



VMware[®] ESXi[™] 4.1 Migration Guide

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® (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™ (ESXi) architecture represents a similar leap forward in reliability and virtualization management. Less than 5 percent as large as ESX, ESXi runs independently of an operating system 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 4.0, 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 ESXi architecture.

This paper describes the architecture of ESXi and then explains how various management tasks are performed in ESXi. This information can be used to help plan a migration to the ESXi architecture from the legacy ESX framework.

Architecture

In the original 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 into the COS to run configuration and diagnostic commands and scripts.

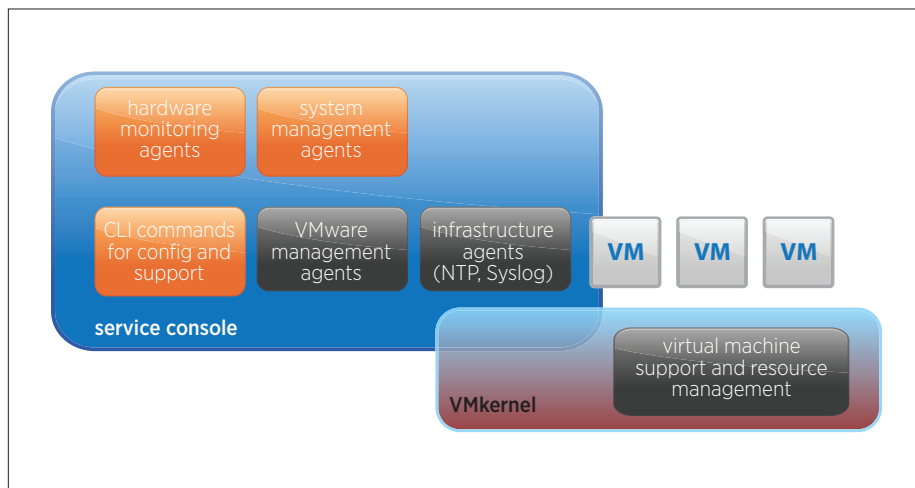


Figure 1. Architecture of ESX

In the 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 ESXi host greatly improves the security of the system.

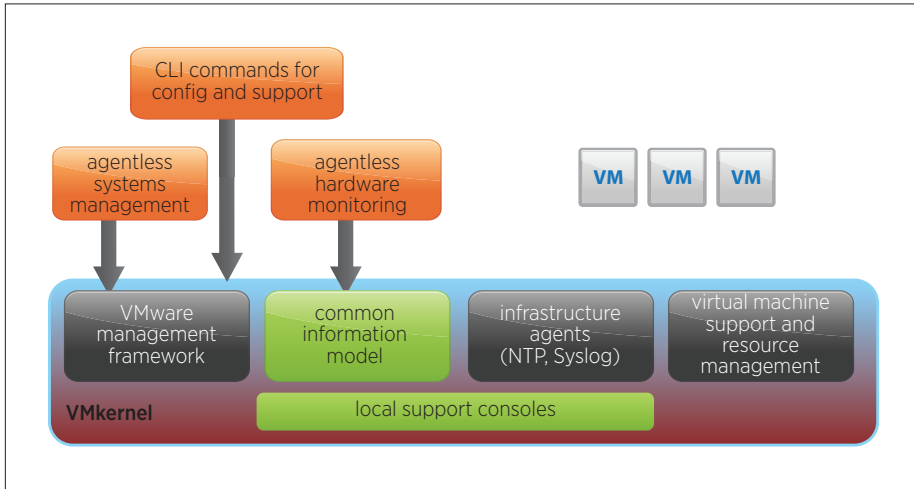


Figure 2. Architecture of ESXi

Management

The management functionality that was provided by agents in the ESX architecture is now exposed via APIs in the ESXi architecture. This allows for an “agent-less” 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 Power CLI (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 ESXi architecture.

