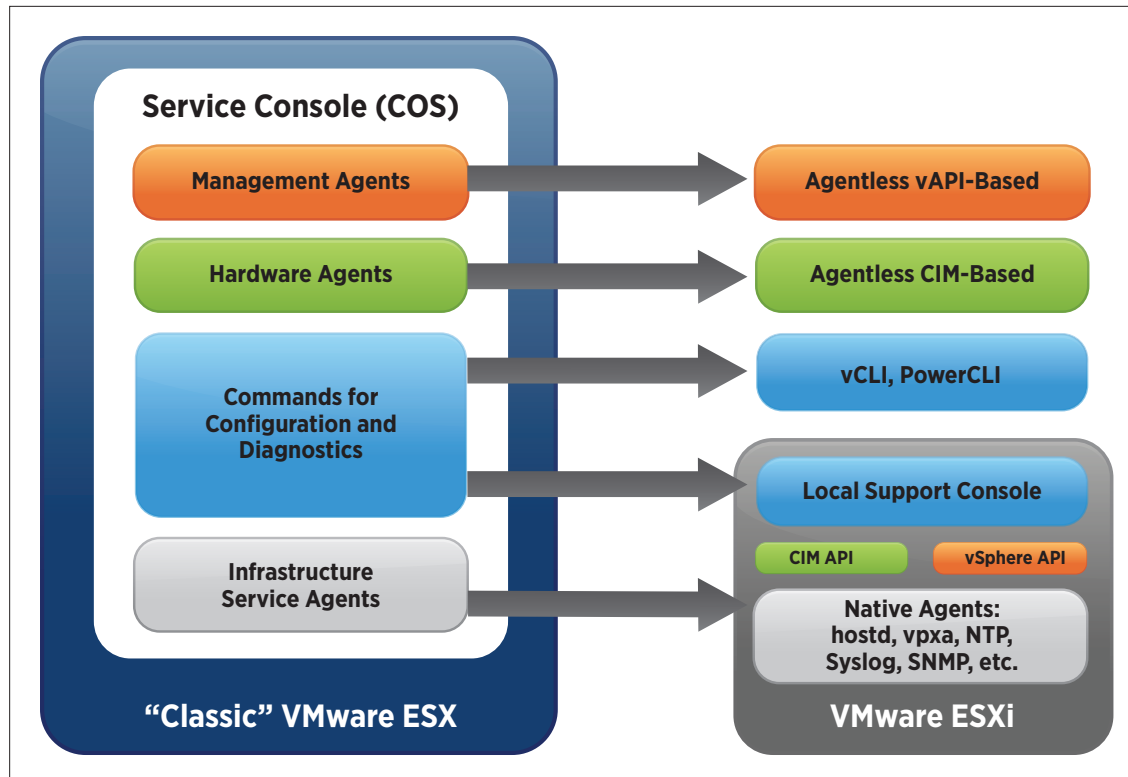


Figure 3. New and Improved Paradigm for VMware ESX Management

Automation

To automate the management of a VMware ESXi deployment, VMware has created easy-to-use scripting tools. Users can write scripts with the same functionality as the vSphere client to automate manual tasks, enabling efficient management and configuration of small- to large-scale environments. These tools work well with both VMware ESXi and VMware ESX hosts, empowering users to administer mixed environments easily.

PowerCLI is a robust command-line tool for automating all aspects of vSphere management, including host, network, storage, virtual machine, guest OS and more. PowerCLI is distributed as a Windows PowerShell snap-in. It includes more than 150 PowerShell cmdlets, along with documentation and samples. PowerCLI seamlessly blends the vSphere platform with Windows and .NET, which means you can use PowerCLI by itself or within many different third-party tools.

vCLI is a set of more than 30 command-line utilities that help users provision, configure, troubleshoot and maintain VMware ESX and VMware ESXi hosts. Where PowerCLI is better suited for large-scale automation, vCLI is aimed at users who feel more comfortable with the legacy COS commands. There are commands that can completely automate the initial configuration of a VMware ESXi host, and others that provide troubleshooting and diagnostic capabilities. VMware provides vCLI packages for installation on both Windows and Linux systems and is available prepackaged as part of the VMware vSphere 4.1 Management Assistant (vSphere Management Assistant).

vCLI has numerous commands for troubleshooting, including:

- vmkfstools
- vmware-cmd
- resxtop

In vSphere 4.1, important enhancements make the vCLI more powerful:

- Performs host operations, such as rebooting and entering or exiting maintenance mode, using the “vicfg-hostops” command
- Configures Microsoft Active Directory using the “vicfg-authconfig” command
- Configures IPsec with “vicfg-ipsec”
- Forcibly terminates a virtual machine, even when it is not responding to normal shutdown commands, using the “esxcli vms” command
- Configures storage to a greater extent, including various software iSCSI parameters and storage plug-ins, using a series of new options to the “esxcli” command
- Employs additional diagnostic capabilities for networking and storage, including:
 - The “esxcli network” command, which lists active connections or active ARP table entries
 - New options for “resxtop,” which show NFS statistics

Both PowerCLI and vCLI are built on the same interfaces as the vSphere client. They can be pointed directly at a VMware ESXi host or at VMware vCenter™. When pointed at a host, they can execute commands directly on a VMware ESXi host, similarly to how a command in the COS of VMware ESX operates on only that host. Local authentication is required in this case. Alternatively, when communicating through VMware vCenter, the vCLI and PowerCLI commands benefit from the same authentication (for example, Active Directory) roles and privileges and event logging as vSphere client interactions. This provides for a much more secure and auditable management framework.

The VMware vSphere 4.1 Management Assistant is a virtual appliance that packages the vCLI, the VMware vSphere 4.1 SDK for Perl, as well as a logging module (called “vi-logger”) and authentication modules for unattended script execution (called “vi-fastpass”) into one convenient bundle.

The following table contains different categories of operational procedures and the preferred tool for each category. We have rated each tool per task to classify the level of expertise required:

TASK	vCLI/vMA	POWERCLI
Reporting	Normal	Easy
Monitoring	Hard	Normal
Configuration	Easy	Easy
Automation	Normal	Easy
Troubleshooting	Easy	Hard

Table 1. Ease of Operational Tasks

Although each of the tools can be used to accomplish a given task, the preceding table can be used as an indication of which tools would best meet your requirements.

