

Main Documentation Set (ESX not ESXi)

Introduction to VMware vSphere

- **vCompute** – aggregate resources
- **vStorage** – enables the most efficient use and management of storage
- **vNetwork** – simplify and enhance networking
- **Fault Tolerance** – a secondary copy. Actions completed on the primary VM are also applied to the secondary VM. If the primary VM becomes unavailable, the secondary machine becomes active, providing continual availability.
- **Distributed Virtual Switch (DVS)** – spans many hosts reduction of maintenance and increasing network capacity.
- **Host Profiles** – host configuration management through user-defined configuration policies. Captures the blueprint and monitors compliance.
- **Pluggable Storage Array (PSA)** – greater array certification flexibility and improved array-optimized performance. A multipath I/O framework.
- **Cluster** – aggregate computing and memory resources of a group of physical x86 servers sharing the same network and storage arrays.
- **Resource pools** – partitions of computing and memory resources from a single host or a cluster. can be hierarchical and nested.
- **Storage VMotion** – enables the migration of VMs from one datastore to another datastore without service interruption
- **Fault Tolerance (FT)** – uses vLockstep technology, continuous availability by protecting a VM (the Primary VM) with a shadow copy (Secondary VM) that runs in virtual lockstep on a separate host.
- **vNetwork Distributed Switch (dvSwitch)** – functions as a single virtual switch across all associated hosts.
- **dvPort (distributed virtual port)** – a port on a DVS that connects to a host's service console or VMkernel or to a VM's network adapter.
- **vApp** – has the same basic operation as a VM, but can contain multiple VMs or appliances.
- Web Access cannot be used to access a host running ESXi 4.0.
- Web Access is turned off by default for ESX hosts.

Getting Started with ESX

- ESX Service Console is held in the esxconsole.vmdk partition.
- vCenter Server uses Microsoft SQL Server 2005 Express for small deployments with up to 5 hosts and 50 VMs.
- If SQL Native Client is already installed, uninstall SQL Native Client before you begin the vCenter Server installation.
- vCenter server must belong to a domain rather than a workgroup. Otherwise it's not able to discover – using such features as vCenter Guided Consolidation Service. The computer name cannot be more than 15 characters.
- vCenter Server should not be an Active Directory domain controller.
- The Domain user account should have the following permissions:
 - Member of the Administrators group
 - Act as part of the operating system
 - Log on as a service

ESX and vCenter server installation guide

- ESX4 will only install and run on servers with 64-bit x86 CPUs. They require a 2GB RAM minimum.
- vCenter Server must have 2 CPUs and 3GB RAM.
- The vCenter Server has a service called VMware VirtualCenter Management Webservices. This service requires 128MB to 1.5GB of additional memory.
- You can use a 32-bit Windows for up 200 hosts. A 64-bit Windows can have 200-300 hosts.
- The vSphere Client requires the Microsoft .NET 3.0 SP1 Framework.
- vCenter server required firewall ports:

Port	Description
80	Redirects requests to HTTPS port 443.
389	LDAP port number for the Directory Services for the vCenter Server group. Needs to bind to port 389, even if you are not joining this vCenter Server instance to a Linked Mode group.
443	Listen for connections from the vSphere Client.
636	For vCenter Linked Mode, this is the SSL port of the local instance
902	Uses to send data to managed hosts. Managed hosts also send a regular heartbeat over UDP port 902 to the vCenter
902/903	Ports used by vSphere Client to display VM consoles.
8080	Web Services HTTP. Used for the VMware VirtualCenter Management Webservices.
8443	Web Services HTTPS. Used for the VMware VirtualCenter Management Webservices.

- If ESX will not use an NTP server, make sure that the server hardware clock is set to UTC
- IPv6 is not supported for ESX installation
- The service console must be installed on a VMFS datastore that is resident on a host's local disk or on a SAN disk that is masked and zoned to that particular host only.
- The evaluation period is 60 days and begins as soon as you power on the ESX machine, even if you start in license mode initially.

- The installer creates three basic partitions: /boot, vmkcore and VMFS. The service console VMDK file contains swap, and /var/log, by default, and any other partitions that you specify.
- The media depot is a network-accessible location that contains the ESX installation media. You can use HTTP/ HTTPS, FTP, or NFS to access the depot.
- Scripted installation – you must point to the media depot in the script by including the install command with the nfs or url option.
- Interactive installation – include the askmedia boot option.
- The boot options list appears when you boot the installer and press F2.
- Bootstrap Commands for ESX Installation

Command	Description
askmedia	Allows you to interactively select the location of the ESX installation media. This option is required if the image is hosted at an HTTP, FTP, or NFS location.
BOOTIF	Accepts the format for the boot network adapter as supplied by PXELINUX.
gateway=<ip address>	Sets this network gateway as the default gateway during the install.
ip=<ip address>	Specifies a static IP address to be used for downloading the script and the installation media. The IPAPPEND option is also supported if you PXE boot the installer.
ks=cdrom:</path>	Performs a scripted installation with the script at <path>, which resides on the DVD in the DVD-ROM drive.
ks=file:</path>	Performs a scripted installation with the script at <path>, which resides inside the initial ramdisk image.
ks=ftp:</server>/<path>/	Performs a scripted installation with a script located at the given URL.
ks=http:</server>/<path>	Performs a scripted installation with a script located at the given URL.
ks=https:</server>/<path>	Performs a scripted installation with a script located at the given URL.
ks=nfs:</server>/<path>	Performs a scripted installation with the script located at <path> on a given NFS server.
ks=usb	Performs a scripted installation with the ks.cfg script in the root directory of the USB flash drive attached to the host. If multiple flash drives are attached, the installer cycles through each one, mounting and unmounting them until the file named ks.cfg is found.
ks=UUID:<partition-UUID>:</path>	Performs a scripted installation with a script located on the ext partition with the given UUID.
ksdevice=<device>	Same as netdevice
nameserver=<ip address>	Specifies a domain name server as the nameserver during the install.
netdevice=<device>	Tries to use a network adapter <device> when looking for an installation script and installation media. Specify as a MAC address (for example, 00:50:56:C0:00:01). If not specified and files need to be retrieved over the network, the installer defaults to the first discovered network adapter. The IPAPPEND option is also supported if you PXE boot the installer.
netmask=<subnet mask>	Specifies subnet mask for the network interface that downloads the installation media.
noapic	Flags the kernel to use the XTPIC instead of the APIC.
text	Starts the ESX installer in text mode.
url=<url>	Looks for the installation media at the specified URL. When you are PXE booting the installer, the url= command only works with earlier versions of SYSLINUX. The command does not work with SYSLINUX/PXELINUX version 3.70 and higher.
vlanid=<vlanid>	Configures the VLAN for the network card.

- PXE Boot the ESX Installer:
 1. Install TFTP server software that supports PXE booting.
 2. Put menu.c32 file in an accessible place
 3. Install PXELINUX.
 4. Configure the DHCP server.
 5. Create the kernel image and ramdisk directory by copying the vmlinuz and initrd.img files from the /isolinux directory on the ESX installation DVD to a supported location.
 6. Create the /tftboot/pxelinux.cfg directory on your TFTP server.
 7. Create a PXE configuration file. PXE configuration file in /tftboot/pxelinux.cfg
- In an interactive installation, omit the ks= option.
- ESX 3.x supported a hybrid installation. You could supply an incomplete ESX installation script, and the installer prompts you for the missing parts. ESX 4.0 does not support this.
- Install ESX interactively or by using a script. For interactive installation, you can use graphical mode or text mode.
- The installer erases all content on the selected storage device.
- Installing ESX on a USB device is not supported.
- VMFS2 volumes are not recognized by ESX 4.0.
- The installation log is /var/log/esx_install.log.
- The installation script can reside in one of the following locations:
 - Default installation script
 - HTTP/HTTPS
 - USB flash drive
 - FTP
 - NFS
 - Local disk
- The installer creates a /root/ks.cfg script, which reflects the choices you made in the interactive installation.
- Installation media contains the following default installation scripts:
 - **ks-first-safe.cfg** Installs ESX on the first detected disk and preserves the VMFS datastores on the disk.

- **ks-first.cfg** Installs ESX on the first detected disk.
- The default root password is mypassword.
- /boot and vmkcore are physical partitions. /, swap, /var/log, and all the optional partitions are stored on a virtual disk called *esxconsole-
<system-uuid>/esxconsole.vmdk*. The virtual disk is stored in a VMFS volume.
- You cannot define the sizes of the /boot, vmkcore, and /vmfs partitions when you use the graphical or text installation modes. You can define these partition sizes when you do a scripted installation.
- ESX Required Partitions

Mount Point	Type	Size	Location
/boot	ext3	1.25GB of free space and includes the /boot and vmkcore partitions. The /boot partition alone requires 1100MB.	Physical partition
N/A	swap	600MB recommended minimum 1600MB maximum.	Virtual disk in a VMFS volume
/	ext3	Based on the size of the /usr partition. By default, the minimum size is 5GB and no /usr partition is defined.	Virtual disk in a VMFS volume
N/A	VMFS3	For VMFS volumes hosting esxconsole.vmdk: 1200MB and an additional 10GB. VMFS2 is supported in read-only mode to import legacy VMs.	Physical partition.
N/A	vmkcore	See /boot	Physical partition

- ESX Optional Partitions

Mount Point	Type	Recommended Size	Location
/home	ext3	512MB	Virtual disk in a VMFS volume
/tmp	ext3	1024MB	Virtual disk
/usr	ext3	Missing in PDF	Virtual disk
/var/log	ext3	2000MB	Virtual disk. The graphical and text installers create this partition by default.

- vihostupdate command applies software updates to ESX4/ESXi4 hosts and installs and updates ESX/ESXi extensions (use vihostupdate35 on ESX 3.5/ESXi 3.5 hosts.)
- The esxupdate utility is for ESX only.
- You can use the vihostupdate utility in conjunction with offline bundles or with a depot
- vSphere Databases:
 - Microsoft SQL Server 2005 Express – up to 5 hosts and 50 VMs. If the machine has Microsoft SQL Native Client installed, remove it before installing vCenter Server with the bundled database. If the machine has MSXML Core Services 6.0 installed, remove it before installing
 - Microsoft SQL Server 2005 – Windows XP, apply MDAC 2.8 SP1 to the client. Use the SQL Native Client driver (version 9.x) for the client. Ensure that the machine has a valid ODBC DSN entry. Remove MSXML Core Services 6.0 before
 - Microsoft SQL Server 2008 – Windows XP, apply MDAC 2.8 SP1 to the client. Use the SQL Native Client driver (version 10.x) for the client. Ensure that the machine has a valid ODBC DSN entry.
 - Oracle 10g – If necessary, first apply patch 10.2.0.3 (or later) to the client and server. Then apply patch 5699495 to the client. Ensure that the machine has a valid ODBC DSN entry.
 - Oracle 11g – Ensure that the machine has a valid ODBC DSN entry.
- Even though vCenter Server is supported on 64-bit operating systems, the vCenter Server system must have a 32-bit DSN. This requirement applies to all supported databases. By default, any DSN created on a 64-bit system is 64 bit. On a 64-bit system use C:\WINDOWS\SYSWOW64\odbc32.exe.
- vCenter Server must have a computer name that is 15 characters or fewer. The data source name (DSN) and remote database systems can have names with more than 15 characters.
- To prepare a SQL Server database to work with vCenter Server, you generally need to create a SQL Server database user with database operator (DBO) rights.
- If you use SQL Server for vCenter Server, do not use the master database.
- When using Microsoft SQL Server 2008 Standard Edition with vCenter Server, do not name the instance MSSQLSERVER.
- The vCenter Server performs a silent installation of vCenter Orchestrator. If you install vCenter Server on an IPv6 operating system, the vCenter Orchestrator module is not supported.
- The vSphere Host Update Utility is for updating and patching ESXi 4.0 hosts and upgrading ESX 3.x/ESXi 3.5 hosts to ESX 4.0/ESXi 4.0.
- You can join multiple vCenter Server systems to form a Linked Mode group.
- Linked Mode global data includes:
 - Connection information (IP and ports)
 - Certificates
 - Licensing information
 - User roles
- vCenter Server instances in a Linked Mode group can be in different domains if the domains have a two-way trust relationship.