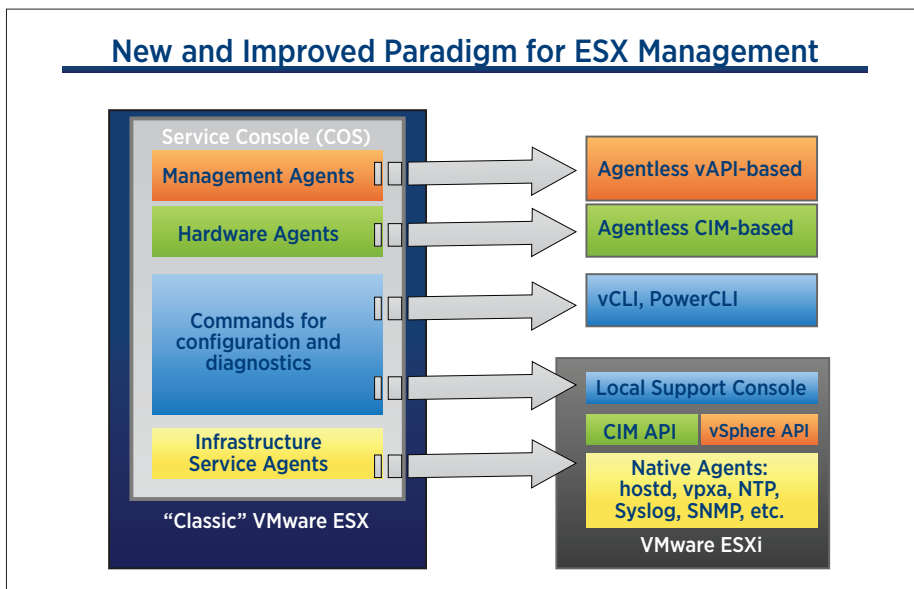


Figure 3. New and Improved Paradigm for ESX Management

Scripting

To automate the management of an ESXi deployment, VMware has created easy-to-use scripting tools for managing day-to-day operations. Users can write scripts with the same functionality as the vSphere client to automate manual tasks, enabling efficient management of small- to large-scale environments. These tools work well with both ESXi and 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, and 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 and maintain ESX and ESXi hosts. There are commands that can completely automate the initial configuration of an ESXi host, and others that provide troubleshooting and diagnostic capabilities. VMware provides vCLI packages for installation on both Windows and Linux systems.

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

NOTE: *Certain commands can be executed only directly on an ESXi host, not through vCenter Server. These are documented in the vSphere Command Line Interface Installation and Scripting Guide.*

The VMware vSphere 4.1 Management Assistant (vMA) is a virtual appliance that brings together all the tools users need to manage vSphere. vMA packages the vCLI, the VMware vSphere 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.

Hardware Monitoring

The Common Information Model (CIM) is an open standard that defines a framework for agentless, standards-based monitoring of hardware resources for 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, ESXi storage infrastructure and virtualization-specific resources. These providers run inside the 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 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 ESXi host in their environment, providing a single view of the physical and virtual health of their systems. Users can also set vCenter alarms to be triggered on certain hardware events, such as temperature or power failure and warning states.