Deployment

Various deployment methods are supported for VMware ESXi, such as booting the installer off of a DVD or over PXE, and deploying the VMware ESXi image onto a local disk over the network using a variety of protocols, including secure HTTP. VMware ESXi 4.1 enables users to do a scripted installation of the VMware ESXi software onto the local disk of a server, analogous to the kick-start mechanism used for VMware ESX architecture. The scripted installation configuration file (typically named “ks.cfg”) can also specify the following scripts to be executed during the installation:

- Preinstall
- Postinstall
- First-boot

These scripts are run locally on the VMware ESXi host and can perform various tasks, such as configuring the host’s virtual networking and joining it to VMware vCenter Server. These scripts can be written in either the Tech Support Mode shell or Python.

Support for boot from SAN has been added to VMware ESXi 4.1. This support includes Fibre Channel SAN, as well as iSCSI and FCoE for certain storage adapters that have been qualified for this capability.

VMware ESXi 4.1 is still available preinstalled on Flash drives on certain server models available from a number of hardware OEM vendors. (Consult the server HCL to determine which combinations of server and USB or Flash drive are supported.)

As stated, with vSphere 4.1, VMware has added scripted installation capabilities to VMware ESXi. A basic scripted CD-ROM-based install entails the following procedure:

1. Boot from the VMware ESXi CD-ROM
2. Press “Tab” when the “VMware VMvisor Boot Menu” is displayed
3. Edit the string so that it includes the location of your script:

```
> mboot.c32 vmkboot.gz ks=http://<ip-address>/ks.cfg --- vmkernel.gz --- sys.vgz ---
cim.vgz --- ienviron.vgz --- install.vgz
```

When including the location of your script, ensure to append it after “vmkboot.gz” and before “--- vmkernel.gz” for the script to function correctly. The <ip-address> should be replaced with the ip-address of the Web server hosting the configuration file. The ks.cfg configuration file can also be located on other types of media such as CD-ROM or an FTP server. For more details, refer to the *VMware vSphere 4.1 ESXi Installable and vCenter Server Setup Guide*.

It is also possible to PXE boot the VMware ESXi installer. This however requires a TFTP server that supports PXE boot, gPXE and a modification to your DHCP server to allow the DHCP server to send the host the correct TFTP and PXE information. For more details, refer to the *VMware vSphere 4.1 ESXi Installable and vCenter Server Setup Guide*, where this procedure is fully documented.

Whether using a PXE mechanism to facilitate the installation or a CD-ROM, a so-called answer script is required. The script follows a standardized format to supply the installer with the correct parameters. The following example includes a postinstallation action, and actions on the first boot, to demonstrate the endless capabilities the VMware ESXi installer offers. These will be explained in detail as follows:

```
# Accept the VMware End User License Agreement
vmaccepteula

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

```

# Choose the first discovered disk to install onto
autopart --firstdisk --overwritevmfs

# The installation media is in the CD-ROM drive
install cdrom

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

# A sample post-install script
%post --unsupported --interpreter=busybox --ignorefailure=true

# Download drivers required to access the network after a reboot
wget http://192.168.1.1/drivers.zip-0/vmfs/volumes/datastore1/drivers.zip

# A sample of the script that will run on first boot
%firstboot ---unsupported ---interpreter=busybox

# Installation of the drivers for network access
esxupdate --bundle=/vmfs/volumes/datastore1/drivers.zip update

# Configuration of NTP Servers
echo restrict default kod nomodify notrap noquerynopeer > /etc/ntp.conf
echo restrict 127.0.0.1 >> /etc/ntp.conf
echo server 10.0.0.11 >> /etc/ntp.conf
echo server 10.0.0.12 >> /etc/ntp.conf
echo driftfile /var/lib/ntp/drift >> /etc/ntp.conf
/sbin/chkconfig --level 345 ntpd on
/etc/init.d/ntpd stop
/etc/init.d/ntpd start

# Rename the local datastore so that it includes the hostname
vim-cmd hostsvc/datastore/rename datastore1 $(hostname -s)-datastore01

# Add an extra nic to vSwitch0 (vmnic0 and vmnic2) and a VLAN ID
esxcfg-vswitch -L vmnic2 vSwitch0
esxcfg-vswitch -v 10 -p 'Management Network' vSwitch0

# Add vMotion Portgroup to vSwitch0
esxcfg-vswitch -A vMotion vSwitch0

# Assign an ip-address to the vMotion VMkernel and a VLAN ID to the Portgroup
esxcfg-vswitch -v 20 -p vMotion vSwitch0
esxcfg-vmknic -a -i 192.168.2.41 -n 255.255.255.0 vMotion

# Wait to ensure everything has been created and refresh the network stack
sleep 5
vim-cmd hostsvc/net/refresh

# Enable vMotion on the newly created VMkernel vmk1
vim-cmd hostsvc/vmotion/vnic_set vmk1

```

```

# Add new vSwitch for VM traffic (vmnic1 and vmnic3)
esxcfg-vswitch -a vSwitch1

#Create a standard portgroup for VMs to vSwitch1 and set a VLAN ID
esxcfg-vswitch -A Production_VLAN5 vSwitch1
esxcfg-vswitch -v 30 -p Production_VLAN5 vSwitch1

