Here are my condensed notes for the vSphere 5.0 documentation. They're excerpts taken directly from VMware's own official documentation. These notes aren't meant to be comprehensive, or for a beginner; just my own personal notes. They're snippets I found interesting while reviewing the official VMware documentation. I hope you find them useful.

VMware vSphere Basics (ESXi 5.0 & vCenter 5.0)

- Virtualization Layer — includes infrastructure services and application services.
  - Infrastructure services:
    - Compute services
    - Storage services
    - Network services
  - Application services are the set of services provided to ensure availability, security, and scalability for applications. Examples include vSphere High Availability and Fault Tolerance.
- In addition to the Virtualization Layer, there is also the Management Layer and the Interface Layer.
- VMware vSphere Web Client — A Web interface that enables users to connect remotely to vCenter Server from a variety of Web browsers and operating systems.
- VMware Storage DRS — Allocates and balances storage capacity and I/O dynamically across collections of datastores. This feature includes management capabilities that minimize the risk of running out of space and the risk of I/O bottlenecks slowing the performance of virtual machines.
- A datastore cluster is an aggregation of multiple datastores into a single logical, load-balanced pool.
- Storage DRS helps you manage multiple datastores as a single compute resource, called a datastore cluster. A datastore cluster is an aggregation of multiple datastores into a single logical, load-balanced pool. You can treat the datastore cluster as a single flexible storage resource for resource management purposes. In effect, a datastore cluster is the storage equivalent of an ESXi compute cluster. You can dynamically populate datastore clusters with datastores of similar characteristics. You can assign a virtual disk to a datastore cluster and Storage DRS finds an appropriate datastore for it. The load balancer manages initial placement and future migrations based on workload measurements. Storage space balancing and I/O balancing minimize the risk of running out of space and the risk of I/O bottlenecks slowing the performance of virtual machines.
- Network resource pools determine the priority different network traffic types are given on a vDS. When network resource management is enabled, vDS traffic is divided into the following network resource pools: FT traffic, iSCSI traffic, vMotion traffic, management traffic, NFS traffic, and virtual machine traffic (*Forbes*: There are now 2 newly defined pools, see the Networking sections for details). You can control the priority for the traffic from each of these network resource pools by setting the physical adapter shares and host limits for each network resource pool.
- VMware vShield is a suite of security virtual appliances built to work with vSphere.
  - *vShield Zones* — provides firewall protection for traffic between virtual machines.
  - *vShield Edge* — provides network edge security and gateway services to isolate the virtual machines in a port group, vDS port group, or Cisco Nexus 1000V.
  - *vShield App* — an interior, vNIC-level firewall that allows you to create access control policies regardless of network topology.
  - *vShield Endpoint* — delivers an introspection-based antivirus solution.
- Administrative functions are available to the vSphere user through the vSphere Client, with full functionality, and the vSphere Web Client, with a subset of that functionality.
- vSphere Client:
  - For infrastructure configuration and day-to-day operations:
    - Locally installed application
    - Windows operating system only
    - Can connect to vCenter Server or directly to hosts
    - Full range of administrative functionality
  - Users: Virtual infrastructure administrators for specialized functions
- vSphere Web Client:
  - For day-to-day operations
    - Web application
    - Cross platform
    - Can connect to vCenter Server only
    - Subset of full functionality, focused on virtual machine deployment as well as basic monitoring functions. Cannot configure hosts, clusters, networks, or datastores
    - Extensible plug-in-based architecture
  - Users: Virtual infrastructure administrators, helpdesk, network operations center operators, virtual machine owners

vSphere Installation and Setup (vSphere 5.0, ESXi 5.0 & vCenter 5.0)