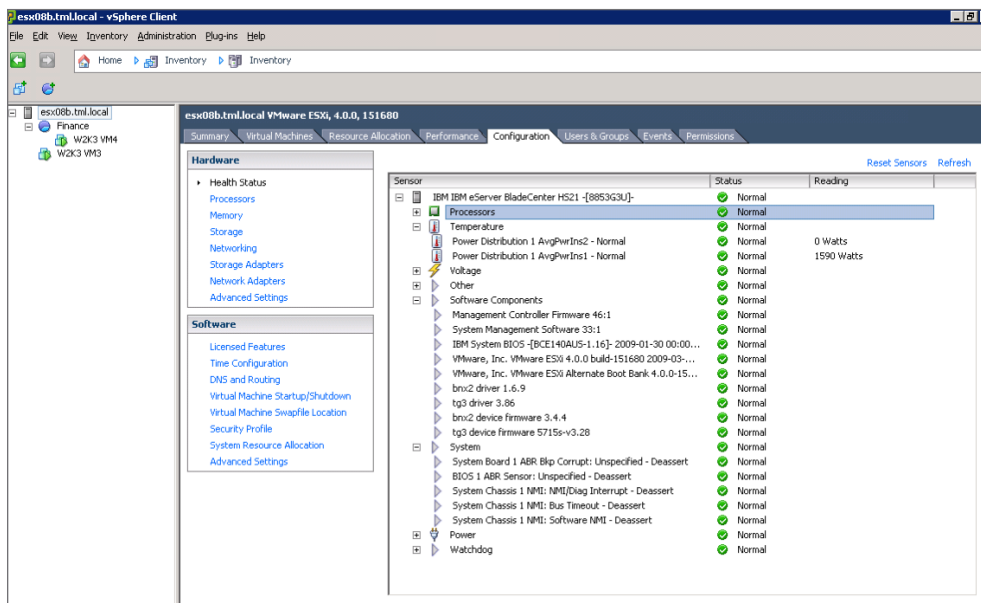


Figure 4. Hardware Monitoring in vCenter Server

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

Systems Management and Backup

Systems management and backup products integrate with 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 ESXi. As a result, the majority of systems management and backup vendors in the VMware ecosystem support ESXi today. Partners such as BMC, CA, HP, IBM, EMC, NetIQ, Quest Software, Commvault, Vizioncore, Double-Take Software, SteelEye and Symantec are among the many partners that have systems management or backup products that support ESXi. If you are using an agent-based partner solution to integrate with ESX, check with your vendor to see if a newer version of the product supports 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 using when they virtualize their infrastructure 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

Patching and updating of ESXi allows flexibility and control. During the patching process, only the specific modules being updated are changed, letting the administrator preserve any previous updates to other components. Whether installed on disk or embedded flash memory, 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 the administrator wishes to revert to the prior one, the host is simply rebooted off the recent good image.

VMware vCenter Update Manager (Update Manager) is a vCenter plug-in patch-management solution for vSphere. Update Manager enables centralized, automated patch and version management for vSphere and offers support for 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 operating systems, as well as to applications on 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 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 ESX/ESXi images, and installs and updates 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 vCenter Server to function. Table 1 gives a summary of ESXi patching and updating options.

PATCHING AND UPDATING TOOL	WHEN TO USE
vCenter Update Manager	<ul style="list-style-type: none"> • Use when hosts are managed by vCenter Server as Update Manager is integrated with vCenter • 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 vCenter Server is not accessible. • Use when ESX/ESXi hosts not managed by vCenter Server.

Table 1.

User Authentication

Although day-to-day operations are done on 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 an ESXi system. With 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 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 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. ESXi 4.1 also automatically grants administrator access to the Active Directory group named “ESX Admins,” which allows the creation of a global administrators group.

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 only defined locally; in other words, the root password is not managed by Active Directory.

Logging

Logging is important for both troubleshooting and compliance. ESXi exposes logs from the host agent (hostd), 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 ESXi host; in ESXi 4.1, the system is automatically configured to write log files to the scratch partition of the host. Users can also configure syslog to forward log messages to a syslog server for enterprise central logging.

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 ESXi host’s in-memory filesystem. They can be downloaded from the host by using the vSphere client option “Export Diagnostic Data.”

Keeping the ESXi host in synch with an accurate time source is very important for ensuring log accuracy, and is required for compliance. It is also important if you are using the host to maintain accurate time on the guest virtual machines. ESXi has built-in NTP capabilities for synchronizing with NTP time servers.

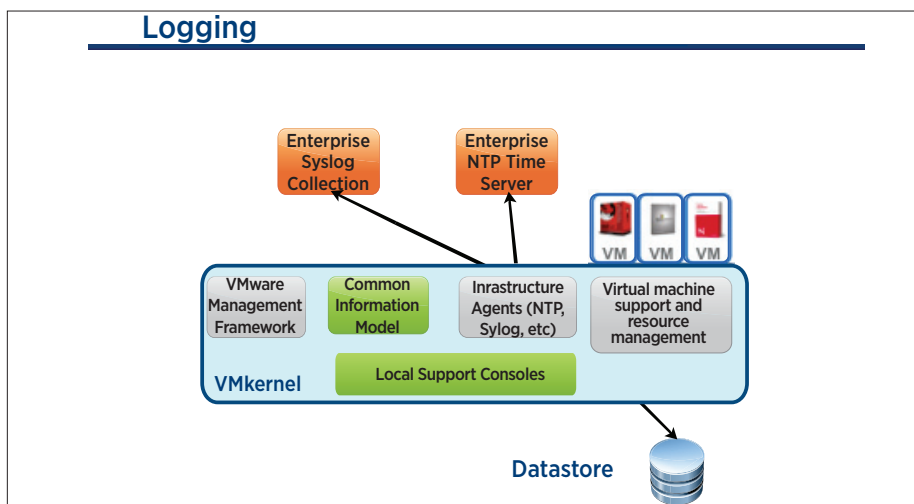


Figure 5. Logging in ESXi

Local Shell Access

Tech support mode is a simple shell for advanced technical support. With situations in which remote scripting tools are not capable of addressing some particular issue, tech support mode provides an alternative. Similar to the way 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 run certain configuration and utility commands in order to diagnose and fix problems. Note that tech support mode is not based on Linux; rather, it is a limited-capability shell compiled especially for ESXi.