# Add NICs to the new vSwitch
esxcfg-vswitch -L vmnic1 vSwitch1
esxcfg-vswitch -L vmnic3 vSwitch1

# Wait to ensure everything has been created and refresh the network stack
sleep 5
vim-cmd hostsvc/net/refresh

```

This example script shows how to automate the installation of a VMware ESXi host that requires the download of a driver package before the host is rebooted by the installation process, and the installation of this driver package after the first boot. Although this scenario is rare, it shows the flexibility you have when developing these scripts.

The tool “wget” and the use of persistent storage on a datastore enable you to download drivers, additional scripts and much more. Furthermore we display how to enable and configure NTP and how to create additional vSwitches and port groups including VLAN IDs.

Of course, there are many more manual steps that can be automated through the use of standard CLI commands such as (but not limited to) `esxcli`, `esxcfg-*` and `vim-cmd`.

It is important to recognize the difference between the `%post` and the `%firstboot` section. The `%firstboot` section is the section that is most commonly used for configuring the VMware ESXi host. It is executed during the first boot after the installer has completed. The following diagram depicts the process of a scripted installation where both the `%post` and `%firstboot` section are used:

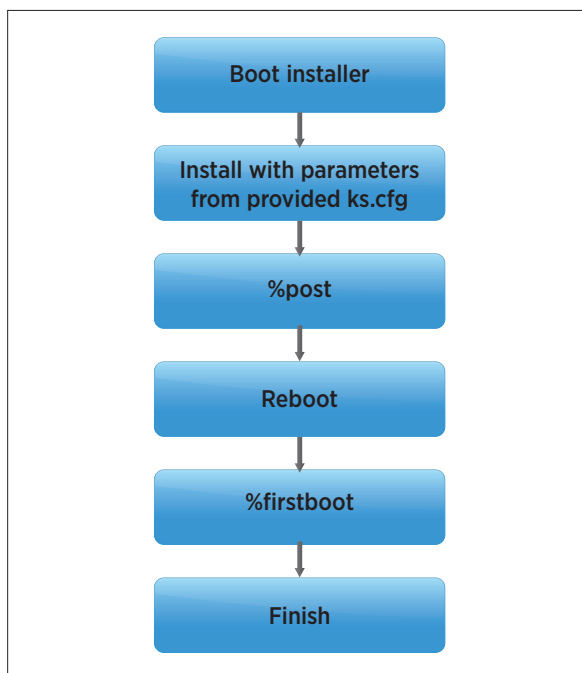


Figure 4. Scripted Installation Process

Installation Destination

When planning the implementation of—or migration to—VMware ESXi, one of the first decisions that must be made concerns the type of installation destination to be used. The form factor of VMware ESXi enables it to be installed on multiple different installation destination types, which include the following:

- Local disk (including SSD)
- Removable media
 - SB
 - SD
- Boot from SAN
 - FC
 - iSCSI

Local disks are a popular installation destination. Local disk installations provide two advantages over removable devices: resiliency and the level of automation. Resiliency refers to the ability to run two local disks in RAID-1. Although VMware ESXi is loaded into memory, it must write its configuration once every 10 minutes on average. In the case of boot media failure, this might be at risk, possibly resulting in a loss of configuration changes. Local disk installations also enable you to implement a scripted installation mechanism. This method is also supported for boot from SAN (iSCSI/FC) but currently not supported for removable devices such as USB and SD media. There must be at least 5GB of disk space available for a local disk installation.

Removable devices such as USB and SD have always been one of the top VMware ESXi installation destinations due to the flexibility and cost factors associated with them. These devices typically have a shorter life span than hard disks and therefore impose a minor risk. Hardware vendors have found a solution that increases resiliency by offering a “dual-SD module” configuration. And many customers have mitigated the risk by using enterprise-grade USB/SD modules and having one or more of them on hand.

From an operational perspective, the flexibility and resiliency offer many benefits, although there is one restriction. Scripted installation of VMware ESXi on a removable device is currently not supported. This can be mitigated through the use of host profiles or by automating the configuration through PowerCLI.

Requirements:

- VMware supports removable devices only under these conditions:
 - The server on which you want to install VMware ESXi 4.x is in the [VMware ESXi 4.x Hardware Compatibility Guide](#).

and

- You have purchased a server with VMware ESXi 4.x embedded on the server from a certified vendor.

or

- You have used a USB or SD Flash device that is approved by the server vendor for the particular server model on which you want to install VMware ESXi 4.x on a USB or SD Flash storage device.

As of vSphere 4.1, support for boot from SAN—both FC and iSCSI—has been included. Boot from SAN gives you resiliency and enables you to leverage the flexibility of a diskless server while still providing you with the option to do a scripted installation.

Requirement:

- Support for boot from SAN for storage device and adapters

Each type of installation media has its benefits. Depending on the environment, all media should be considered. Based on requirements and constraints regarding budget, licensing and array capabilities, a decision must be made on a per-case basis. Generally speaking, using “local disks” is the most compelling option because it enables you to fully automate your installation, in comparison to USB/SD, and it is relatively inexpensive, in comparison to boot from SAN.

Hardware Monitoring

The Common Information Model (CIM) is an open standard that defines a framework for agentless, standards-based monitoring of hardware resources for VMware ESXi. This framework consists of a CIM object manager, often called a CIM broker, and a set of CIM providers.

CIM providers are the mechanisms that provide management access to device drivers and underlying hardware. Hardware vendors, including server manufacturers and specific hardware device vendors, can write providers to supply monitoring and management of their particular devices. VMware also writes providers that implement monitoring of server hardware, VMware ESXi storage infrastructure and virtualization-specific resources. These providers run inside the VMware ESXi system and are designed to be extremely lightweight and focused on specific management tasks. The CIM broker takes information from all CIM providers and presents it to the outside world via standard APIs, the most common one being WS-MAN. Any software tool compatible with one of these APIs, such as HP SIM or Dell OpenManage, can read this information, monitoring the hardware of the VMware ESXi host.

One consumer of the CIM information is VMware vCenter. Through a dedicated tab in the vSphere client, users can view the hardware status of any VMware ESXi host in their environment, providing a single view of the physical and virtual health of their systems. Users can also set VMware vCenter alarms to be triggered on certain hardware events, such as temperature or power failure and warning states.

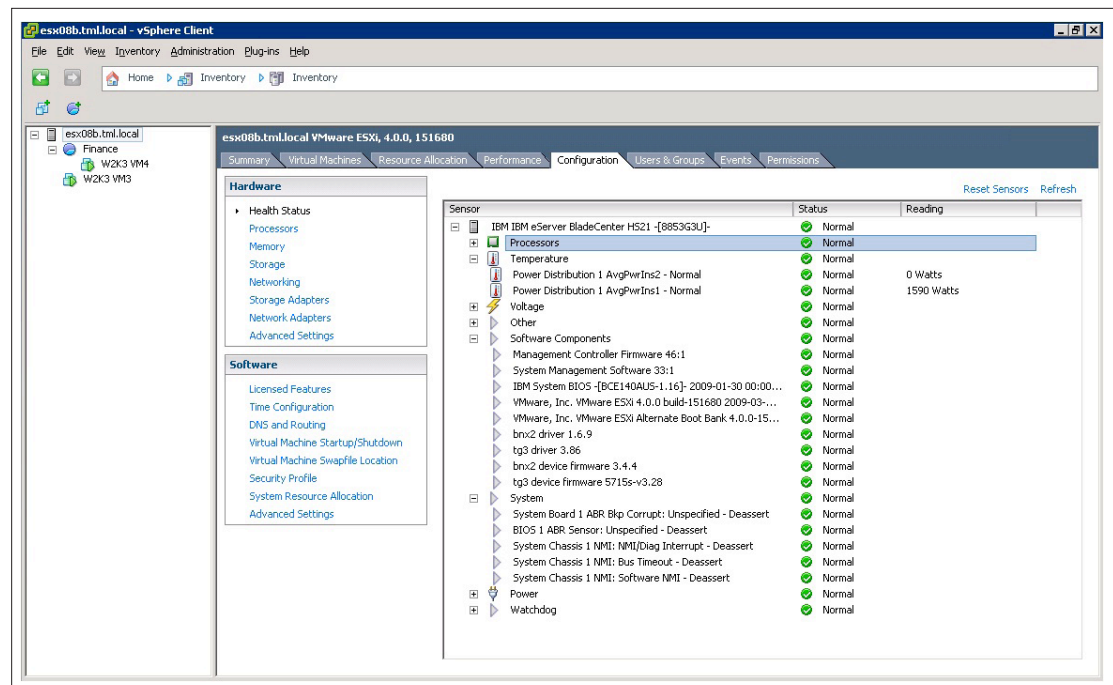


Figure 5. Hardware Monitoring in VMware vCenter Server

VMware ESXi also exposes hardware status information via SNMP for other management tools that rely upon that standard. SNMP traps are available from both the VMware ESXi host and VMware vCenter. VMware ESXi 4.1 currently supports SNMPv2, and it can be configured using the vCLI command “vicfg-snmp.”

Firmware Upgrades

Upgrading firmware on any platform is a cumbersome task. Historically, customers who have used the COS have upgraded the firmware with tools provided by the respective vendor. With VMware ESXi, that approach will no longer work, due to the absence of the COS. Firmware upgrades will, however, still periodically need to be applied. Currently, VMware ESXi offers no native functionality. The following work-arounds exist to solve this problem:

1. Hardware vendor bootable upgrade CD-ROM/DVD
2. PXE boot of vendor's upgrade CD-ROM/DVD
3. Hardware vendor VMware vCenter plug-in or management application
4. PXE boot of small Linux distribution

Many vendors offer a bootable CD-ROM/DVD that contains all drivers and firmware code required. These are typically categorized per server model and can be used to boot a host off and manually upgrade the appropriate devices. This solution typically is used in environments of up to 10 hosts. For larger environments, we have found that using a PXE boot configuration in conjunction with the vendor-provided upgraded CD-ROM/DVD can be a flexible alternative. The overall transfer size of the total package might be a constraint.

Several hardware vendors provide management plug-ins for VMware vCenter Server that enable you to manage firmware upgrades from within the vSphere Client. These plug-ins often offer reporting capabilities that reduce the chances of inconsistency across the virtual infrastructure. Large environments typically use a centralized management application to manage hardware end to end, which also includes the capabilities to upgrade firmware.

Finding a unified solution to manage firmware and patches in an environment where multiple types of hardware from different vendors are used can be a challenge. Creating a custom, slimmed-down Linux appliance that identifies the hardware configuration and updates the firmware accordingly can solve this problem. Solutions such as these typically use a PXE boot configuration with a central repository for the different types of firmware for this environment. This does require extensive knowledge of the various components and a substantial effort with regard to development, but it ultimately leads to a highly flexible and scalable solution that enables you to update any of the hardware components.

We advise managing the firmware level consistently and following the hardware vendor's recommendations, to avoid running into any interdependency issues. We also recommend that when you are acquiring new hardware, you look into the level of integration and the mechanisms that can be leveraged around managing your hardware. Especially in converged shared platforms, availability and manageability are key to the success of your IT department.

Systems Management and Backup

Systems management and backup products integrate with VMware ESXi via the vSphere APIs, which have been significantly enhanced in vSphere 4.1 through agentless partner integration. The API-based partner integration model significantly reduces management overhead by eliminating the need to install and manage agents in the COS.

VMware has worked extensively with our ecosystem to transition all partner products to the API-based integration model of VMware ESXi. As a result, BMC, CA, HP, IBM, EMC, NetIQ, Quest Software, Commvault, Vizioncore, Double-Take Software, SteelEye and Symantec are among the majority of systems management and backup vendors in the VMware ecosystem that have products that support VMware ESXi today. If you are using an agent-based partner solution to integrate with VMware ESX, check with your vendor to see if a newer version of the product supports VMware ESXi.

VMware also includes backup capability with the vSphere product suite. VMware Data Recovery is a robust, simple-to-deploy backup and recovery solution that businesses should consider to provide the first line of data protection for their virtual environment.

VMware Data Recovery enables:

- Full image backup of virtual machines
- Full and incremental recovery of virtual machines, plus recovery of individual files and directories

Patching and Updating

The patching and updating of VMware ESXi enables flexibility and control. During the patching process, only the specific modules being updated are changed. The administrator can preserve any previous updates to other components. Whether installed on disk or embedded Flash memory, VMware ESXi employs a “dual-image” approach, with both the current and prior version present. When a patch is installed, the new image is constructed and overwrites the prior image. The current version becomes the prior version and the system boots off the newly written image. If there is a problem with the image, or if the administrator wishes to revert to the prior one, the host is simply rebooted off the recent, good image.

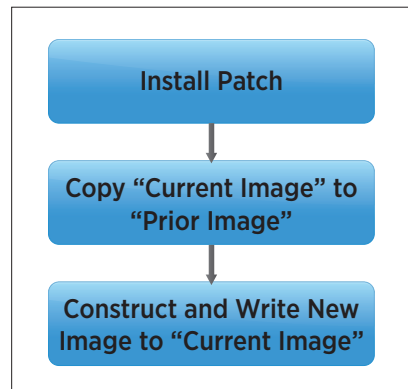


Figure 6. Workflow for Installing Patches

VMware vCenter Update Manager (Update Manager) is a VMware vCenter plug-in patch-management solution for vSphere. Update Manager enables centralized, automated patch and version management for vSphere. It offers support for VMware ESX/ESXi hosts, virtual machines and virtual appliances, enabling administrators to make their virtual infrastructure compliant with baselines they define. Updates that users specify can be applied to guest operating systems as well as to VMware ESX/ESXi hosts, virtual machines and virtual appliances that can be scanned. With Update Manager, users can perform the following tasks:

- Scan for compliance and apply updates for guests, appliances and hosts
- Directly upgrade hosts, virtual machine hardware, VMware Tools and virtual appliances
- Install and update third-party software on hosts

Update Manager 4.1 empowers users to apply offline bundle patches. These are patches that are downloaded manually from a VMware or third-party Web site, not hosted in an online depot. This is especially relevant to VMware ESXi, because many important components, such as third-party driver updates and CIM provider updates, are often distributed only as offline bundles.

An alternative to Update Manager is the vCLI command “vihostupdate.” This command applies software updates to VMware ESX/ESXi images, and installs and updates VMware ESX/ESXi extensions such as VMkernel modules, drivers and CIM providers. Unlike Update Manager, “vihostupdate” works only on an individual host and does not monitor for compliance to baselines. However, “vihostupdate” does not require VMware vCenter Server to function. Table 2 gives a summary of VMware ESXi patching and updating options.

PATCHING AND UPDATING TOOL	WHEN TO USE
VMware vCenter Update Manager	<ul style="list-style-type: none"> • Use when hosts are managed by VMware vCenter Server. Update Manager is integrated with VMware vCenter and provides a single pane of glass. • Use when monitoring for compliance against patching baselines is required. • Use when coordination with host maintenance mode is needed for VMware DRS to perform an orderly evacuation of virtual machines from existing hosts.
"vihostupdate"	<ul style="list-style-type: none"> • Use for one-off host upgrades. • Use in remote situation in which VMware vCenter Server is not accessible. • Use when VMware ESX/ESXi hosts are not managed by VMware vCenter Server.