- The installer must be run by a domain user who is an administrator on both the machines.
- The vCenter Server installer validates that the machine clocks are not more than 5 minutes apart.
- The instances can run under different domain accounts.
- Windows Server 2008 automatically configures the firewall to permit access. Launch firewall.cpl and add an exception for C:\Windows\ADAM\dsamain.exe
- When you install vCenter Server in Linked Mode, the firewall configuration on any network-based firewalls must be modified.
- Configure Windows RPC ports to generically allow selective ports for machine-to-machine RPC communication.
- vCenter Collector service – uses port 8181 and 8182, by default.
- vCenter Web Server – uses ports 8080 and 8443, by default.
- All product licenses are encapsulated in 25-character license keys that you can manage and monitor from vCenter Server.
- vSphere Licenses – For ESX/ESXi.
- vCenter Server Licenses – For vCenter Server.
- vCenter Server 4.0 does not require a license server to manage ESX 4.0/ESXi 4.0 hosts. However, vCenter Server 4.0 does require a license server to manage ESX 3.x/ESXi 3.5 hosts.
- If a vCenter Server license expires, the managed hosts become disconnected.
- If an ESX/ESXi host license expires, the VMs that reside on the host continue to run, but you cannot power on the VMs or reset them.

vSphere upgrade guide

- When you upgrade from ESX 3.x/ESXi 3.5 to ESX 4.0/ESXi 4.0, you can use either the vSphere Host Update Utility or vCenter Update Manager.
- **EDIT** – you can also use the upgrade script (esxupgrade.sh – KB 1009440). VMware doesn't mention the option of a fresh install at any point; they are pushing the upgrade option as much as possible. You can't seem to do an upgrade with the ESX4 install CD.
- Host Update Utility is used for upgrading ESX 3.x/ESXi 3.5 standalone hosts to ESX 4.0/ESXi 4.0 and for patching ESXi 4.0 standalone hosts. It is a standalone Microsoft Windows application recommended for smaller deployments with fewer than 10 ESX/ESXi hosts, without vCenter Server or Update Manager.
- vCenter Update Manager is for upgrading and updating ESX/ESXi hosts that are managed in vCenter Server.
- Orchestrated upgrades can be used to upgrade the VM hardware and VMware Tools.
- No VMFS upgrade is required if you are upgrading from ESX 3.x
- You must upgrade VMware Tools before upgrading virtual hardware.
- After you upgrade to vCenter Server, you cannot revert to VirtualCenter 2.x. Take appropriate backups before starting the upgrade.
- Upgrade VirtualCenter 2.x to vCenter Server 4.0:
 - Make sure your database is compatible
 - Have required permissions
 - Take a full backup of the VirtualCenter 2.x database
 - Back up the VirtualCenter 2.x SSL certificates
 - Install the vSphere Client
 - vCenter Converter, upgrade
 - vCenter Guided Consolidation upgrade
 - Upgrade to vCenter Update Manager 4.0.
 - Use vCenter Update Manager to upgrade ESX 3.x hosts to ESX 4.0. (or use host update utility)
 - Use vCenter Update Manager to upgrade your VMs. vCenter Update Manager ensures that the VMware Tools upgrade and the virtual hardware upgrade happen in the correct order
 - Upgrade your product licenses
- ESX 2.x hosts cannot be added to clusters.
- Oracle 9i is no longer supported.
- Microsoft SQL Server 2000 is no longer supported.
- To Back Up VirtualCenter 2.x:
 - Make a full backup of the VirtualCenter 2.x database.
 - Back up the VirtualCenter 2.x SSL certificates. %ALLUSERSPROFILE%\Application Data\VMware\VMware VirtualCenter or %ALLUSERSPROFILE%\VMware\VMware VirtualCenter\.
 - Note any non-default settings, such as the IP address, the database DSN, user name, password, and assigned ports.
 - Create a backup copy of vpxd.cfg.
- The database upgrade log is at %TEMP%\VCDatabaseUpgrade.log.
- By default, a vCenter Server creates a maximum of 10 simultaneous database connections.
- Datastores and networks have their own set of privileges that control access to them.
- Users are initially granted the No Access role on all new managed objects, including datastores and networks. All existing objects in vCenter maintain their permissions after the upgrade.
- The upgrade process uses the datacenter's "Read-only" privilege.
 - If the "Read-only" privilege is nonpropagating (not inherited by child objects), VMware assumes access privileges should not be assigned to datastores and networks.

- If the “Read-only” privilege is propagating (inherited by child objects), VMware assumes access privileges should be assigned to datastores and networks so users can view them and perform basic operations that require access.
- The “Read-only” propagating permission on a datacenter, as well as all other permissions you have set, will continue to work as expected after the upgrade.
- You must change “Read-only” nonpropagating datastore permissions to propagating datastore permissions in order for users to access the datastore. You can assign datastore permissions on datastore or folders containing datastore. The same goes for Network permissions.
- The “Database Consumer” sample role assigns the “Allocate Space” privilege to the datastore. “Network Consumer” sample role assigns the “Assign Network” privilege.
- Host Update Utility does not upgrade VMFS datastores or VM
- Update Manager supports mass remediation.
- During host upgrades using the Update Manager, static IP addresses are a requirement.
- The upgrade to ESXi4 & ESX4 preserves almost all configuration data, including your networking, security, and storage configuration. The only configuration not preserved is related to licensing.
- For ESX only, the upgrade reuses the existing /boot partition to hold the ESX 4.0 boot files. After the upgrade, the ESX 3.x installation is mounted in the new ESX 4.0 installation under the /esx3-installation directory.
- Backup the ESX Host Configuration:
 - Back up the files in the /etc/passwd, /etc/groups, /etc/shadow, and /etc/gshadow directories.
 - Back up any custom scripts.
 - Back up your .vmx files.
 - Back up local images, such as templates, exported VMs, and .iso files.
- Backup the ESXi Host Configuration: `vicfg-cfgbackup --server <ESXi-host-ip> --portnumber <port_number> --protocol <protocol_type> --username username --password <password> -s <backup-filename>`
- You cannot perform an in-place upgrade from ESX to ESXi (or from ESXi to ESX).
- The only ESX 2.x version that has upgrade support is ESX 2.5.5 – you can perform a migration upgrade with or without VMotion.
- Direct, in-place upgrade from ESX 2.5.5 to ESX 4.0 is not supported, even if you upgrade to ESX 3.x as an intermediary step. The default ESX 2.5.5 installation creates a /boot partition that is too small to enable upgrades to ESX 4.0. As an exception, if you have a non-default ESX 2.5.5 installation on which at least 100MB of space is available on the /boot partition, you can upgrade ESX 2.5.5 to ESX 3.x and then to ESX 4.0.
- The upgrade logs are here:
 - /esx3-installation/esx4-upgrade/
 - /var/log/vmware/
- For an unsuccessful ESX upgrade:
 - /esx4-upgrade/
 - /var/log/vmware/
- You may need to reconnect the host to vCenter Server and assign an upgraded license to your product within 60 days after the upgrade.
- The host sdX devices might be renumbered after the upgrade.
- You must convert LUN masking to the claim rule format. Run the `esxcli corestorage claimrule convert` command. It converts the /adv/Disk/MaskLUNs advanced configuration entry in esx.conf to claim rules with MASK_PATH as the plug-in.
- Web Access service is disabled after you upgrade the host.
- vSphere Web Access is not supported on ESXi hosts.
- 60-day evaluation count down starts even if the host is licensed and you are not using evaluation mode.
- One advantage of using evaluation mode is that it offers full feature functionality.
- After you determine that the ESX 4.0 upgrade is stable, you can remove the ESX 3.x boot option from the ESX 4.0 bootloader menu to disable the ability to roll back to ESX 3.x. If you deselected the rollback option, this procedure is not applicable. Run the `cleanup-esx3` command with the optional -f (force) flag.
- The cleanup-esx3 script removes the following files and references from the ESX 4.0 host:
 - ESX 3.x references in the /etc/fstab directory
 - ESX 3.x boot files in the /boot directory
 - The rollback-to-esx3 script in the /usr/sbin/ directory
- After you upgrade all of your hosts to ESX4/ESXi4, you can optionally uninstall your license server and remove the license server configuration from vCenter Server.
- For ESXi, the vSphere Host Update Utility does not support roll back.
- You can remove the ESX 4.0 boot option from the ESX bootloader menu and perform a complete roll back to ESX 3.x. Any changes made to the ESX 4.0 service console are lost after the rollback.
 - Any changes made to VMs will persist after the rollback.
 - If you upgraded the VM hardware, the VMs will not work after you perform the ESX rollback. To avoid this situation, take a snapshot of the VM before you upgrade the VM hardware.
 - Run the rollback-to-esx3 command – reconfigures the bootloader.
 - Reboot the server.
 - After the host boots into ESX 3.x, delete the ESX 4.0 service console VMDK folder from the VMFS datastore.
- Each time you update an ESXi host, a copy of the ESXi build is saved on your host.

- ESXi permits only one level of rollback.
- When you manually boot into the standby build instead of the current boot build, an irreversible rollback occurs.
- When the page that displays the current boot build appears, press Shift+r to select the standby build.
- Do not use vmware-vmupgrade.exe to upgrade VMs.
- Some VMs that you create on ESX4 hosts are supported on ESX 3.x/ESXi 3.5 hosts. If you create a VM on ESX4 and select the typical path, the virtual hardware version is version 7. Virtual machines with virtual hardware version 7 are not supported on ESX 3.x/ESXi 3.5 hosts. Select the custom path and select virtual hardware version 4
- If you create VMs that use paravirtualization (VMI) or an enhanced networking device (vmxnet), VMotion is not supported. In this case, you can only move the VM to the ESX 3.x host if the VM is powered off.
- On Windows guest operating systems, you must reboot the VM a total of three times when you upgrade VMware Tools and the virtual hardware:
 1. Power on the VM.
 2. Upgrade VMware Tools.
 3. Reboot the VM at the end of the VMware Tools upgrade.
 4. Power off the VM.
 5. Upgrade the virtual Hardware.
 6. Power on the VM.
 7. The Windows operating system detects new devices and prompts you to reboot the VM.
 8. Reboot the VM to make the devices work properly.
- During the virtual hardware upgrade, the VM must be shut down for all guest operating systems.
- When you upgrade from virtual hardware version 3 to version 7, the upgrade is irreversible, even if you take a VM backup or snapshot before performing the upgrade. 4 to version 7 the upgrade is reversible if you take a VM backup or snapshot
- Tools for installing updates and patches to ESX/ESXi hosts:
 - Host Update Utility – graphical utility for ESXi only.
 - Update Manager – for ESX and ESXi,
 - vihostupdate – command-line utility for ESX and ESXi.
 - esxupdate – command-line utility for ESX only.
- If the Host Update Utility does not meet the needs of your environment, you can customize the application the settings.config XML file, located in %PROGRAMFILES%\VMware\Infrastructure\VIUpdate 4.0
- Run vihostupdate on ESX 4.0/ESXi 4.0 hosts. Run vihostupdate35 on ESX 3.5/ESXi 3.5 hosts.
- The esxupdate utility is supported as well. It is for ESX only.
- The vihostupdate command works with bulletins.
- You can use the vihostupdate utility in conjunction with offline bundles or with a depot.
- To get the build number of an ESX4 host: # vmware -l (from KB 1001179).

vSphere Basic System Administration

- VMware modules (plugins) include:
 - Update Manager
 - Converter Enterprise
 - vShield Zones – an application-aware firewall built for vCenter Server integration. It inspects client-server communications and inter-virtual-machine communication to provide detailed traffic analytics and application-aware firewall partitioning.
 - Orchestrator – a workflow engine that enables you to create and execute automated workflows.
 - Data Recovery – a disk-based backup and recovery solution. Enable centralized and efficient management of backup jobs and includes data de-duplication.
- Tomcat Web server is installed as part of the vCenter Server installation. The components that require Tomcat are:
 - Linked Mode
 - CIM/Hardware Status tab
 - Performance charts
 - WebAccess
 - vCenter Storage Monitoring/Storage Views tab
 - vCenter Service Status
- When a server is connected to other vCenter Server systems using Linked Mode, you can connect to that vCenter Server system and view and manage the inventories of all the vCenter Server systems that are linked.
- Linked Mode uses Microsoft Active Directory Application Mode (ADAM) to store and synchronize data.
- The following requirements apply to each vCenter Server system that is a member of a Linked Mode group:
 - DNS must be operational.
 - Can be in different domains if the domains have a two-way trust relationship.
 - When adding a vCenter instance to a Linked Mode group, the installer must be run by a domain user who is an administrator on both the machines.