In 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 vCenter Server.
- Tech support mode may 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, allowing for a full audit trail. If a syslog server is configured, then 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 the next section. All other uses of tech support mode, including running custom scripts, are not recommended for most cases.

Deployment

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

- Pre-install
- Post-install
- First-boot

These scripts are run locally on the ESXi host and can perform various tasks, such as configuring the host's virtual networking and joining it to 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 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.

ESXi 4.1 is still available pre-installed on flash drives on certain server models available from a number of hardware OEM vendors. Scripted installations can also be used to deploy ESXi to a supported USB or flash drive on a server. (Please consult the server HCL to determine which combinations of server and USB or flash drive are supported.)

Diagnostics and Troubleshooting

With 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 more appropriate depending upon the situation, and VMware issues 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 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 vNetwork Distributed 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

Credentials of a user with an administrator role must be provided in the browser in order to use this feature.

Finally, tech support mode provides another means for more advanced troubleshooting and support, as mentioned earlier. Some new commands added to tech support mode in 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

Local Access and Lockdown Mode

ESXi 4.1 provides the ability to fully control all direct access to the host via vCenter Server. Once a host has been joined to 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.

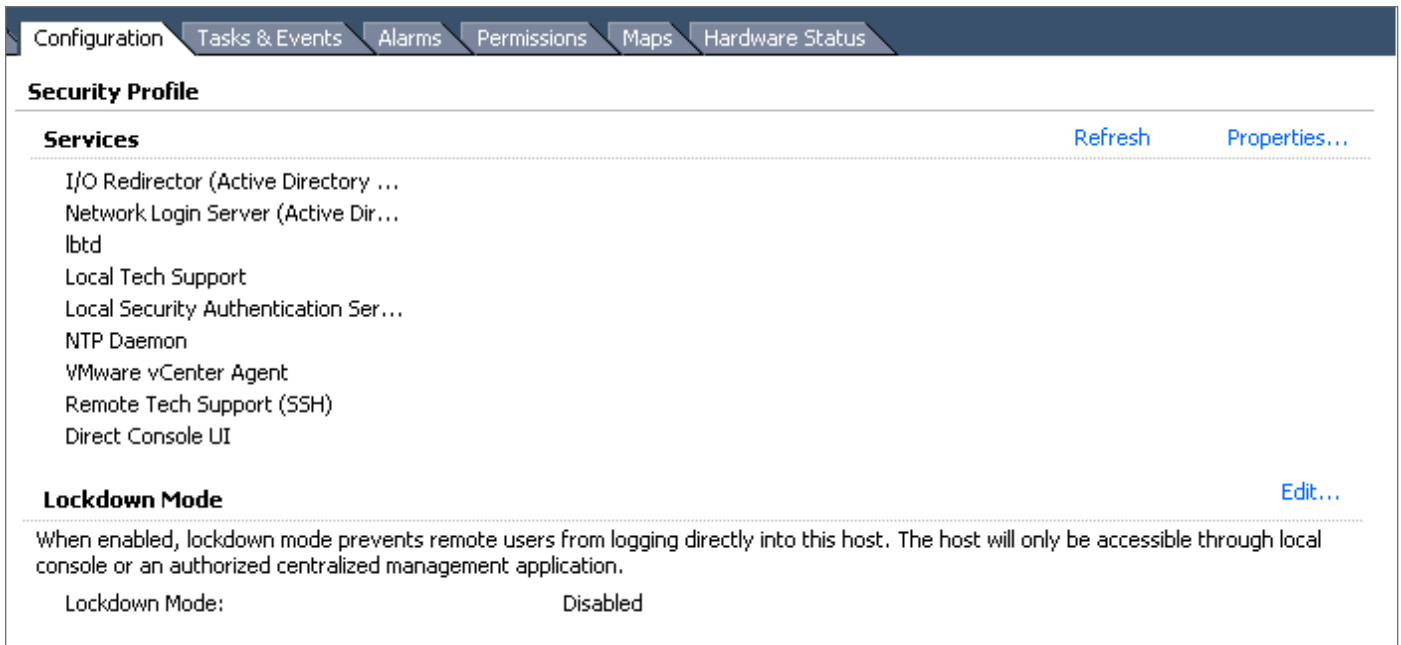


Figure 6. Local Access Services

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 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 vCenter Server. The interaction between the host and 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. In order to do this, the monitoring software must be programmed to obtain a special authentication ticket from vCenter Server. This ticket allows the software to obtain the information from the host through the 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 vCenter Server. In addition, the DCUI can also turn off lockdown mode, disabling it without going through vCenter Server. This might be useful if vCenter Server is down or otherwise unavailable, and you wish to revert to direct management of the host. In order 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 may 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 vCenter Server. After this is done, no direct interaction with the host, local or remote, is possible. It can be managed only through vCenter Server. If vCenter Server is down or otherwise unavailable, you cannot revert to direct management, because logging into the DCUI is no longer possible. If the vCenter Server cannot be restored, then the only way to revert to direct management is to reinstall the ESXi software on the host.