Table 2. Considerations for Patching and Update Tool

User Authentication

Although day-to-day operations are done on VMware vCenter, there are instances when users must work with ESXi directly, such as with configuration backup and log file access. To control access to the host, you can have local users on a VMware ESXi system. With VMware ESXi 4.1, you can configure the host to join an Active Directory domain, and any user trying to access the host will automatically be authenticated against the centralized user directory. You can also have local users defined and managed on a host-by-host basis and configured using the vSphere client, vCLI or PowerCLI. This second method can be used in place of, or in addition to, the Active Directory integration.

Users can also create local roles, similar to VMware vCenter roles, which define things that the user is authorized to do on the host. For instance, a user can be granted read-only access, which allows them only to view host information. Or they can be granted administrator access, which allows them both to view and to modify host configuration. If the host is integrated with Active Directory, local roles can also be granted to Active Directory users and groups. For example, an Active Directory group can be created to include users who should have an administrator role on a subset of VMware ESXi servers. On those servers, the administrator role can be granted to that Active Directory group. For all other servers, those users would not have an administrator role. If your AD administrator creates a group with the name "VMware ESX Admins," VMware ESXi 4.1 automatically grants administrator access to this group, enabling the creation of a global administrators group. This operation can be overridden on individual VMware ESXi hosts by assigning the "No Access" role to the group "ESX Admins."

The only user defined by default on the system is the root user. The initial root password is typically set using the direct console user interface (DCUI). It can be changed afterward using the vSphere client, vCLI or PowerCLI. The root user is defined only locally. In other words, the root password is not managed by Active Directory. It is possible to exclude the root user access by enabling Lockdown Mode. This is addressed in a later section of this paper.

Logging

Logging is important for both troubleshooting and compliance. VMware ESXi exposes logs from the host agent (hostd), VMware vCenter agent (vpxa) and VMkernel (messages) by using a host syslog capability. Users can configure syslog to write logs onto a file on any datastore accessible to the VMware ESXi host. In VMware ESXi 4.1, the system is automatically configured to write log files to the scratch partition of the host, depending on the type of device used for installation. For installations to local disks, the installer requires a minimum of 5GB available disk space to guarantee that the 4GB scratch partition can be created. For USB/SD or boot-from-SAN installations, we recommend using a shared VMFS volume of 20GB in total, regardless of the cluster size. Monitoring the available disk space on this volume, using the VMware vCenter-provided alarm functionality, is also recommended. Users can also configure syslog to forward log messages to a syslog server for enterprise central logging.

The VMware ESXi log file structure is different from that of ESX. Due to the fact that there is no service console, there is also no need to have the same collection of files. With VMware ESXi, the following log files are used:

PATH + LOG FILE	DESCRIPTION
/var/log/messages	This log file includes the VMkernel, vmkwarning, and hostd logs
/var/log/vmware/hostd.log	Host management service (hostd = Host daemon) log
/var/log/sysboot.log	System boot log
/var/log/vmware/aam/vmware_hostname-xxx.log	VMware HA log file

Table 3. Summary of Log Files

Log files for certain capabilities, such as VMware HA, are not managed through the syslog facility. These log files are stored only on the local VMware ESXi host's in-memory file system. They can be downloaded from the host by using the vSphere client option "Export Diagnostic Data."

It is a best practice to leverage the syslog capabilities that ESXi 4.1 offers. Using a syslog server will simplify troubleshooting and ensure that log files are accessible even when a VMware ESXi host has physically failed. Many syslog servers also enable easier correlation between events.

VMware offers a syslog solution as part of the vSphere Management Assistant, which is designed primarily for VMware ESXi logs. Other alternatives enable you to do advanced event correlation between many types of devices. Configuring the syslog client is straightforward and can be done in seven simple steps:

1. In the **vSphere Client** inventory, left-click the host.
2. Click the **Configuration** tab.
3. Click **Advanced Settings** under **Software**.
4. Select **Syslog** in the tree control.
5. In the **Syslog.Remote.Hostname** text box, enter the name of the remote host where syslog data will be forwarded. If no value is specified, no data is forwarded.
6. In the **Syslog.Remote.Port** text box, enter the port on the remote host where syslog data will be forwarded. By default **Syslog.Remote.Port** is set to 514, the default UDP port used by syslog. Changes to **Syslog.Remote.Port** take effect only if **Syslog.Remote.Hostname** is configured.
7. Click **OK**.

A second capability that VMware ESXi offers is specifying a location for local log files. Local refers to a non-syslog solution. Local log files don't need to be stored on a local drive, but rather on any datastore. It is a best practice for environments without a syslog server to specify a remote VMFS datastore, to ensure that log files will be available when a VMware ESXi host has physically failed to allow for a root cause analysis. This can be configured through the vSphere Client as follows:

1. In the **vSphere Client** inventory, left-click the host.
2. Click the **Configuration** tab.
3. Click **Advanced Settings** under **Software**.
4. Select **Syslog** in the tree control.
5. In the **Syslog.Local.DatastorePath** text box, enter the datastore path to the file where syslog will log messages. If no path is specified, the default path is /var/log/messages. In addition, if pointing at a datastore, ensure that the directory has been created previously.

The datastore path format is `/vmfs/volumes/<datastore>/<folder>/filename`

NOTE: You might need to reboot the host for the changes to take effect.

You can also include the server name in the “folder” name.

Both the syslog advanced setting and the local datastore path setting can be done during a scripted installation through the use of `vim-cmd`. The following command is an example of how to set the local datastore path to a datastore named “vmfs01”; it includes a variable that inserts the hostname into the patch:

```
vim-cmd hostsvc/advopt/update Syslog.Local.DatastorePath string "[vmfs01] /$(hostname -s) /
logfiles/messages"
```

Keeping the VMware ESXi host in sync with an accurate time source is very important for ensuring log accuracy, and it is required for compliance. It is also important if you are using the host to maintain accurate time on the guest virtual machines. However, VMware recommends synchronizing virtual machines with an NTP or w32tm server as described in [VMware knowledge base article 1006427](#) and [VMware knowledge base article 1318](#). VMware ESXi has built-in NTP capabilities for synchronizing with NTP time servers, which can be configured through the vSphere Client or through the shell, as shown in the automated installation script.

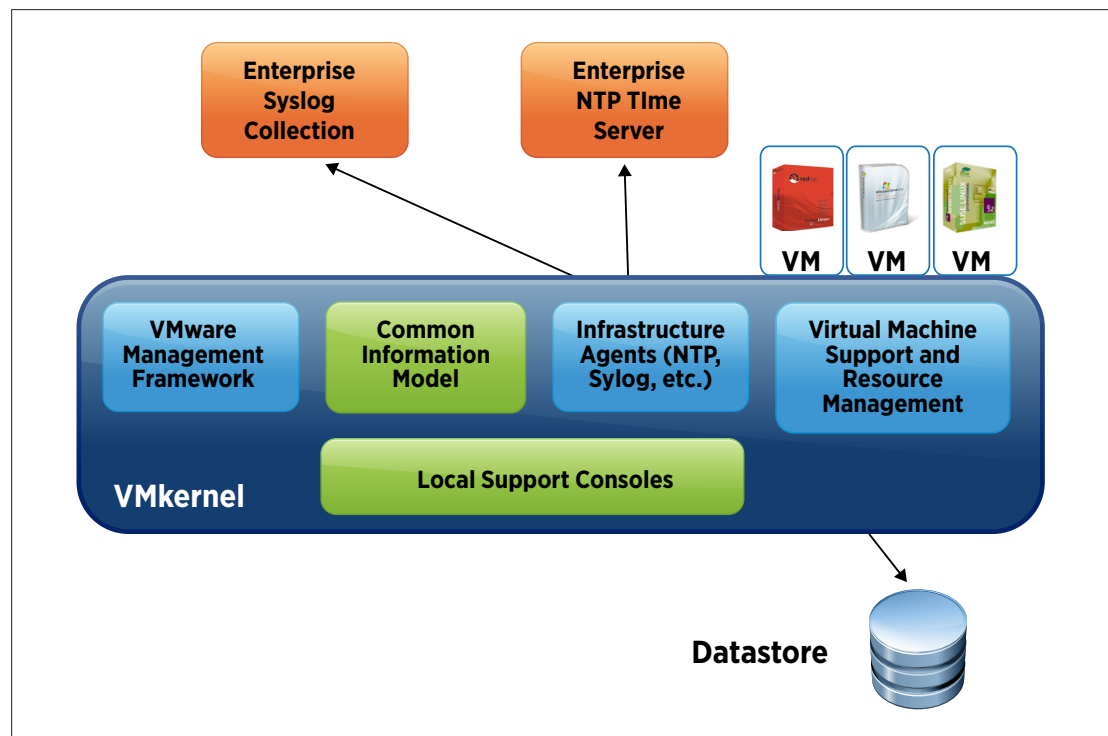


Figure 7. Logging in VMware ESXi

Local Shell Access

Tech Support Mode is a simple shell for advanced technical support. When remote command-line tools are not capable of addressing a particular issue, Tech Support Mode provides an alternative. Similarly to how the COS is used to execute diagnostic commands and fix certain low-level problems, Tech Support Mode enables users to view log and configuration files, as well as to run certain configuration and utility commands to diagnose and fix problems. Tech Support Mode is not based on Linux. Rather, it is a limited-capability shell compiled especially for VMware ESXi.

In VMware ESXi 4.1, Tech Support Mode is fully supported for use by end users and is enhanced in several ways. In addition to being available on the local console of a host, it can also be accessed remotely through SSH. Access to Tech Support Mode is controlled in the following ways:

- Both local and remote Tech Support Mode can be enabled and disabled separately in both the DCUI and vSphere Client.
- Tech Support Mode can be used by any authorized user, not just root users. Users become authorized when they are granted the administrator role on a host (through Active Directory membership in a privileged group and through other methods).
- All commands issued in Tech Support Mode are logged through syslog, providing a full audit trail. If a syslog server is configured, this audit trail is automatically included in the remote logging.
- A timeout can be configured for Tech Support Mode (both local and remote), so that after being enabled, it will automatically be disabled after the configured time.

Tech Support Mode is recommended for use primarily for support, troubleshooting and break-fix situations. It also can be used as part of a scripted installation, as described in a previous section. All other uses of Tech Support Mode, including running custom scripts, are not recommended in most cases.

Diagnostics and Troubleshooting

With VMware ESXi 4.1, there are a variety of options for diagnosing problems with the server configuration or operation, as well as for fixing them. Different methods will be appropriate depending upon the situation. There are also VMware knowledge base articles with instructions on various issues.

The DCUI is the menu-driven interface available at the console of the physical server on which VMware ESXi is installed or embedded. Its main purpose is to perform the initial configuration of the host (IP address, host name, root password) and diagnostics.

The DCUI has several diagnostic menu items:

Restart all management agents, including

- hostd
- vpxa

Reset configuration settings, for example,

- Fix a misconfigured switch
- Reset all configurations to factory defaults

Enable Tech Support Mode (shell access), including

- Local Tech Support Mode
- Remote Tech Support Mode (SSH-based)

Users can also point an ordinary Web browser to the host and view files, including

- Log files
- Configuration files
- Virtual machine files

As an example, we will demonstrate how to view the log files of any given virtual machine. A user with an administrator role must provide credentials to use this feature. The procedure is as follows:

1. Open a browser and enter the URL `http://<vCenter hostname>`, where `<vCenter hostname>` is the IP or fully qualified domain name for the VMware vCenter Server.
2. Provide administrative credentials when prompted.
3. Click the **Browse datastores in the VMware vCenter inventory** link.
4. Navigate the Web pages until you reach the appropriate datacenter, datastore and folder, as noted in step 1.
5. Click the link to the appropriate log file, and open it with your preferred editor.

Tech Support Mode provides another means for more advanced troubleshooting and support, as mentioned earlier. Some new commands added to Tech Support Mode in VMware ESXi 4.1 include

- `vscsiStats`, which provides detailed information on SCSI performance
- `nc`, which is based on the standard netcat utility
- `tcpdump-uw`, which is based on the standard tcpdump utility

Some commands that are used in troubleshooting scenarios are listed as follows for your convenience. This is not a comprehensive list. Rather, these are just a few of the capabilities that VMware ESXi Tech Support Mode offers:

- `vmkping -s 9000 <ipaddress>`

The command `vmkping` can be used to do basic network troubleshooting, but it is more often used to validate the operation of jumbo frames by adding the size of the packet, as shown in our example.

- `fdisk -l`

This enlists all partitions and includes the type of the partition, where VMFS partitions are labeled as "fb."

- `vim-cmd hostsvc/maintenance _mode _enter`

Maintenance mode can be entered from the command line by using `vim-cmd`.

- `vim-cmd hostsvc/maintenance _mode _exit`

Maintenance mode can be exited using this command.

- `vim-cmd vmsvc/getallvms`
- `vim-cmd vmsvc/poweroff <vm id>`

The first command provides a list of all the virtual machines currently registered on the host. The second command enables you to power off a virtual machine.