- Must have network time synchronization. The installer validates that the machine clocks are not more than 5 minutes apart.
- The roles defined on each vCenter Server system in a linked mode group are replicated to the other systems in the group.
- Troubleshooting:
 - Conflicts can occur, when you clone a vCenter Server instance that is running in a VM and you do not use sysprep or a similar utility to ensure that the cloned vCenter Server instance has a globally unique identifier (GUID).
 - The DNS name of the machine must match with the actual machine name.
 - Joining a Linked Mode group:
 1. Verify that the vCenter Server domain name matches the machine name.
 2. Update the URLs to make them compatible with the new domain name and machine name (if required).
 3. Join the vCenter Server system to a Linked Mode group.
- When you are logged in to a vCenter Server system that is part of a connected group, you can monitor the health of services running on each server in the group.
- Client uses ports 80 and 443 to communicate with vCenter Server and ESX/ESXi hosts. These ports cannot be changed.
- The SNMP traps sent by vCenter Server are defined in VMWARE-VC-EVENT-MIB.mib
- ESX/ESXi includes an SNMP agent embedded in hostd that can both send traps and receive polling requests such as GET requests. This agent is referred to as the embedded SNMP agent. Versions of ESX prior to ESX 4.0 included a Net-SNMP-based agent. You can continue to use this Net-SNMPbased agent in ESX 4.0 with MIBs supplied by your hardware vendor and other third-party management applications. However, to use the VMware MIB files, you must use the embedded SNMP agent. By default, the embedded SNMP agent is disabled. To enable it, use the vSphere CLI command `vicfg-snmp`.
- Both the embedded SNMP agent and the Net-SNMP-based agent available in the ESX service console listen on UDP port 161 by default. If you enable both of these agents for polling on an ESX host, you must change the port used by at least one of them.
- Can use SNMP to monitor guest operating systems or applications running in VMs. Do not install agents in the VM that are intended to monitor physical hardware.
- VMware MIB Files

MIB File	Description
VMWARE-ROOT-MIB.mib	Contains VMware's enterprise OID and top level OID assignments.
VMWARE-AGENTCAP-MIB.mib	Defines the capabilities of the VMware agents by product versions.
VMWARE-ENV-MIB.mib	Defines variables and trap types used to report on the state of physical hardware components of the host computer.
VMWARE-OBSOLETE-MIB.mib	Defines OIDs that have been made obsolete to maintain backward compatibility with earlier versions of ESX/ESXi. Includes variables formerly defined in the files VMWARE-TRAPS-MIB.mib and VMWARE-VMKERNEL-MIB.mib.
VMWARE-PRODUCTS-MIB.mib	Defines OIDs to uniquely identify each SNMP agent on each VMware platform by name, version, and build platform.
VMWARE-RESOURCES-MIB.mib	Defines variables used to report information on resource usage of the VMkernel, including physical memory, CPU, and disk utilization.
VMWARE-SYSTEM-MIB.mib	The VMWARE-SYSTEM-MIB.mib file is obsolete. Use the SNMPv2-MIB to obtain information from sysDescr.0 and sysObjec ID.0.
VMWARE-TC-MIB.mib	Defines common textual conventions used by VMware MIB files.
VMWARE-VC-EVENTS-MIB.mib	Defines traps sent by vCenter Server. Load this file if you use vCenter Server to send traps.
VMWARE-VMINFO-MIB.mib	Defines variables for reporting information about VMs, including VM traps.

- ESX/ESXi System Logs

Component	Location
ESX Server 2.x Service log	<code>/var/log/vmware/vmware-serverd.log</code>
ESX Server 3.x or ESX Service log	<code>/var/log/vmware/hostd.log</code>
vSphere Client Agent log	<code>/var/log/vmware/vpx/vpxa.log</code>
Virtual Machine Kernel Core file	<code>/root/vmkernel-core.<date></code> and <code>/root/vmkernel-log.<date></code> present after rebooting
Syslog log	<code>/var/log/messages</code>
Service Console Availability report	<code>/var/log/vmkernel</code>
VMkernel Messages	<code>/var/log/vmkernel</code>
VMkernel Alerts & Availability report	<code>/var/log/vmkernel</code>
VMkernel Warning	<code>/var/log/vmkernelwarning</code>
Virtual Machine log file	The same directory as the .vmx file for the VM.

- vSphere Client System Logs

Component	Location
vSphere Client Installation log	Temp directory on the vSphere Client machine. e.g.: C:\Documents and Settings\ <user name="">\Local Settings\Temp\vmmsi.log</user>
vSphere Client Service log	\vpx directory in the Application Data directory on the vSphere Client machine. e.g.: C:\Documents and Settings\ <user name="">\Local Settings\Application Data \vpx\viclient-x.log x(=0, 1, ... 9)</user>

- VMware Server System Logs

Component	Operating System	Location
VM Console log	Windows	Temp directory e.g.: C:\Documents and Settings\ <username>\Local Settings\Temp\vmware-<username>-<PID>.log</username>
	Linux	Temp directory e.g.: /tmp/vmware-<username>/ui-<PID>.log
VM log	Windows & Linux	vmware.log Located in the same directory as the VM .vmx file.
VM Event log	Windows	C:\Program Files\VMware\VMware Virtual Infrastructure\vmserverdRoot\eventlog\vent- <path_to_configuration_file>.vmx.log
	Linux	/var/log/vmware/event-<path_to_configuration_file>.vmx.log
VM Conf file	Windows	<virtual_machine_name>.vmx Located in the folder where VMs are stored.
	Linux	<virtual_machine_name>.vmx Located in the folder where VMs are stored.

- All ESX/ESXi hosts run a syslog service (syslogd).
- ESXi hosts can use the vSphere Client or the vSphere CLI command `vicfg-syslog`
- Cannot use the vSphere Client or `vicfg-syslog` to configure syslog behavior for an ESX host. To configure syslog for an ESX host, you must edit the `/etc/syslog.conf` file.
- Libraries – Central repositories for VM provisioning media such as VM templates, ISO images, floppy images, VMDK files, guest customization files, and so on.
- Guided Consolidation ports

Port	Protocol	Service	Description	MS Windows
135	TCP/UDP	Loc-srv/epmap	Microsoft DCE Locator service (End-point Mapper).	DHCP, DNS & WINS Server
137	TCP/UDP	Netbios-ns	NetBIOS names service.	WINS & DNS Server
138	TCP/UDP	Netbios-dgm	NetBIOS datagram	
139	TCP/UDP	Netbios-ssn	NetBIOS Session Windows File and Printer sharing.	
445	TCP/UDP	DNS	DNS Direct Hosting port.	Active Directory

- Guided Consolidation can be installed together with vCenter Server, or can be installed on a separate host. For best performance, install Guided Consolidation on a separate host. It includes the following services:
 - vCenter Collector Service – Discovers domains and systems within domains. Collects performance data on those systems.
 - vCenter Provider Service – Helper service to the Collector Service. Communicates with target systems and passes the data back.
 - vCenter Guided Consolidation – Coordinates all communication among Guided Consolidation components.
- You can analyze up to 100 systems simultaneously.
- The following formula is used to resize converted disks:
 - amount of space used on physical disk * 1.25 = resultant virtual disk size
- Virtual disks are set to a size of 4GB or larger.
- Disconnecting a managed host does not remove it from vCenter Server; it temporarily suspends all monitoring activities performed by vCenter Server.
- Deploying an OVF template is similar to deploying a VM from a template. However, you can deploy an OVF template from any local file system accessible from the vSphere Client machine, or from a remote web server.
- OVF advantages:
 - OVF files are compressed, allowing for faster downloads.
 - The vSphere Client validates an OVF file before importing it.
- A vApp is a container, like a resource pool and can contain one or more VMs. A vApp can power on and power off, and can also be cloned.
- The vApp metadata resides in the vCenter Server's database
- You can add an object, such as a VM or another vApp, to an existing vApp.
- An IP pool is a network configuration that is assigned to a network used by a vApp. The vApp can then leverage vCenter Server to automatically provide an IP configuration to its VMs.
- Each application within the service will be powered on according to how the startup order is set. When powering on a vApp within a DRS cluster in manual mode, no DRS recommendations are generated for VM placements. The power on operation performs as if DRS is run in a semi-automatic or automatic mode for the initial placements of the VMs. This does not affect VMotion recommendations. Recommendations for individual powering on and powering off of VMs are also generated for vApps that are running
- A VM's name can be up to 80 characters long. Names are case insensitive.
- Virtual machine version 4 — Compatible with ESX 3.0 and greater hosts and VMware Server 1.0 and greater hosts.

- Virtual machine version 7 — Compatible with ESX 4.0 and greater hosts. Provides greater VM functionality.
- Paravirtual SCSI (PVSCSI) adapters are high-performance storage adapters that can result in greater throughput and lower CPU utilization.
 - Best suited for high performance storage environments.
 - Not suited for DAS environments. VMware recommends that you create a primary adapter (LSI Logic by default) for use with a disk that will host the system software (boot disk) and a separate PVSCSI adapter for the disk that will store user data, such as a database.
- Paravirtual SCSI adapters are available for VMs running hardware version 7 and greater. They are supported on the following guest operating systems:
 - Windows Server 2008
 - Windows Server 2003
 - Red Hat Linux (RHEL) 5
- Features not supported with Paravirtual SCSI adapters:
 - Boot disks
 - Record/Replay
 - Fault Tolerance
 - MSCS Clustering
- SCSI controller types:
 - BusLogic Parallel
 - LSI Logic SAS
 - LSI Logic Parallel
 - VMware Paravirtual
- Thin Provisioned Format – Use this format to save storage space. If a virtual disk supports clustering solutions such as Fault Tolerance, you cannot make the disk thin. You can manually convert the thin disk into thick.
- Thick Format – This is the default virtual disk format. It is not possible to convert the thick disk into thin. (EDIT: you can via Storage VMotion)
- Automatic VMware Tools upgrade is not supported for VMs with Solaris or Netware guest operating systems.
- If you are using a WYSE thin client device to conduct remote desktop sessions using VMware VDI, installing WYSE Multimedia Support in the guest operating system improves the performance of streaming video. WYSE Multimedia Support is supported on the Windows 2003 and Windows XP guest operating systems only. WYSE Multimedia Support is installed as part of a VMware Tools installation or upgrade.
- Virtual machines with hardware versions lower than 4 can run on ESX4 hosts but have reduced performance and capabilities. In particular, you cannot add or remove virtual devices on VMs with hardware versions lower than 4 when they reside on an ESX4 host.
- Virtual Machine Hardware Versions:

	Version 7	Version 4	Version 3	Compatible with vCenter Server version
ESX/ESXi 4.x	create, edit, run	create, edit, run	run	vCenter Server 4.x
ESX Server 3.x	–	create, edit, run	run	VirtualCenter Server 2.x and higher
ESX Server 2.x	–	–	create, edit, run	VirtualCenter Server 1.x and higher

- SCSI Bus Sharing list:

Option	Description
None	Virtual disks cannot be shared by other VMs.
Virtual	Virtual disks can be shared by VMs on same server.
Physical	Virtual disks can be shared by VMs on any server.

- Memory/CPU Hotplug – VMware Tools must be installed for hotplug functionality to work properly.
- VMI – A paravirtualization standard that enables improved performance for VMs capable of utilizing it.
- Enabling paravirtualization utilizes one of the VM's six virtual PCI slots
- A VM with paravirtualization enabled and that is powered off can be moved manually to a host that does not support paravirtualization. However, this can result in reduced performance.
- N-port ID virtualization (NPV) – Provides the ability to share a single physical Fibre Channel HBA port among multiple virtual ports, each with unique identifiers. This allows control over VM access to LUNs on a per-VM basis.
- NPV support is subject to the following limitations:
 - NPV must be enabled on the SAN switch.
 - NPV is supported only for VMs with RDM disks.
 - The physical HBAs on the ESX host must have access to a LUN using its WWNs in order for any VMs on that host to have access to that LUN using their NPV WWNs.
 - The physical HBAs on the ESX host must support NPV.
 - Each VM can have up to 4 virtual ports. NPV-enabled VMs are assigned exactly 4 NPV-related WWNs. Can utilize up to 4 physical HBAs for NPV purposes.
- A VM with WWNs that are already in use on the storage network is prevented from powering on.
- While hyperthreading does not double the performance of a system, it can increase performance by better utilizing idle resources.
- The advanced CPU settings are useful only for fine-grained tweaking of critical VMs.
- NUMA memory node affinity enables fine-grained control over how VM memory is distributed to host physical memory.

- Specify nodes to be used for future memory allocations only if you have also specified CPU affinity.
- The following NIC types are supported:

Flexible	Supported on VMs that were created on ESX Server 3.0 or greater and that run 32-bit guest operating systems. The Flexible adapter functions as a vance adapter if VMware Tools is not installed in the VM and as a vmxnet driver if VMware Tools is installed in the VM.
e1000	Emulates the functioning of an E1000 network card. It is the default adapter type for VMs that run 64-bit guest operating systems.
Enhanced vmxnet	An upgraded version of the vmxnet device with enhanced performance. It requires that VMware Tools be installed in the VM.
vmxnet 3	Next generation vmxnet device with enhanced performance and enhanced networking features. It requires that VMware Tools be installed in the VM, and is available only on VMs with hardware version 7 and greater.