Note that lockdown mode is not permanent; it can be disabled for any individual ESXi host at any time (provided that 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 vCenter Server are not sufficient, then lockdown mode should be disabled and more extensive diagnostics should be performed, using tech support mode, for example.

Table 2 presents a summary of lockdown mode and its interaction with the various host access services.

ACCESS MODE	NORMAL	LOCK DOWN	LOCK DOWN + DCUI DISABLED
vSphere API (e.g., vSphere client, PowerCLI, vCLI, etc)	Any user, based on local roles/privileges	None (except vCenter "vpxuser")	None (except vCenter "vpxuser")
CIM	Any user, based on local role/privilege	None (except via vCenter ticket)	None (except via vCenter ticket)
DCUI	Root and users with admin privileges	Root only	None
Tech support mode (local)	Root and users with admin privileges	None	None
Tech support mode (remote)	Root and users with admin privileges	None	None

Table 2. Summary of Lockdown Mode Effect on Local Access

Summary

Table 3 provides a summary of the tasks traditionally performed in the service console of ESX and the functional equivalents for ESXi.

TASK	ESX	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 interfaces commands to list and retrieve 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 interfaces commands to list and retrieve 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 • 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 • ESXi backup: single small backup file created via vCLI command "vicfgcgbbackup"
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	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	Editing configuration files (for example, hostd.conf) directly	<ul style="list-style-type: none"> • Remote command-line interfaces commands to list and retrieve ESXi configuration files • Editing 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 • "esxstop" in service console 	<ul style="list-style-type: none"> • vSphere client • vCLI command "resxstop" • "esxstop" 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 interfaces commands to list and retrieve log files, configuration and settings • vSphere client option to export diagnostic data

Table 3. Comparison of Management Capabilities in ESX and ESXi

ESXi Editions

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 ESXi features are limited in this edition, as outlined in Table 4. 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	ENTERPRISE ESXi
SNMP monitoring	Not supported	Full functionality
VMware Consolidated Backup (VCB) and Data Recovery (DR) tool	Not available	Both applications are available
vCLI	Limited to read-only access	Full functionality
PowerCLI and SDK for Perl	Limited to read-only access	Full functionality

Table 4. Comparison of ESX Editions

An administrator who has deployed vSphere Hypervisor can enjoy the benefits of virtualization with 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 ESXi software. The additional capabilities are simply activated when the proper license key is provided, either in the host configuration or in vCenter Server.

References

- *ESXi Configuration Guide*: http://www.vmware.com/pdf/vsphere4/r41/vsp_41_esxi_server_config.pdf
- *vSphere Command-Line Interface Installation and Scripting Guide*: <http://www.vmware.com/support/developer/vcli/>
- *vSphere Command-Line Interface Reference*: <http://www.vmware.com/support/developer/vcli/>
- *ESXi Upgrade Center*: <http://www.vmware.com/UpgradeToESXi>.