- `vdf -ph`

This will provide the utilization of the in-memory file system.

These commands are just examples of what is possible with VMware ESXi Tech Support Mode. We recommend that you avoid enabling access to ESXi Tech Support Mode unless absolutely needed, and disabling access once it is no longer needed. In general, troubleshooting workflows are similar to those with VMware ESX, due to the feature set of Tech Support Mode.

The extensive VMware knowledge base should always be your first resource for any problems that are VMware related. For your convenience, we have listed some of the most common issues and most accessed VMware knowledge base articles for VMware ESXi as follows:

1. Restart the management agents on a VMware ESXi host ([1003490](#))
2. Determining why a single virtual machine is inaccessible ([1018834](#))
3. Determining why a virtual machine was powered off or restarted ([1019064](#))
4. Determining why multiple virtual machines are inaccessible ([1019000](#))
5. Troubleshooting virtual machine network connection issues ([1003893](#))
6. Interpreting virtual machine monitor and executable failures ([1019471](#))
7. Determining why a virtual machine does not respond to user interaction at the console ([1017926](#))
8. Using Tech Support Mode in VMware ESXi 4.1 ([1017910](#))
9. Determining why a VMware ESXi host is inaccessible ([1019082](#))
10. Determining why a VMware ESXi host was powered off or restarted ([1019238](#))
11. Determining why a VMware ESXi host does not respond to user interaction ([1017135](#))
12. Enabling serial-line logging for a VMware ESXi host ([1003900](#))
13. Using performance collection tools to gather data for fault analysis ([1006797](#))
14. Using hardware NMI facilities to troubleshoot unresponsive hosts ([1014767](#))
15. Interpreting a VMware ESXi host purple diagnostic screen ([1004250](#))
16. Troubleshooting VMware High Availability (VMware HA) ([1001596](#))

Local Access and Lockdown Mode

VMware ESXi 4.1 provides the ability to fully control all direct access to the host via VMware vCenter Server. Once a host has been joined to VMware vCenter Server, every direct communication interface with the host is configurable as an independent service in the configuration tab for the host in vSphere Client, including

- DCUI
- Local Tech Support Mode
- Remote Tech Support Mode

Each of these can be turned on and off individually.

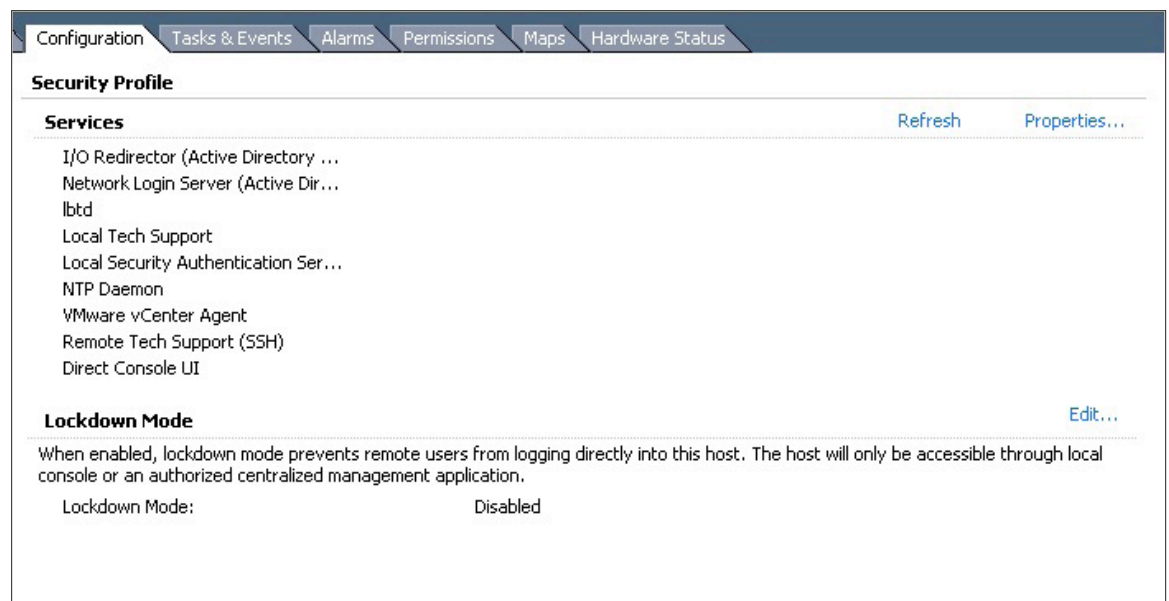


Figure 8. Local Access Services

Access based on the vSphere API—for example, the vSphere Client, PowerCLI, vCLI, and so on—is normally governed by granting local privileges to specific users. The root user is the only one that has a permanent administrator role on the host. All other users must be explicitly granted a local role on the host in order to access it.

There are cases in which you would not want anyone to access the host directly at all, instead managing it exclusively through VMware vCenter Server. Lockdown Mode is a feature designed to provide this capability. When Lockdown Mode is enabled on the host, all direct remote access to the host is blocked, including

- Any vSphere API client
- Local Tech Support Mode
- Remote Tech Support Mode

Even if Tech Support Mode is enabled, Lockdown Mode effectively overrides this by preventing any connection from succeeding. The only way to manage the host remotely is through VMware vCenter Server. The interaction between the host and VMware vCenter Server occurs through a special-purpose account called “vpxuser”; all other ordinary user accounts, including root, can no longer connect remotely.

For the special case of hardware monitoring through the CIM interface, monitoring software must obtain this hardware information directly from the host. To do this, the monitoring software must be programmed to obtain a special authentication ticket from VMware vCenter Server. This ticket allows the software to obtain the information from the host through the VMware vCenter Server “vpxuser” account on a one-time basis.

With Lockdown Mode enabled, the only direct access to the host that remains open is through the DCUI. This provides a way to perform limited administrative tasks outside of VMware vCenter Server. The DCUI can also turn off Lockdown Mode, disabling it without going through VMware vCenter Server. This might be useful if VMware vCenter Server is down or otherwise unavailable, and you wish to revert to direct management of the host. To log in to the DCUI in Lockdown Mode, however, the root password is required. No other user can log in, even if they have been granted an administrator role.

In the extreme case, disabling of all direct access to the host might be desired. For example, you might want to prevent anyone with the root password from disabling Lockdown Mode and managing the host. In this case, you can take the additional step of disabling the DCUI for the host, through VMware vCenter Server. After this is done, no direct interaction with the host, local or remote, is possible. It can be managed only through VMware vCenter Server. If VMware vCenter Server is down or otherwise unavailable, you cannot revert to direct management, because logging in to the DCUI is no longer possible. If the VMware vCenter Server cannot be restored, the only way to revert to direct management is to reinstall the VMware ESXi software on the host.

Lockdown Mode is not permanent. It can be disabled for any individual VMware ESXi host at any time (provided that VMware vCenter Server is running and able to connect to that host). The recommendation is that Lockdown Mode be used in ordinary, day-to-day operations, but that it be disabled for a host if the need arises to interact with it directly. For example, if a troubleshooting situation is encountered and the tools provided by VMware vCenter Server are not sufficient, Lockdown Mode should be disabled and more extensive diagnostics should be performed, using Tech Support Mode, for example.