- Independent disks are not affected by snapshots.
- Two modes for independent disks:
 - Persistent – The disk operates normally except that changes to the disk are permanent even if the VM is reverted to a snapshot.
 - Nonpersistent – The disk appears to operate normally, but whenever the VM is powered off or reverted to a snapshot, the contents of the disk return to their original state. All later changes are discarded.
- VMDirectPath I/O allows a guest operating system on a VM to directly access physical PCI and PCIe devices connected to a host. Each VM can be connected to up to two PCI devices. PCI devices connected to a host can be marked as available for passthrough from the Hardware Advanced Settings in the Configuration tab for the host.
- Paravirtual SCSI (PVSCSI) adapters are high-performance storage adapters that can provide greater throughput and lower CPU utilization. PVSCSI adapters are best suited for environments, especially SAN environments, running I/O-intensive applications. PVSCSI adapters are not suited for DAS environments.
- Hardware requirements for customizing the guest operating system:
 - Must reside on a disk attached as SCSI 0:0 node in the VM configuration.
 - If a VM has mixed IDE and SCSI disks, the first IDE disk is considered the boot disk, and vCenter Server passes it to the customizer.
- If the new VM encounters customization errors while it is booting
- Customization errors are logged to (Windows guest) %WINDIR%\temp\vmware-imc or (Linux guest) /var/log/vmware/customization.log.
- When you migrate a suspended VM, the new host for the VM must meet CPU compatibility requirements, because the VM must resume executing instructions on the new host.
- Use of Jumbo Frames is recommended for best VMotion performance.
- Some restrictions apply when migrating VMs with snapshots. You cannot migrate a virtual machine with snapshots with Storage VMotion.
- You can migrate as long as the VM is being migrated to a new host without moving its configuration file or disks (the VM must reside on shared storage accessible to both hosts).
- Reverting to a snapshot after migration with VMotion might cause the VM to fail, because the migration wizard cannot verify the compatibility of the VM state in the snapshot with the destination host.
- During a migration with Storage VMotion, you can transform virtual disks from thick-provisioned to thin or from thin-provisioned to thick.
- Storage VMotion is subject to the following requirements and limitations:
 - Virtual machines with snapshots cannot be migrated using Storage VMotion.
 - Virtual machine disks must be in persistent mode or be raw device mappings (RDMs). For virtual compatibility mode RDMs, you can migrate the mapping file or convert to thick-provisioned or thinprovisioned disks during migration as long as the destination is not an NFS datastore. For physical compatibility mode RDMs, you can migrate the mapping file only.
 - Must have a license that includes Storage VMotion.
 - ESX/ESXi 3.5 hosts must be licensed and configured for VMotion. ESX/ESXi 4.0 and later hosts do not require VMotion configuration in order to perform migration with Storage VMotion.
 - A particular host can be involved in up to two migrations with VMotion or Storage VMotion at one time.
 - vSphere supports a maximum of eight simultaneous VMotion, cloning, deployment, or Storage VMotion accesses to a single VMFS3 datastore, and a maximum of four simultaneous VMotion, cloning, deployment, or Storage VMotion accesses to a single NFS or VMFS2 datastore. A migration with VMotion involves one access to the datastore. A migration with Storage VMotion involves one access to the source datastore and one access to the destination datastore
- Disks are converted from thin to thick format or thick to thin format only when they are copied from one datastore to another. If you choose to leave a disk in its original location, the disk format is not converted.
- Thin or thick provisioned – not available for RDMs in physical compatibility mode. If you select this option for a virtual compatibility mode RDM, the RDM is converted to a virtual disk. RDMs converted to virtual disks cannot be converted back to RDMs.
- You can run the storage vmotion command in either interactive or noninteractive mode.
 - Interactive mode, type **svmotion --interactive**.
 - Noninteractive mode: `svmotion [Standard CLI options] --datacenter=<datacenter name> --vm '<VM config datastore path>:<new datastore>' [--disks '<virtual disk datastore path>:<new datastore>, <virtual disk datastore path>:<new datastore>']`
- A snapshot captures the entire state of the VM at the time you take the snapshot. This includes:
 - Memory state – The contents of the VM's memory.

- Settings state – The VM settings.
- Disk state – The state of all the VM’s virtual disks.
- Snapshots of raw disks, RDM physical mode disks, and independent disks are not supported.
- Change Disk Mode to independent to Exclude Virtual Disks from Snapshots
- Persistent – Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk.
- Nonpersistent – Changes are discarded when you power off or reset the VM. Nonpersistent mode enables you to restart the VM with a virtual disk in the same state every time. Changes to the disk are actually written to and read from a redo log file that is deleted when you power off or reset.
- Snapshots:
 - Delete – commits the snapshot data to the parent and removes the selected snapshot.
 - Delete All – commits all the immediate snapshots before the You are here current state to the base disk and removes all existing snapshots for that VM.
 - Revert to Snapshot – a shortcut to the parent snapshot of “You are here”.
- If you use Active Directory groups for permissions, make sure that they are security groups and not distribution groups.
- Users who are currently logged in and are removed from the domain retain their vSphere permissions only until the next validation period (the default is every 24 hours).
- A role is a predefined set of privileges. Privileges define basic individual rights required to perform actions and read properties. When you assign a user or group permissions, you pair the user or group with a role and associate that pairing with an inventory object.
- Default roles:
 - System roles – System roles are permanent. You cannot edit the privileges associated with these roles.
 - Sample roles – VMware provides sample roles for convenience as guidelines and suggestions. You can modify or remove these roles.
- You can also create completely new roles.
- All roles permit the user to schedule tasks by default. Users can schedule only tasks they have permission to perform at the time the tasks are created.
- Default roles:

Role	Role Type	Description of User Capabilities
No Access	system	Cannot view or change the assigned object. available in ESX/ESXi and vCenter Server.
Read Only	system	View the state and details about the object. available on ESX/ESXi and vCenter Server.
Administrator	system	All privileges for all objects. available in ESX/ESXi and vCenter Server.
Virtual Machine Power User	sample	allow the user to interact with and make hardware changes to VMs, as well as perform snapshot operations. available only on vCenter Server.
Virtual Machine User	sample	allow the user to interact with a VM’s console, insert media, and perform power operations. available only on vCenter Server.
Resource Pool Administrator	sample	allow the user to create child resource pools and modify the configuration of the children, but not to modify the resource configuration of the pool or cluster on which the role was granted. Also allows the user to grant permissions to child resource pools, and assign VMs to the parent or child resource pools. available only on vCenter Server.
VMware Consolidated Backup User	sample	used by the VMware Consolidated Backup product and should not be modified. available only on vCenter Server.
Datastore Consumer	sample	allow the user to consume space on the datastores on which this role is granted. available only on vCenter Server.
Network Consumer	sample	allow the user to assign VMs or hosts to networks available only on vCenter Server.

- If you create or edit a role on a vCenter Server system that is part of a connected group in Linked Mode, the changes you make are propagated to all other vCenter Server systems in the group. Assignments of roles to specific users and objects are not shared across linked vCenter Server systems.
- Permissions grant users the right to perform the activities specified by the role on the object to which the role is assigned
- By default, all users who are members of the Windows Administrators group on the vCenter Server system have the same access rights as any user assigned to the Administrator role on all objects.
- Propagation is set per permission, not universally applied. Permissions defined for a child object always override those propagated from parent objects.
- You cannot set permissions directly on a vNetwork Distributed Switches. To set permissions for a vNetwork Distributed Switch and its associated dvPort Groups, set permissions on a parent object, such a folder or datacenter, and select the option to propagate these permissions to child objects.
 - If no permission is defined for the user on that object, the user is assigned the union of privileges assigned to the groups for that object.
 - If a permission is defined for the user on that object, the user’s permission takes precedence over all group permissions
- Reports are updated every 30 minutes.

- Map views are updated every 30 minutes
- Alarms are notifications that occur in response to selected events, conditions, and states that occur with objects in the inventory.
- Alarms are composed of a trigger and an action.
- Alarms have two types of triggers: condition/state triggers, and event triggers.
- Condition or State Triggers Monitor the current condition or state of VMs, hosts, and datastores.
- Event Triggers Monitors events that occur in response to operations occurring with any managed object in the inventory, the vCenter Server system, or the license server.
- Condition and state triggers use one of the following operator sets to monitor an object:
 - Is equal to and Is not equal to
 - Is above and Is below
- Event triggers use arguments, operators, and values to monitor operations that occur in the vServer System.
- Alarm actions are operations that occur in response to triggered alarms.
- The default VMware alarms do not have actions associated with them. You must manually associate actions with the default alarms.
- You can disable an alarm action from occurring without disabling the alarm itself.
- You disable alarm actions for a selected inventory object.
- When you disable the alarm actions for an object, they continue to occur on child objects.
- When you disable alarm actions, all actions on all alarms for the object are disabled. You cannot disable a subset of alarm actions.
- The SNMP agent included with vCenter Server can be used to send traps when alarms are triggered on a vCenter Server.
- Alarm reporting can further restrict when a condition or state alarm trigger occurs by adding a tolerance range and a trigger frequency to the trigger configuration.
- The tolerance range specifies a percentage above or below the configured threshold point, after which the alarm triggers or clears.
- Condition threshold + Tolerance Range = Trigger alarm
- The trigger frequency is the time period during which a triggered alarm action is not reported again. By default, the trigger frequency for the default VMware alarms is set to 5 minutes.
- Statistical data consists of CPU, memory, disk, network, system, and VM operations metrics.
- Collection intervals determine the time period during which statistics are aggregated and rolled up, and the length of time the statistics are archived in the vCenter database. By default, vCenter Server has four collection intervals: Day, Week, Month, and Year.
- Real-time statistics are not stored in the database. They are stored in a flat file on ESX/ESXi hosts and in memory on the vCenter Server systems
- Real-time statistics are collected directly on an ESX/ESXi host every 20 seconds (60 seconds for ESX Server 2.x hosts).
 - On ESX hosts, the statistics are kept for one hour, after which 180 data points (15 -20 second samples) will have been collected.
 - On ESXi hosts, the statistics are kept for 30 minutes, after which 90 data points will have been collected.
- Collection Intervals:

1 Day	5 Minutes
1 Week	30 Minutes
1 Month	2 Hours
1 Year	1 Day

- You can change the frequency at which statistic queries occur, the length of time statistical data is stored in the vCenter Server database, and the amount of statistical data collected.
- Not all attributes are configurable for each collection interval.
- You can assign a collection level of 1- 4 to each collection interval, with level 4 having the largest number of counters.
- By default, all collection intervals use collection level 1.
- Generally, you need to use only collection levels 1 and 2 for performance monitoring and analysis.
- By default, statistics are stored in the vCenter Server database for one year. You can increase this to three years.
- You cannot view datastore metrics in the advanced charts. They are only available in the overview charts.
- CPU Performance Enhancement Advice
 1. Verify that VMware Tools is installed on every VM on the host.
 2. Compare the CPU usage value of a VM with the CPU usage of other VMs on the host or in the resource pool. The stacked bar chart on the host's **Virtual Machine** view shows the CPU usage for all VMs on the host.
 3. Determine whether the high ready time for the VM resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the VM.
 4. Increase the CPU shares to give the VM more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority VMs to guarantee that they receive the required CPU cycles.
 5. Increase the amount of memory allocated to the VM. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the ESX/ESXi host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.

6. Reduce the number of virtual CPUs on a VM to only the number required to execute the workload. For example, a single-threaded application on a four-way VM only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
 7. If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more VMs onto the new host.
 8. Upgrade the physical CPUs or cores on the host if necessary.
 9. Use the newest version of ESX/ESXi, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.
- Memory Performance Enhancement Advice
 1. Verify that VMware Tools is installed on each VM. The balloon driver is installed with VMware Tools and is critical to performance.
 2. Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused VM memory by ballooning and swapping. Generally, this does not impact VM performance.
 3. Reduce the memory space on the VM, and correct the cache size if it is too large. This frees up memory for other VMs.
 4. If the memory reservation of the VM is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other VMs on the host.
 5. Migrate one or more VMs to a host in a DRS cluster.
 6. Add physical memory to the host.
 - Disk I/O Performance Enhancement Advice
 1. Increase the VM memory. This should allow for more operating system caching, which can reduce I/O activity. Note that this may require you to also increase the host memory. Increasing memory might reduce the need to store data because databases can utilize system memory to cache data and avoid disk access. To verify that VMs have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
 2. Defragment the file systems on all guests.
 3. Disable antivirus on-demand scans on the VMDK and VMEM files.
 4. Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput.
 5. Use Storage VMotion to migrate I/O-intensive VMs across multiple ESX/ESXi hosts.
 6. Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
 7. Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the VM by adjusting the `Disk.SchedNumReqOutstanding` parameter. For more information, see the *Fibre Channel SAN Configuration Guide*.
 8. For resource-intensive VMs, separate the VM's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use.
 9. On systems with sizable RAM, disable memory trimming by adding the line `MemTrimRate=0` to the VM's `.VMX` file.
 10. If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
 11. For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
 12. Use the most current ESX/ESXi host hardware.
 - Networking Performance Enhancement Advice
 1. Verify that VMware Tools is installed on each VM.
 2. If possible, use `vmxnet3` NIC drivers, which are available with VMware Tools. They are optimized for high performance.
 3. If VMs running on the same ESX/ESXi host communicate with each other, connect them to the same vSwitch to avoid the cost of transferring packets over the physical network.
 4. Assign each physical NIC to a port group and a vSwitch.
 5. Use separate physical NICs to handle the different traffic streams, such as network packets generated by VMs, iSCSI protocols, VMotion tasks, and service console activities.
 6. Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some VMs to a vSwitch with a lighter load or to a new vSwitch.
 7. If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
 8. Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1Gbps are not reset to 100Mbps because they are connected to an older switch.
 9. Verify that all NICs are running in full duplex mode. Hardware connectivity issues might result in a NIC resetting itself to a lower speed or half duplex mode.
 10. Use vNICs that are TSO-capable, and verify that TSO-Jumbo Frames are enabled where possible.
 - Tasks represent system activities that do not complete immediately, such as migrating a VM.