Table 4 presents a summary of Lockdown Mode and its interaction with the various host access services.

ACCESS MODE	NORMAL	LOCKDOWN	LOCKDOWN + DCUI DISABLED
vSphere API (e.g., vSphere Client, PowerCLI, vCLI, and so on)	Any user, based on local roles/privileges	None (except VMware vCenter “vpxuser”)	None (except VMware vCenter “vpxuser”)
CIM	Any user, based on local roles/privileges	None (except via VMware vCenter ticket)	None (except via VMware vCenter ticket)
DCUI	Root and users with administrator privileges	Root only	None
Tech Support Mode (local)	Root and users with administrator privileges	None	None
Tech Support Mode (remote)	Root and users with administrator privileges	None	None

Table 4. Summary of Lockdown Mode Effect on Local Access

Summary

The following table provides a summary of the tasks traditionally performed in the service console of VMware ESX and the functional equivalents for VMware ESXi.

TASK	VMWARE ESX	VMWARE ESXi
Access local files: VMFS files, configuration files, log files	Console commands to browse datastores and virtual machine files	<ul style="list-style-type: none"> Remote command-line interface commands to list and retrieves files vSphere client datastore browser for VMFS files downloads and uploads files
Manipulate virtual machine files (for example, modify .vmx)	<ul style="list-style-type: none"> Advanced configuration done in the vSphere client Console commands to modify virtual machine files 	<ul style="list-style-type: none"> Advanced configuration done in vSphere Client Remote command-line interface commands to list and retrieves virtual machine files
Backup	<ul style="list-style-type: none"> Virtual machine backup: agents in service console, VMware Data Recovery or third-party backup products VMware ESX backup: uses agents in the service console, creates archive of service console files or performs a scripted reinstall 	<ul style="list-style-type: none"> Virtual machine backup: VMware Data Recovery or third-party backup products VMware ESXi backup: single small backup file created via vCLI command "vicfg-cfgbackup"
Hardware monitoring	<ul style="list-style-type: none"> Agents in service console SNMP 	<ul style="list-style-type: none"> CIM-based framework SNMP
Patching and updating	<ul style="list-style-type: none"> Update Manager RPM-based third-party tools 	<ul style="list-style-type: none"> Update Manager vCLI command "vhostupdate"
Automated deployment	Red Hat Kickstart	<ul style="list-style-type: none"> VMware ESXi scripted installation (analogous to Red Hat Kickstart)
Troubleshooting or support	Local esxcfg-* commands	<ul style="list-style-type: none"> Remote command-line interface commands Tech Support Mode
Advanced configuration	Edits configuration files (for example, hostd.conf) directly	<ul style="list-style-type: none"> Remote command-line interface commands to list and retrieves VMware ESXi configuration files Edits files in Tech Support Mode directly
Logging	Remote syslog in service console	Built-in remote syslog client
Performance monitoring	<ul style="list-style-type: none"> vSphere client "esxtop" in service console 	<ul style="list-style-type: none"> vSphere client vCLI command "resxtop" "esxtop" in Tech Support Mode
Reporting and auditing	<ul style="list-style-type: none"> Service console scripts Log files 	<ul style="list-style-type: none"> Remote command-line interface commands to list and retrieves log files, configuration and settings vSphere Client option to export diagnostic data

Table 5. Comparison of Management Capabilities in VMware ESX and VMware ESXi

VMware ESXi Editions

VMware ESXi architecture is offered as a part of all vSphere product editions, with each successive edition offering greater functionality. At the entry level, VMware offers the vSphere Hypervisor, which is a free virtualization product. Certain VMware ESXi features are limited in this edition, as outlined in Table 6. All other paid editions of vSphere lift these feature restrictions. However, even though the host-level features are not limited in all paid editions, many advanced features, such as VMware DRS and VMware HA, are still only available in higher-license versions.

FEATURE	VSPHERE HYPERVISOR	VMWARE ESXi ENTERPRISE
SNMP monitoring	Not supported	Full functionality
VMware Consolidated Backup (VCB) and VMware Data Recovery (vDR) tool	Not available	Both applications are available
vCLI	Limited to read-only access	Full functionality
PowerCLI and vSphere SDK for Perl	Limited to read-only access	Full functionality

Table 6. Comparison of VMware ESXi Editions

An administrator who has deployed vSphere Hypervisor can enjoy the benefits of virtualization with VMware ESXi within the feature limits. However, the deployment can be upgraded to a more fully featured version of vSphere at any time without having to uninstall or reinstall the VMware ESXi software. The additional capabilities are activated simply when the proper license key is provided, either in the host configuration or in VMware vCenter Server.

References

- *VMware ESXi Configuration Guide*
http://www.vmware.com/pdf/vsphere4/r41/vsp_41_esxi_server_config.pdf
- *VMware ESXi Installable and vCenter Server Setup Guide:*
http://www.vmware.com/pdf/vsphere4/r41/vsp_41_esxi_i_vc_setup_guide.pdf
- *VMware vSphere Command-Line Interface Installation and Scripting Guide*
<http://www.vmware.com/support/developer/vcli/>
- *VMware vSphere Command-Line Interface Reference*
<http://www.vmware.com/support/developer/vcli/>
- *VMware ESXi Upgrade Center*
<http://www.vmware.com/go/UpgradeToESXi>
- *VMware ESXi Chronicles Blog*
<http://blogs.vmware.com/esxi/>
- *William Lam's VMware Scripts and Resources*
<http://www.virtuallyghetto.com/>