- If you are logged in to a vCenter Server system that is part of a Connected Group, a column in the task list displays the name of the vCenter Server system on which the task was performed.

Appendix A – Defined privileges

Appendix B – Installing the MS sysprep tools

Appendix C – Performance metrics

ESX Configuration Guide

- A vNetwork Distributed Switch acts as a single vSwitch across all associated hosts on a datacenter. This allows virtual machines to maintain consistent network configuration as they migrate across multiple hosts. A dvPort is a port on a vNetwork Distributed Switch.
- The VMkernel TCP/IP networking stack supports iSCSI, NFS, and VMotion. Virtual machines run their own systems' TCP/IP stacks and connect to the VMkernel at the Ethernet level through virtual switches.
- TCP Segmentation Offload (TSO), allows a TCP/IP stack to emit very large frames (up to 64KB) even though the maximum transmission unit (MTU) of the interface is smaller. The network adapter then separates the large frame into MTU-sized frames and prepends an adjusted copy of the initial TCP/IP headers.
- The default number of logical ports for a vSwitch is 56.
- Each uplink adapter associated with a vSwitch uses one port.
- You can create a maximum of 127 vSwitches on a single host.
- Maximum of 512 port groups on a single host.
- For a port group to reach port groups located on other VLANs, the VLAN ID must be set to 4095. If you enter 4095, the port group can see traffic on any VLAN while leaving the VLAN tags intact.
- VLAN ID is a number between 1 and 4094.
- ESX supports only NFS version 3 over TCP/IP.
- You can create a maximum of 16 service console ports in ESX.
- CDP advertisements typically occur once a minute.
- dvPort group properties include:
 - Port Binding - when ports are assigned to virtual machines connected to this dvPort group.
 - Static binding - to assign a port to a virtual machine when the virtual machine is connected to the dvPort group.
 - Dynamic binding - to assign a port to a virtual machine the first time the virtual machine powers on after it is connected to the dvPort group.
 - Ephemeral - for no port binding.
 - Whether to allow live port moving.
 - Config reset at disconnect to discard per-port configurations when a dvPort is disconnected from a virtual machine.
 - Binding on host allowed to specify that when vCenter Server is down, ESX can assign a dvPort to a virtual machine.
 - Port name format to provide a template for assigning names to the dvPorts in this group.
- Private VLANs are used to solve VLAN ID limitations.
- A private VLAN is identified by its primary VLAN ID. A primary VLAN ID can have multiple secondary VLAN IDs associated with it. Primary VLANs are Promiscuous, so that ports on a private VLAN can communicate with ports configured as the primary VLAN. Ports on a secondary VLAN can be either:
 - Isolated - communicating only with promiscuous ports
 - Community - communicating with both promiscuous ports and other ports on the same secondary VLAN.
- Only one VMotion and IP storage port group for each ESX host.
- You can enable or disable IPv6 support on the host.
- The following networking policies can be applied:
 - Security
 - Promiscuous Mode - In non-promiscuous mode, a guest adapter listens only to traffic forwarded to own MAC address. In promiscuous mode, it can listen to all the frames. By default, guest adapters are set to non-promiscuous mode.
 - MAC Address Changes - the guest OS changes the MAC address of the adapter to anything other than what is in the .vmx
 - Forged Transmits - Outbound frames with a source MAC address that is different from the one set on the adapter are dropped.
 - Traffic shaping
 - Traffic shaping policy is defined by three characteristics: average bandwidth, peak bandwidth, and burst size.
 - ESX shapes outbound network traffic on vSwitches and both inbound and outbound traffic on a vNetwork Distributed Switch.
 - Peak bandwidth cannot be less than the specified average bandwidth.
 - NIC Teaming (Load balancing and failover)
 - Load Balancing
 1. Route based on the originating port ID — Choose an uplink based on the virtual port where the traffic entered the virtual switch.
 2. Route based on ip hash — Choose an uplink based on a hash of the source and destination IP addresses of each packet.
 3. Route based on source MAC hash — Choose an uplink based on a hash of the source Ethernet.

4. Use explicit failover order — Always use the highest order uplink from the list of Active adapters which passes failover detection criteria.
 - o IP-based teaming requires that the physical switch be configured with etherchannel. For all other options, etherchannel should be disabled.
 - o Incoming traffic is controlled by the load balancing policy on the physical switch
 - o Network failover detection
 - o Link Status only
 - o Beacon probing - Do not use beacon probing with IP-hash load balancing.
 - o Notify Switches - a notification is sent out over the network to update the lookup tables on physical switches. In almost all cases, this process is desirable for the lowest latency of failover occurrences and migrations with VMotion. Do not use this option when the virtual machines using the port group are using Microsoft Network Load Balancing in unicast mode.
 - o Failback - determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to Yes (default), the adapter is returned to active duty immediately upon recovery.
 - o Failover Order
 1. Active Uplinks
 2. Standby Uplinks
 3. Unused Uplinks
 - o When using IP-hash load balancing, do not configure standby uplinks.
 - o VLAN - The VLAN policy allows virtual networks to join physical VLANs - vNetwork Distributed Switch only (dvPorts).
 - o Port blocking policies - vNetwork Distributed Switch only (dvPorts).
- o VMware uses the Organizationally Unique Identifier (OUI) 00:50:56 for manually generated addresses. You must set them in a virtual machine's configuration file: `ethernet<number>.addressType="static"`
- o Jumbo frames must be enabled at the host level using the command-line interface to configure the MTU size for each vSwitch.
- o TCP Segmentation Offload (TSO) is enabled on the VMkernel interface by default, but must be enabled at the virtual machine level.
- o To enable TSO at the virtual machine level, you must replace the existing vmxnet or flexible virtual network adapters with enhanced vmxnet virtual network adapters. This might result in a change in the MAC address of the virtual network adapter.
- o To check whether TSO is enabled on a particular VMkernel networking interface use the `esxcfg-vmknic -l` command. The list shows each TSO-enabled VMkernel interface with TSO MSS set to 65535.
- o If TSO is not enabled for a particular VMkernel interface, the only way to enable it is to delete the VMkernel interface and recreate the interface.
- o Jumbo frames up to 9kB (9000 bytes) are supported.
- o Use the `vicfg-vswitch -m <MTU> <vSwitch>` command to set the MTU size for the vSwitch.
- o Enabling jumbo frame support on a virtual machine requires an enhanced vmxnet adapter for that virtual machine.
- o NetQueue in ESX takes advantage of the capability of some network adapters to deliver network traffic to the system in multiple receive queues that can be processed separately. This allows processing to be scaled to multiple CPUs, improving receive-side networking performance.
- o NetQueue is enabled by default.
- o ESX supports a direct PCI device connection for virtual machines running on Intel Nehalem platforms. Each virtual machine can connect to up to 2 passthrough devices.
- o The following features are unavailable for virtual machines configured with VMDirectPath:
 - o VMotion
 - o Hot adding and removing of virtual devices
 - o Suspend and resume
 - o Record and replay
 - o Fault tolerance
 - o High availability
 - o DRS (limited availability; the virtual machine can be part of a cluster, but cannot migrate across hosts)
- o Software-initiated iSCSI is not available over 10GigE network adapters in ESX.
- o You cannot use IDE/ATA drives to store virtual machines.
- o Use local SATA storage, internal and external, in unshared mode only.
- o Some SAS storage systems can offer shared access
- o You can have up to 256 VMFS datastores per system, with a minimum volume size of 1.2GB.
- o Grow the existing datastore extent if the storage device where your datastore resides has free space. You can grow the extent up to 2 TB.
- o You can connect up to 32 hosts to a single VMFS volume.
- o Perform a rescan each time you:
 - o Create new LUNs on a SAN.
 - o Change the path masking on a host.
 - o Reconnect a cable.
 - o Make a change to a host in a cluster.
- o Do not rescan when a path is unavailable.

- To rescan adapters on all hosts managed by vCenter by right-clicking a datacenter, cluster, or folder and selecting **Rescan for Datastores**.
- ESX does not support the delegate user functionality that enables access to NFS volumes using non-root credentials
- Disk format on a NAS device is dictated by the NFS server, typically a thin format that requires on-demand space allocation.
- When your host accesses a virtual machine disk file on an NFS-based datastore, a .lck-XXX lock file is generated to prevent other hosts from accessing this file.
- If the underlying NFS volume, is read-only, make sure that the volume is exported as a read-only share by the NFS server, or configure it as a read-only on the ESX host.
- A diagnostic partition cannot be located on an iSCSI LUN accessed through a software iSCSI initiator.
- You can query and scan the host's diagnostic partition using the `vicfg-dumppart -l` command
- You can group datastores into folders.
- You can unmount:
 - NFS datastores
 - VMFS datastore copies mounted without resignaturing
- You can have up to 32 extents.
- You can grow an extent in an existing VMFS datastore. Only extents with free space immediately after them are expandable.
- If a shared datastore has powered on virtual machines and becomes 100% full, you can increase the datastore's capacity only from the host, with which the powered on virtual machines are registered.
- You can mount a VMFS datastore only if it does not collide with an already mounted VMFS datastore that has the same UUID (signature).
- When resignaturing a VMFS copy, ESX assigns a new UUID and a new label to the copy, and mounts the copy as a datastore distinct from the original.
- The default format of the new label assigned to the datastore is `snap-<snapID>-<oldLabel>`, where `<snapID>` is an integer and `<oldLabel>` is the label of the original datastore.
- Datastore resignaturing is irreversible.
- A spanned datastore can be resignatured only if all its extents are online.
- Pluggable Storage Architecture (PSA) is an open modular framework that coordinates the simultaneous operation of multiple multipathing plugins (MPPs). The VMkernel multipathing plugin that ESX provides by default is the VMware Native Multipathing Plugin (NMP). Two types of NMP subplugins, Storage Array Type Plugins (SATPs), and Path Selection Plugins (PSPs).
- The VMware NMP supports all storage arrays listed on the VMware storage HCL and provides a default path selection algorithm based on the array type.
- ESX offers an SATP for every type of array that VMware supports.
- By default, the VMware NMP supports the following PSPs:
 - Most Recently Used (MRU)
 - Fixed - with active-passive arrays that have a Fixed path policy, path thrashing might be a problem.
 - Round Robin (RR) - Uses a path selection algorithm that rotates through all available paths enabling load balancing across the paths.
- Claim rules defined in the `/etc/vmware/esx.conf` file, the host determines which multipathing plugin (MPP) should claim the paths.
- By default, the host performs a periodic path evaluation every 5 minutes.
- Active multiple working paths currently used for transferring data are marked as Active (I/O). In ESX 3.5 or earlier, the term active means the only path that the host is using to issue I/O to a LUN.
- Standby path is operational and can be used for I/O if active paths fail.
- If you created a virtual disk in the thin format, you can later inflate it to its full size.
- RDM offers several benefits. User-Friendly Persistent Names, Dynamic Name Resolution, Distributed File Locking, File Permissions, File System Operations, Snapshots, vMotion, SAN Management Agents and N-Port ID Virtualization(NPIV).
- Certain limitations exist when you use RDMs:
 - Not available for block devices or certain RAID devices.
 - Available with VMFS-2 and VMFS-3 volumes only.
 - No snapshots in physical compatibility mode.
 - No partition mapping. It requires a whole LUN.
- Key contents of the metadata in the mapping file include the location of the mapped device (name resolution), the locking state of the mapped device, permissions, and so on.
- You cannot perform vMotion or Storage vMotion between datastores when NPIV is enabled.
- VMware protects the service console with a firewall. It also mitigates risks using other methods:
 - Only services essential to managing its functions.
 - By default, installed with a high-security setting. All outbound ports are closed.
 - By default, all ports not specifically required for management access to the service console are closed.
 - By default, weak ciphers are disabled and all communications from clients are secured by SSL. Default certificates created on ESX use SHA-1 with RSA encryption as the signature algorithm.
 - The Tomcat Web service, has been modified to run only those functions required.
 - VMware monitors all security alerts (for the RHEL5 distribution and 3rd party software).
 - Insecure services such as FTP and Telnet are not installed.
 - The number of applications that use a `setuid` or `setgid` flag is minimized.

- ESX can automate whether services start based on the status of firewall ports, but this only applies to service settings configured through the vSphere Client or applications created with the vSphere Web services SDK. Doesn't apply to changes made with the `esxcfg-firewall` utility or configuration files in `/etc/init.d/`.

Port	Purpose	Interface	Traffic type
22	SSH Server	Service Console	Incoming TCP
80	HTTP access and WS-Management	Service Console	Incoming TCP
123	NTP Client	Service Console	Outgoing UDP
427	The CIM client SLPv2 to find CIM servers.	Service Console	Incoming and outgoing UDP
443	HTTPS access - vmware-hostd vCenter Server access to ESX hosts Client access to vCenter Server and ESX hosts WS-Management Client access to vSphere Update Manager Converter access to vCenter Server Web Access to vCenter Server and ESX hosts	Service Console	Incoming TCP
902	Host access to other hosts for migration and provisioning Authentication traffic for ESX (xinetd/vmware-authd) Client access to virtual machine consoles (UDP) Status update (heartbeat) connection from ESX to vCenter Server	Service Console	Incoming TCP, outgoing UDP
903	Remote console traffic from VI client & Web Access (xinetd/vmware-authd-mks)	Service Console	Incoming TCP
2049	Transactions from NFS storage devices	VMkernel	Incoming and outgoing TCP
2050-2250	Between ESX hosts for HA and EMC Autostart Manager	Service Console	Outgoing TCP, incoming and outgoing UDP
3260	Transactions to iSCSI storage devices	VMkernel & Service Console	Outgoing UDP
5900-5964	RFB protocol, which is used by management tools such as VNC	Service Console	Incoming and outgoing TCP
5989	CIM XML transactions over HTTPS	Service Console	Incoming and outgoing TCP
8000	VMotion requests	VMkernel	Incoming and outgoing TCP
8042-8045	Between ESX hosts for HA and EMC Autostart Manager	Service Console	Outgoing TCP, incoming and outgoing UDP
8100, 8200	Between ESX hosts for Fault Tolerance	Service Console	Outgoing TCP, incoming and outgoing UDP

PLUS installed management agents and supported services such as NFS.

- Create a separate VLAN for communication with the service console.
- Configure network access for connections with the service console through a single virtual switch and one or more uplink ports.
- Set up a separate VLAN or virtual switch for vMotion and network attached storage.
- The iSCSI initiator relies on being able to get MAC address changes from certain types of storage. If you are using ESX iSCSI and have iSCSI storage, set the MAC Address Changes option to Accept.
- A legitimate need for more than one adapter to have the same MAC address, is if you are using Microsoft Network Load Balancing in unicast mode. When NLB is used in the standard multicast mode, adapters do not share MAC addresses.
- ESX uses the Pluggable Authentication Modules (PAM) structure for authentication. The PAM configuration in `/etc/pam.d/vmware-authd`, ESX uses `/etc/passwd` authentication, but you can configure ESX to use another distributed authentication mechanism.
- CIM transactions also use ticket-based authentication in connecting with the `vmware-hostd` process.
- Management functions with `username/password > vmware-hostd > Service Console`
- VM console with `ticket > vmkauthd > vm` in VMkernel
- `vicfg` commands do not perform an access check.
- The `vpxuser` is used for vCenter Server permissions.
- The root user and `vpxuser` permissions are the only users not assigned the No Access role by default.

- ESX supports SSL v3 and TLS v1.
- All network traffic is encrypted as long as:
 - Did not change the Web proxy service to allow unencrypted traffic for the port.
 - Service console firewall is configured for medium or high security.
- The default location for your certificate is /etc/vmware/ssl/ on the ESX host. The certificate consists of two files: the certificate itself (rui.crt) and the private-key file (rui.key).
- The ESX host generates certificates the first time the system is started.
- Each time you restart the vmware-hostd process, the mgmt-vmware script searches for existing certificate files (rui.crt and rui.key). If it cannot find them, it generates new certificate files.
- SSL timeout settings are set in /etc/vmware/hostd/config.xml.
- Do not set up certificates using passphrases.
- For certificates in a location other than the default location, set the location in /etc/vmware/hostd/proxy.xml.
- If you are performing activities that require root privileges, log in to the service console as a recognized user and acquire root privileges through the `sudo` command, which provides enhanced security compared to the `su` command.
- The service console firewall is configured to block all incoming and outgoing traffic, except for ports 22, 123, 427, 443, 902, 5989, 5988, pings (ICMP) and communication with DHCP and DNS (UDP only) clients.
- Medium security - All incoming traffic is blocked, except on the default ports and any ports you specifically open. Outgoing traffic is not blocked.
- Low security - There are no blocks on either incoming or outgoing traffic. This setting is equivalent to removing the firewall.
- Password aging restrictions are enabled for user logins by default.
- Maximum days - By default, passwords are set to never expire.
- Minimum days - The default is 0, meaning that the users can change their passwords any time.
- Warning time - The default is seven days.
- To change this for hosts use `esxcfg-auth`. Change for users use the command `chage`.
- By default, ESX uses the `pam_cracklib.so` plug-in. There is no restrictions on the root password, but the defaults for non-root users is:
 - minimum password length is nine
 - password length algorithm allows shorter passwords if the user enters a mix of character classes. M – CC = E where the Character Classes are upper, lower, digits and other.
 - retries is set to three
- The `pam_passwdqc.so` provides a greater number of options for fine-tuning password strength and performs password strength tests for all users, including the root user.
- `setuid` allows an application to temporarily change the permissions of the user running the application.
- `setgid` changes the permissions of the group running the application.
- Default `setuid` applications: `crontab`, `pam_timestamp_check`, `passwd`, `ping`, `pwdb_chkpwd`, `ssh-keysign`, `su`, `sudo`, `unix_chkpwd`, `vmkload_app`, `vmware-authd`, `vmware-vmx`. Default `setgid` Applications: `wall`, `lockfile`.
- Virtual Machine Recommendations:
 - Install Antivirus Software
 - Disable Copy and Paste Operations Between the Guest Operating System and Remote Console
 - Removing Unnecessary Hardware Devices
 - Limiting Guest Operating System Writes to Host Memory
 - Configuring Logging Levels for the Guest Operating System
- Host profiles eliminates per-host, configuration and maintain configuration consistency and correctness across the datacenter.
- Only supported for VMware vSphere 4.0 hosts.
- Host Profiles are only available when the appropriate licensing is in place.
- You can export a profile to a file that is in the VMware profile format (.vpf).

Appendix A – ESX Technical Support Commands

Command	Purpose
<code>esxcfg-advcfg</code>	advanced options
<code>esxcfg-auth</code>	Configures authentication
<code>esxcfg-boot</code>	bootstrap settings
<code>esxcfg-dumppart</code>	Configures a diagnostic partition
<code>esxcfg-firewall</code>	service console firewall ports
<code>esxcfg-info</code>	Information about the state of the service console, VMkernel, various subsystems in the virtual network, and storage resource hardware.
<code>esxcfg-init</code>	Internal initialization routines. Used for the bootstrap process you should not use it under any circumstances.
<code>esxcfg-module</code>	Sets driver parameters and modifies which drivers are loaded during startup.
<code>esxcfg-mpath</code>	multipath settings for your Fibre Channel or iSCSI disks.
<code>esxcfg-nas</code>	Manages NFS mounts

esxcfg-nics	physical network adapters
esxcfg-resgrp	resource group settings
esxcfg-route	default VMkernel gateway route
esxcfg-swiscsi	software iSCSI software adapter.
esxcfg-upgrade	Upgrades from ESX Server 2.x to ESX.
esxcfg-scsidevs	Prints a map of VMkernel storage devices to service console devices.
esxcfg-vmknic	VMkernel TCP/IP settings for VMotion, NAS, and iSCSI.
esxcfg-vswif	service console network settings.
esxcfg-vswitch	virtual machine network settings.

Appendix B – Linux Commands Used with ESX

Appendix C – Using vmkfstools

- `vmkfstools` utility is used to create and manipulate virtual disks, file systems, logical volumes, and physical storage devices.

Fibre Channel SAN Configuration Guide

- Zones define which HBAs can connect to which SPs.
- Zoning is similar to LUN masking, which is commonly used for permission management. Usually, LUN masking is performed at the SP or server level.
- WWPN (World Wide Port Name) is a globally unique identifier for a port.
- Port ID (or port address) enables routing. FC switches assign the port ID when the device logs in to the fabric.
- When N-Port ID Virtualization (NPIV) is used, a single FC HBA port (N-port) can register with the fabric by using several WWPNs.
- active-active - access to the LUNs simultaneously through all the storage ports that are available, without significant performance degradation.
- active-passive - one port is actively providing access to a given LUN. The other ports act as backup
- Disk shares are relevant only within a given ESX/ESXi host.
- Virtual machine I/O might be delayed for up to sixty seconds while path failover takes place. I/O delays might be longer on active-passive arrays.
 - On virtual machines running Microsoft Windows, increase the value of the SCSI TimeoutValue parameter to 60.
- Only one VMFS volume per LUN.
- Unless you are using diskless servers, do not set up the diagnostic partition on a SAN LUN.
- ESX/ESXi does not support FC connected tape devices.
- You cannot use virtual machine logical-volume manager software to mirror virtual disks. Dynamic disks on a Microsoft Windows virtual machine are an exception, but require special configuration.
- You should not mix FC HBAs from different vendors in a single server.
- Use a dedicated SCSI adapter for any tape drives that you are connecting to an ESX/ESXi system.
- You should not use boot from SAN in the following situations:
 - If you are using Microsoft Cluster Service.
 - If I/O contention might occur between the service console and VMkernel.
- Proper LUN masking is critical in boot from SAN mode.
- Runtime Name - the name of the first path to the device. Created by the host. Is not a reliable identifier for the device, and is not persistent.
 - `vmhba#:C#:T#:L#`, where:
 - `vmhba#` is the name of the storage adapter
 - `C#` is the storage channel number.
 - `T#` is the target number.
 - `L#` is the LUN number
- If a target has only one LUN, the LUN number is always zero (0).
- On ESXi, it is not possible to rescan a single storage adapter.
- You can modify the `Disk.MaxLUN` parameter to improve LUN discovery speed.
- You cannot discover LUNs with a LUN ID number that is greater than 255.
- You can disable the default sparse LUN support to decrease the time ESX/ESXi needs to scan for LUNs.
- The sparse LUN support enables the VMkernel to perform uninterrupted LUN scanning when a storage system presents LUNs with nonsequential LUN numbering.
- NPIV enables a single FC HBA port to register several unique WWNs with the fabric, each of which can be assigned to an individual virtual machine.
- The virtual machine's configuration file (.vmx) is updated to include a WWN pair (consisting of a World Wide Port Name and a World Wide Node Name).
- If NPIV is enabled, four WWN pairs (WWPN & WWNN) are specified for each virtual machine at creation time. All physical paths must be zoned to the virtual machine.
- NPIV can only be used for virtual machines with RDM disks. Physical HBAs, must have access to all LUNs that are to be accessed by virtual machines running on that host.
- By default, the host performs a periodic path evaluation every 5 minutes causing any unclaimed paths to be claimed by the appropriate MPP.

- Make sure read/write caching is enabled.
- Dynamic load balancing is not currently supported with ESX/ESXi.
- Path thrashing only occurs on active-passive arrays

Appendix A – Multipathing Checklist

Appendix B – Managing Storage Paths and Multipathing Plugins

- Claim rules indicate which multipathing plugin, the NMP (Native MP) or any third-party MPP, manages a given physical path.
- List claim rules `esxcli corestorage claimrule list`
- To list all multipathing modules: `vicfg-mpath --server <server> --list-plugins`
- List all VMware SATPs: `esxcli nmp satp list`
- List all storage devices: `esxcli nmp device list`

iSCSI SAN Configuration Guide

- There is no mention of requiring a Service Console connection for iSCSI anymore.
- Virtual SCSI controllers - BusLogic Parallel, LSI Logic Parallel, LSI Logic SAS, and VMware Paravirtual.
- iSCSI Name identifies a particular iSCSI element. The iSCSI name can use either IQN or EUI format.
 - IQN (iSCSI qualified name) - can be up to 255 characters long and has the following format: `iqn.yyyy-mm.naming-authority:unique_name`
 - EUI (extended unique identifier) - takes the form `eui.<16 hex digits>`
- iSCSI aliases - not unique, and are intended to be just a friendly name to associate with the node.
- You must enable your software iSCSI initiator so that ESX can use it to access iSCSI storage.
- Dynamic Discovery - Also known as Send Targets discovery responds by supplying a list of available targets to the initiator. The names and IP addresses of these targets appear on the Static Discovery tab. If you remove a static target added by dynamic discovery, the target might be returned to the list the next time a rescan happens, the HBA is reset, or the host is rebooted.
- Static Discovery - The initiator does not have to perform any discovery.
- Dynamic discovery obtains a list of accessible targets from the iSCSI storage system, while static discovery can only try to access one particular target by target name.
- You cannot change the IP address, DNS name, or port number of an existing Send Targets server. To make changes, delete the existing server and add a new one.
- To protect the integrity of iSCSI headers and data, the iSCSI protocol defines error correction methods known as header digests and data digests. Both parameters are disabled by default, but you can enable them.
- Check the end-to-end, noncryptographic data integrity beyond the integrity checks that other networking layers provide.
- Enabling header and data digests does require additional processing for both the initiator and the target. Intel Nehalem processors offload the iSCSI digest calculations.
- Use the `esxcli` command to connect the VMkernel ports to the software iSCSI initiator.
- Jumbo Frames up to 9kB (9000 Bytes) are supported.
- You cannot change the IP address, DNS name, iSCSI target name, or port number of an existing target. To make changes, remove the existing target and add a new one.
- iSCSI requires that all devices on the network implement Challenge Handshake Authentication Protocol (CHAP), which verifies the legitimacy of initiators that access targets on the network. (**EDIT** - I don't think its "required")
- ESX/ESXi supports one-way CHAP for both hardware and software iSCSI, and mutual CHAP for software iSCSI only.
- For software iSCSI only, you can set one-way CHAP and mutual CHAP for each initiator or at the target level.
- Hardware iSCSI supports CHAP only at the initiator level.
- For software iSCSI, the CHAP name should not exceed 511 and the CHAP secret 255 alphanumeric characters.
- For hardware iSCSI, the CHAP name should not exceed 255 and the CHAP secret 100 alphanumeric characters.
- Boot from a SAN - ensure that the LUN is presented to the ESX system as LUN 0. The host can also boot from LUN 255.
- Perform a rescan each time you make one of the following changes:
 - Create new LUNs on a SAN.
 - Change the path masking on a host.
 - Reconnect a cable.
 - Make a change to a host in a cluster.
 - Change CHAP settings or add new discovery addresses.
- If you notice unsatisfactory performance for your software iSCSI LUNs, you can change their maximum queue depth by using the `vicfg-module` command.
- The `iscsi_max_lun_queue` parameter is used to set the maximum outstanding commands, or queue depth, for each LUN accessed through the software iSCSI adapter. The default is 32, and the valid range is 1 to 255.
- Setting the queue depth higher than the default can decrease the total number of LUNs supported.

Appendix A - iSCSI SAN Configuration Checklist

- Set the following Advanced Settings for the ESX/ESXi host:
 - Set `Disk.UseLunReset` to 1
 - Set `Disk.UseDeviceReset` to 0

Appendix B - VMware vSphere Command-Line Interface

- The `resxtop` command provides a detailed look at ESX/ESXi resource use in real time.
- The `vicfg-iscsi` command allows you to configure software or hardware iSCSI on ESX/ESXi hosts, set up CHAP parameters, and set up iSCSI networking.
- Use the `vicfg-mpath` command to view information about storage devices, paths, and multipathing plugins.
- Use the `esxcli corestorage claimrule` command to manage claim rules. Claim rules determine which multipathing module should claim paths to a particular device and manage the device.
- The `vmkping` command allows you to verify the VMkernel networking configuration.

Appendix C - Managing Storage Paths and Multipathing Plugins (Same as Appendix B in the Fibre Channel SAN Configuration Guide)

Resource Management Guide

- The need for resource management arises from the over-commitment of resources.
- Resources include CPU, memory, power, storage, and network resources.
- The guide focuses primarily on CPU and memory. Power resource consumption can also be reduced with the Distributed Power Management (DPM) feature.
- ESX/ESXi manages network bandwidth and disk resources on a per-host basis, using network traffic shaping and a proportional share mechanism, respectively.
- Shares specify the relative priority or importance of a virtual machine (or resource pool).
- Shares are typically specified as **High**, **Normal**, or **Low** and these values specify share values with a 4:2:1 ratio. (or can set **Custom** value)
- A reservation specifies the guaranteed minimum allocation for a virtual machine.
- The reservation is expressed in concrete units (megahertz or megabytes).
- Reservation defaults to 0.
- Limit specifies an upper bound for CPU or memory resources that can be allocated to a virtual machine. It never gets more than this.
- A limit is expressed in concrete units (megahertz or megabytes).
- CPU and memory limit default is unlimited.
- Using limits can be beneficial if you want to manage user expectations, but might waste idle resources.
- Expandable Reservation defines whether reservations are considered during admission control.
- Overhead Reservation is the amount of the "Reserved Capacity" field that is being reserved for virtualization overhead.
- Worst Case Allocation is the amount of (CPU or memory) resource that is allocated to the virtual machine based on user-configured resource allocation policies (for example, reservation, shares and limit), and with the assumption that all virtual machines in the cluster consume their full amount of allocated resources.
- Admission Control - If enough unreserved CPU and memory are available, or if there is no reservation, the virtual machine is powered on. Otherwise, an Insufficient Resources warning appears.
- ESX/ESXi cannot enable hyperthreading on a system with more than 32 physical cores, because ESX/ESXi has a logical limit of 64 CPUs.
- Hyperthreaded Core Sharing Modes:
 - Any - can freely share cores with other virtual CPUs.
 - None - each virtual CPU should always get a whole core to itself, with the other logical CPU on that core being placed into the halted state.
 - Internal - cannot share cores with vCPUs from other virtual machines. Can share cores with the other virtual CPUs from the same virtual machine. Only for SMP virtual machines.
- For the best performance, when you use manual affinity settings, include at least one additional physical CPU in the affinity setting to allow at least one of the virtual machine's threads to be scheduled at the same time as its virtual CPUs.
- Dynamic Voltage and Frequency Scaling (DVFS) - You can configure your hosts to dynamically switch CPU frequencies based on workload demands.
- ESX/ESXi uses at least 50MB of system memory for the VMkernel. This is not configurable. It depends on the number and type of PCI devices. An ESXi host uses additional system memory for management agents.
- The service console typically uses 272MB.
- Memory activity is monitored to estimate the working set sizes for a default period of 60 seconds.
- ESX/ESXi charges more for idle memory than for memory that is in use. This is done to help prevent virtual machines from hoarding idle memory.
- Hosts can reclaim memory from virtual machines using:
 - Memory balloon driver (`vmmemctl`) - collaborates with the server to reclaim pages that are considered least valuable by the guest operating system. Closely matches the behavior of a native system under similar memory constraints. Causes the guest to use its own native memory management algorithms. You must configure the guest operating system with sufficient swap space.
 - Swap Files - hosts use swapping to forcibly reclaim memory from a virtual machine when the `vmmemctl` driver is not available or is not responsive. You must reserve swap space for any unreserved virtual machine memory (the difference between the reservation and the configured memory size) on per-virtual machine swap files.

- If you are overcommitting memory, to support the intra-guest swapping induced by ballooning, ensure that your guest operating systems also have sufficient swap space. This guest-level swap space must be greater than or equal to the difference between the virtual machine's configured memory size and its Reservation.
- Many workloads present opportunities for sharing memory across virtual machines.
- To determine the effectiveness of memory sharing use `resxtop` or `esxtop` to observe the actual savings. The `PSHARE` field of the interactive mode in the Memory page.
- You measure guest physical memory using the Memory Granted metric (for a virtual machine) or Memory Shared (for an ESX/ESXi host). To measure machine memory, however, use Memory Consumed (for a virtual machine) or Memory Shared Common (for an ESX/ESXi host).
- The VMkernel maps guest physical memory to machine memory.
- Multiple regions of guest physical memory might be mapped to the same region of machine memory (in the case of memory sharing) or specific regions of guest physical memory might not be mapped to machine memory (when the VMkernel swaps out or balloons guest physical memory)
- Resource Pool Hierarchy can have Parents, Children, and Siblings.
- Resource Pool Admission Control - Before you power on a virtual machine or create a resource pool, check the CPU Unreserved and Memory Unreserved fields in the resource pool's Resource Allocation tab to determine whether sufficient resources are available.
- A **group power on** will power on multiple virtual machines at the same time.
- VMotion does not support raw disks or migration of applications clustered using Microsoft Cluster Service (MSCS).
- Other VMware products or features, such as VMware vApp and VMware Fault Tolerance, might override the automation levels of virtual machines in a DRS cluster.
- An affinity rule specifies that two or more virtual machines be placed on the same host. An anti-affinity DRS rule is limited to two virtual machines,
- If two rules conflict, the older one will take precedence, and the newer rule is disabled.
- Disabled rules are ignored. DRS gives higher precedence to preventing violations of anti-affinity rules than violations of affinity rules.
- When a host machine is placed in standby mode, it is powered off.
- Hosts are placed in standby mode by the VMware DPM feature
- A cluster becomes overcommitted (yellow) when the cluster does not have the capacity to support all resources reserved by the child resource pools. Typically this happens when cluster capacity is suddenly reduced.
- A cluster enabled for DRS becomes invalid (red) when the tree is no longer internally consistent, that is, resource constraints are not observed.
- VMware DPM can use one of three power management protocols
 - IPMI - Intelligent Platform Management Interface
 - iLO - Hewlett-Packard Integrated Lights-Out
 - WOL - Wake-On-LAN
- If a host supports multiple protocols, they are used in the following order: IPMI, iLO, WOL.
- The VMotion NIC on each host must support WOL to use that protocol.
- The DRS threshold and the VMware DPM threshold are essentially independent. You can differentiate the aggressiveness of the migration and host-power-state recommendations.
- Verify that DPM is functioning properly by viewing each host's **Last Time Exited Standby** information.
- The most serious potential error you face when using VMware DPM is the failure of a host to exit standby mode when its capacity is needed by the DRS cluster. Use the preconfigured **Exit Standby Error** alarm for this error.
- DRS Recommendations have 5 levels (1-5). Priority 1, the highest, indicates a mandatory move because of a host entering maintenance or standby mode or DRS rule violations. Other priority ratings denote how much the recommendation would improve the cluster's performance;
- Prior to ESX/ESXi 4.0, recommendations received a star rating (1 to 5 stars) instead of a priority level.
- Non-Uniform Memory Access (NUMA) systems are advanced server platforms with more than one system bus.
- Some virtual machines are not managed by the ESX/ESXi NUMA scheduler: if you manually set the processor affinity for a virtual machine, or virtual machines that have more virtual processors than the number of physical processor cores available on a single hardware node.
- When a virtual machine is powered on, ESX/ESXi assigns it a home node. This is initially assigned to home nodes in a round robin fashion.
- Periodically (every two seconds by default), the system examines the loads of the various nodes and determines if it should rebalance the load by moving a virtual machine from one node to another.
- Transparent page sharing has also been optimized for use on NUMA systems.
- The `VMkernel.Boot.sharePerNode` option controls whether memory pages can be shared (de-duplicated) only within a single NUMA node or across multiple NUMA nodes. It is turned on by default.
- If you turn off the option, identical pages can be shared across different NUMA nodes. In memory-constrained environments, such as VMware View this could be very beneficial.
- The systems that offer a NUMA platform include AMD CPUs or the IBM Enterprise X-Architecture.
- You must manually select the boxes for all processors in the NUMA node. CPU affinity is specified on a per-processor, not on a per-node, basis.
- Specify nodes to be used for future memory allocations only if you have also specified CPU affinity.

Appendix A - Performance Monitoring Utilities: `resxtop` and `esxtop`

- The `esxtop` utility reads its default configuration from `.esxtop4rc`.
- Do not edit the `.esxtop4rc` file. Instead, select the fields and the order in a running `esxtop` process, make changes, and save this file using the `W` interactive command.

Appendix B – Advanced attributes

vSphere Availability Guide

- The first five hosts added to the cluster are designated as primary hosts, and all subsequent hosts are designated as secondary hosts. The primary hosts maintain and replicate all cluster state and are used to initiate failover actions. If a primary host is removed from the cluster, VMware HA promotes another host to primary status.
- One of the primary hosts is also designated as the active primary host and its responsibilities include:
 - Deciding where to restart virtual machines.
 - Keeping track of failed restart attempts.
 - Determining when it is appropriate to keep trying to restart a virtual machine.
- If the active primary host fails, another primary host replaces it.
- If a host stops receiving heartbeats from all other hosts in the cluster for more than 12 seconds, it attempts to ping its isolation addresses. If this also fails, the host declares itself as isolated from the network.
- Three types of admission control:
 - Host
 - Resource pool
 - HA
- Only VMware HA admission control can be disabled.
- Slot size is a logical representation of the memory and CPU resources that satisfy the requirements for any powered-on virtual machine in the cluster.
- The maximum Configured Failover Capacity that you can set is four.
- If your cluster contains any virtual machines that have much larger reservations than the others, they will distort slot size calculation. To avoid this, you can specify an upper bound for the CPU or memory component of the slot size by using the `das.slotCpuInMHz` or `das.slotMemInMB` advanced attributes, respectively.
- You can configure VMware HA to perform admission control by reserving a specific percentage of cluster resources for recovery from host failures.
- You can configure VMware HA to designate a specific host as the failover host.
- To ensure that spare capacity is available on the failover host, you are prevented from powering on virtual machines or using VMotion to migrate virtual machines to the failover host.
- When choosing an admission control policy, you should consider a number of factors:
 - Avoiding Resource Fragmentation - The Host Failures Cluster Tolerates policy avoids resource fragmentation by defining a slot as the maximum virtual machine reservation. The Percentage of Cluster Resources policy does not address the problem of resource fragmentation. With the Specify a Failover Host policy, resources are not fragmented because a single host is reserved for failover.
 - Flexibility of Failover Resource Reservation - The Host Failures Cluster Tolerates policy allows you to set the failover level from one to four hosts. The Percentage of Cluster Resources policy allows you to designate up to 50% of cluster resources for failover. The Specify a Failover Host policy only allows you to specify a single failover host.
 - Heterogeneity of Cluster - In a heterogeneous cluster, the Host Failures Cluster Tolerates policy can be too conservative because it only considers the largest virtual machine reservations when defining slot size and assumes the largest hosts fail when computing the Current Failover Capacity. The other two admission control policies are not affected by cluster heterogeneity.
- VMware HA Attributes:

Attribute	Description
<code>das.isolationaddress[...]</code>	Address to ping to determine if a host is isolated from the network. If not specified, the default gateway of the console network is used. Can specify multiple isolation addresses (up to 10).
<code>das.usedefaultisolationaddress</code>	Specifies whether or not this default is used (true false).
<code>das.failedetectiontime</code>	Failure detection time for host monitoring. The default is 15000 milliseconds
<code>das.failedetectioninterval</code>	Heartbeat interval among VMware HA hosts. Default is 1000 milliseconds.
<code>das.defaultfailoverhost</code>	The host that VMware HA tries to fail virtual machines over to.
<code>das.isolationShutdownTimeout</code>	Time the system waits for a VM to shut down before powering it off. Default is 300 seconds.
<code>das.slotMemInMB</code>	The maximum bound on the memory slot size.
<code>das.slotCpuInMHz</code>	The maximum bound on the CPU slot size.
<code>das.vmMemoryMinMB</code>	Memory resource value assigned to a VM if it's not specified or zero. Default is 0 MB.
<code>das.vmCpuMinMHz</code>	Default CPU resource value assigned to a VM if it's not specified or zero. Default is 256MHz.
<code>das.iostatsInterval</code>	I/O stats interval for VM monitoring sensitivity. Default is 120 (seconds).

- If you change the value of any of the following advanced attributes, you must disable and then re-enable VMware HA before your changes take effect.
 - `das.isolationaddress[...]`
 - `das.usedefaultisolationaddress`
 - `das.failedetectiontime`

- das.failedetectioninterval
- das.isolationShutdownTimeout
- On ESX hosts, HA communications travel over service console networks.
- On ESXi hosts, HA communications travel over VMkernel networks.
- HA needs and automatically opens the following firewall ports:
 - Incoming port: TCP/UDP 8042-8045
 - Outgoing port: TCP/UDP 2050-2250
- VMware Fault Tolerance provides continuous availability for virtual machines by creating and maintaining a Secondary VM that is identical to, and continuously available to replace, the Primary VM in the event of a failover situation.
- A fault tolerant VM and its secondary copy are not allowed to run on the same host. Fault Tolerance uses anti-affinity rules.
- Fault Tolerance prerequisites:
 - VMware HA must be enabled on the cluster. Host Monitoring should also be enabled.
 - Host certificate checking must be enabled for all hosts.
 - Each host must have a VMotion and a Fault Tolerance Logging NIC configured.
 - Hosts must have the same version and patch level.
 - Hosts must have processors from the FT-compatible processor group.
 - Hosts must have Hardware Virtualization (HV) enabled in the BIOS.
 - Virtual machines must be stored in virtual RDM or VM disk (VMDK) files that are thick provisioned with the Cluster Features option.
 - Virtual machines must be running on one of the supported guest operating systems.
- Not supported for fault tolerant virtual machines.

○ Snapshots.	○ SMP virtual machines	○ NPIV
○ Storage VMotion	○ Physical RDMs	○ NIC passthrough
○ DRS features	○ Paravirtualized guests	○ EPT/RVI
- You should have no more than four fault tolerant virtual machines (primaries or secondaries) on any single host.
- The recommendation is that you use a maximum of 16 virtual disks per fault tolerant virtual machine.

Appendix – Fault Tolerance Error Messages

vSphere Web Access Administrator's Guide

- Supported Operating Systems:
 - Windows 2003 SP1, XP Pro SP3, XP Home SP2, 2000 SP4
 - Linux with standard libraries, but requiring GTK+ 2.
- Supported browsers:
 - IE 6, 7 or later
 - Firefox 2, 3.0 or later
- The console webAccess service is now `vmware-webAccess`
- The Alarms tab is available only when you use vSphere Web Access to connect to vCenter Server.
- You can view the assigned tasks for a virtual machine, but you cannot assign tasks using Web Access.
- New disk policy (not available option from vSphere Client)
 - Optimize for safety – Saves all changes to the virtual disk before notifying the system.
 - Optimize for performance – Acknowledges changes to the virtual disk immediately, but saves them at a later time.
- VMI currently supports only 32-bit guests.

Additional Resources

Setup for Failover Clustering and Microsoft Cluster Service

- Support for Windows Server 2003 SP2, 2000 SP4, or 2008.
- Only two-node clustering.

Storage Type	Cluster in a Box	Cluster Across Boxes	Standby Host (n+1)
Virtual disks	Yes (recommended)	No	No
Pass-through RDM (physical compatibility mode)	No	Yes (recommended)	Yes
Non-pass-through RDM (virtual compatibility mode)	Yes	Yes	No

- Clusters across physical machines with non-pass-through RDM is supported only for clustering with Windows 2000 Server or Windows Server 2003. It is not supported for clustering with Windows Server 2008.
- Not supported:
 - Clustering on iSCSI or NFS disks.
 - Clustered virtual machines as part of VMware clusters (DRS or HA).
 - MSCS in conjunction with VMware Fault Tolerance.

- Migration with VMotion.
- N-Port ID Virtualization (NPIV)
- With native multipathing (NMP), clustering is not supported when the path policy is set to round robin.
- You must use hardware version 7 with ESX4.
- You can put the boot disk of a virtual machine on a SAN-based VMFS volume.
- Clustered continuous replication (CCR) environment for Microsoft Exchange - Use physical compatibility mode RDMs.
- For boot disks - Select "Support clustering features such as Fault Tolerance" to create a disk in eagerzeroedthick format.
- SCSI Controller Type:
 - Windows 2000 and 2003 - Server LSI Logic Parallel (download from the LSI web site).
 - Windows 2008 - LSI Logic SAS
- Cluster across boxes:
 - Shared storage must be on an FC SAN.
 - RDM in physical compatibility (pass-through) or virtual compatibility (non-pass-through) mode. VMware recommends physical compatibility mode. The cluster cannot use virtual disks for shared storage.
 - Failover clustering with Windows Server 2008 is not supported with virtual compatibility mode
- Standby Host (n+1): Use only a single physical path from the host to the storage arrays.

vSphere Command-Line Interface Installation and Reference Guide

○

License Server Configuration for vCenter Server 4.0 – ESX 3.x licensing

- Using vCenter Server 4.0 to manage ESX 3.x/ESXi 3.5 hosts, you have several options:
 - A single license server for vCenter Server 4.0 and the ESX 3.x/ESXi 3.5 hosts.
 - One license server for vCenter Server 4.0, and use another license server for the ESX 3.x/ESXi 3.5 hosts.
 - A license server for vCenter Server 4.0, and use host-based licensing for the ESX 3.x/ESXi 3.5 hosts.
 - Do not use a license server. Upgrade all of your hosts to ESX 4.0/ESXi 4.0.

ESX4 Patch Management Guide

- **Bulletin** - grouping of one or more VIBs (vSphere Installation Bundle).
- **Depot** - logical grouping of VIBs – online.
- **Offline Bundle zip** - archive that encapsulates VIBs.
- **Patch** - bulletin that groups one or more VIBs together to address a particular issue or enhancement.
- **Roll-up** - collection of patches that is grouped for ease of download and deployment.
- **Update** - periodic release of an ESX image.
- **VIB** - single software package.
- A record of each installed bulletin is written to the `/etc/vmware/esxupdate`
- Four basic modes of `esxupdate`:
 - **Inspection mode**
 - `esxupdate query` - display a list of bulletins installed.
 - `esxupdate info` - display information on the contents of one or more bulletins.
 - **Scan mode** - determines which bulletins are applicable.
 - **Test mode** - go through all installation operations without installing.
 - **Update mode** - installs bulletins.
- The installation process is recorded in the `esxupdate.log` file. By default, this file is located in the `/var/log/vmware` directory.
- `esxupdate` never reboots your host.

Optional vSphere Products and Modules

vCenter Update Manager Administration Guide

○

vCenter Converter Administration Guide

○

vCenter Orchestrator 4.0 – Installation and Configuration Guide

- vCenter Orchestrator is a development and process-automation platform that provides a library of extensible workflows.

- Orchestrator exposes every operation in the vCenter Server API.
- Three global user roles: Administrators, Developers, and End Users.
- Support for several OS:
 - Server 2003 R2, 32bit, 64bit
 - Server 2008, 32bit, 64bit
- Supported directory service types:
 - Active Directory
 - OpenLDAP
 - eDirectory
 - Sun Java Directory Server
- Supported browsers:
 - IE 6.0 and 7.0
 - Firefox 3.0.6 or later
 - Safari 3.x (experimental)
- The database is separate from the standard vCenter database.
 - Microsoft SQL Server 2005 Enterprise (SP2) and x64 and (SP1)
 - Microsoft SQL Server 2005 Standard (SP2), (SP1)
 - Microsoft SQL Server 2000 Enterprise (SP4)
 - Microsoft SQL Server 2000 Standard (SP4)
 - Oracle 10g Enterprise Release 2 (10.2.0.3.0) x32 and x64
- Orchestrator does not support IPv6 operating systems.
- Orchestrator Configuration Service startup type is set to Manual by default.
- `http://<computer_DNS_name_or_IP_address>:8282`.
- HTTPS connection through port 8283, you must configure Jetty to use SSL.
- You cannot change the vmware default user name.
- Orchestrator Default Ports:

Purpose	Port	Description
Lookup	8230	The main port to communicate with the Orchestrator server (JNDI port).
Command	8240	The application communication port (RMI container port) used for remote invocations
Messaging	8250	The Java messaging port used to dispatch events.
Data	8244	The port used to access all Orchestrator data models, such as workflows and policies.
HTTP server	8280	The port for the HTTP connector used to connect to the Web front-end.
HTTPS server	8281	To connect to the Web front-end and to communicate with vCenter API.
HTTP access	8282	The access port for the Web UI of Orchestrator configuration.
HTTPS access	8283	The SSL access port for the Web UI of Orchestrator configuration. To enable configure Jetty to use SSL.
LDAP	389	The look up port of your LDAP Authentication server.
LDAP using SSL	636	The look up port of your secure LDAP Authentication server.
PostgreSQL	5432	PostgreSQL Server for Orchestrator database.
SQL Server	1433	Microsoft SQL Server for Orchestrator database.
Oracle	1521	Oracle Database for Orchestrator database.
MySQL	3306	MySQL for Orchestrator database.
SMTP Server	25	Email notifications.

- The Orchestrator configuration interface uses a secure connection to communicate with vCenter. You can import the required SSL certificate from a URL or file.
- The users in the Developers role have editing privileges on all elements.
- Plug-in file extensions are .vmoapp (can contain several .dar files) and .dar (contains all the resources associated with one plug-in).
- You must import the vCenter Server license. The set of plug-ins delivered with Orchestrator do not require a license

vCenter Orchestrator 4.0 – Administration Guide

○

VMware Consolidated Backup - Virtual Machine Backup Guide

○