- The following changes from vSphere 4.x affect vSphere installation and setup:
  - Service Console is removed
  - ESXi does not have a graphical installer
  - vSphere Auto Deploy and vSphere ESXi Image Builder CLI — With ESXi 5.0, you can load an ESXi image directly into memory by using vSphere Auto Deploy. You can provision and reprovision large numbers of ESXi hosts efficiently with vCenter Server, and manage ESXi updates and patching by using an image profile. You can save host configuration such as network or storage setup as a host profile and apply it to the host by using Auto Deploy. You can use ESXi Image Builder CLI to create ESXi installation images with a customized set of updates, patches, and drivers.
  - Changes in the ESXi installation and upgrade process — The `vihostupdate` and `esxupdate` utilities are not supported for ESXi 5.0. You cannot upgrade or migrate to ESXi 5.0 by using any command-line utility. After you have upgraded or migrated to ESXi 5.0, you can upgrade or patch ESXi 5.0 hosts using vCLI `esxcli` commands. After you upgrade or migrate your host to ESXi 5.0, you cannot roll back to your version 4.x ESX or ESXi software.
  - Installer caching — Instead of using a binary image to install the system, whatever bits were used at boot time are cached to the system. This caching reduces installation problems caused by accessing installation files across networks that are under load. Scripted installations cannot PXE boot a server and then obtain the binary image from some other form of media.
  - Changes to partitioning of host disks — All freshly installed hosts in vSphere 5.0 use the GUID Partition Table format instead of the MSDOS-style partition label. This change supports ESXi installation on disks larger than 2TB. Unlike earlier versions, ESXi 5.0 does not create VMFS partitions in second and successive disks.
  - VMware vCenter Server Appliance — As an alternative to installing vCenter Server on a Windows machine, vSphere 5.0 provides the VMware vCenter Server Appliance. The vCenter Server Appliance is a preconfigured Linux-based virtual machine optimized for running vCenter Server and associated services.
  - vSphere Web Client — The vSphere Web Client is a server application that provides a browser-based alternative to the traditional vSphere Client. You can use a Web browser to connect to the vSphere Web Client to manage an ESXi host through a vCenter Server.
- Interactive installations are recommended for small deployments of fewer than five hosts.
- You boot the installer from a CD or DVD, from a bootable USB device, or by PXE booting the installer from a location on the network.
- Auto Deploy does not store the ESXi state on the host disk. vCenter Server stores and manages ESXi updates and patching through an image profile, and, optionally, the host configuration through a host profile.
- You can use ESXi Image Builder CLI to create ESXi installation images with a customized set of updates, patches, and drivers.
- ESXi Image Builder CLI is a PowerShell CLI command set that you can use to create an ESXi installation image with a customized set of ESXi updates and patches. You can also include third-party network or storage drivers that are released between vSphere releases.
- You can deploy an ESXi image created with Image Builder in either of the following ways:
  - Burning it to an installation DVD.
  - Through vCenter Server, using the Auto Deploy feature.
- The installable version of ESXi is always installed in evaluation mode. ESXi Embedded is preinstalled on an internal USB device by your hardware vendor. It might be in evaluation mode or prelicensed.
- The following boot media are supported for the ESXi installer:
  - Boot from a CD/DVD.
  - Boot from a USB flash drive.
  - PXE boot from the network.
  - Boot from a remote location using a remote management application.
- You can customize the standard ESXi installer ISO image with your own installation or upgrade script.
- Modify the boot.cfg file to specify the location of the installation or upgrade script using the `kernelopt` option.
- TFTP traffic uses UDP port 69.
- ESXi install location must be at least 5GB if you install the components on a single disk.
- The root password must contain between 6 and 64 characters.
- Hardware and system resources to install and use ESXi 5.0:
  - Check HCL link
  - ESXi 5.0 will install and run only on servers with 64-bit x86 CPUs.
  - ESXi 5.0 requires a host machine with at least two cores.
  - 2GB RAM minimum.
  - SATA disks will be considered remote, not local. These disks will not be used as a scratch partition by default because they are seen as remote.
- You cannot connect a SATA CD-ROM device to a virtual machine on an ESXi 5.0 host. To use the SATA CD-ROM device, you must use IDE emulation mode.
- VMware Auto Deploy requires the legacy BIOS firmware and is not available with UEFI.

- Changing the boot type from legacy BIOS to UEFI after you install ESXi 5.0 might cause the host to fail to boot. In this case, the host displays an error message similar to: *Not a VMware boot bank*. Changing the host boot type between legacy BIOS and UEFI is not supported after you install ESXi 5.0.
- vCenter min requirements:
  - Two 64-bit CPUs or one 64-bit dual-core processor (Itanium (IA64) is not supported)
  - 4GB of RAM
  - 4 GB of disk space — In vCenter Server 5.0, the default size for vCenter Server logs is 450MB larger than in vCenter Server 4.x.
- vCenter Server Appliance:
  - At least 7GB, and a maximum of 80GB
  - Very small inventory (10 or fewer hosts, 100 or fewer virtual machines): at least 4GB
  - Small inventory (fewer than 100 hosts or 1000 virtual machines). At least 8GB
  - Medium inventory (100-400 hosts or 1000-4000 virtual machines). At least 12GB
  - Large inventory (More than 400 hosts or 4000 virtual machines). At least 16GB
- vSphere client — 1CPU and 1GB RAM
- The recommended disk sizes assume default log levels. If you configure more detailed log levels, more disk space is required.
- Medium Deployment of Up to 50 Hosts and 500 Powered-On Virtual Machines:

Product	Cores	Memory	Disk
vCenter server	2	4GB	5GB
vSphere client	1	200MB	1.5GB

- Large Deployment of Up to 300 Hosts and 3,000 Powered-On Virtual Machines:

Product	Cores	Memory	Disk
vCenter server	4	8GB	10GB
vSphere client	1	500MB	1.5GB

- Extra-Large Deployment of Up to 1,000 Hosts and 10,000 Powered-On Virtual Machines:

Product	Cores	Memory	Disk
vCenter server	8	16GB	10GB
vSphere client	2	500MB	1.5GB

- vCenter Server requires a 64-bit operating system (*Forbes*: currently Windows 2003 SP2, 2003 R2 SP1, 2008 SP2 or 2008 R2), and the 64-bit system DSN is required for vCenter Server to connect to its database.
- vCenter Server requires the Microsoft .NET 3.5 SP1 Framework.
- If you plan to use the Microsoft SQL Server 2008 R2 Express database that is bundled with vCenter Server, Microsoft Windows Installer version 4.5 (MSI 4.5) is required on your system.
- The VMware vCenter Server Appliance can be deployed only on hosts that are running ESX version 4.x or ESXi version 4.x or later.
- Browsers supported by the vSphere Web Client:
  - Microsoft Internet Explorer 7 and 8
  - Mozilla Firefox 3.6
  - Requires the Adobe Flash Player version 10.1.0 or later to be installed with the appropriate plug-in for your browser
- The default values for log capacity in this infrastructure vary, depending on the amount of storage available and on how you have configured system logging. Hosts that are deployed with Auto Deploy store logs on a RAM disk, which means that the amount of space available for logs is small.
- If your host is deployed with Auto Deploy, reconfigure your log storage in one of the following ways:
  - Redirect logs over the network to a remote collector.
  - Redirect logs to a NAS or NFS store.
- You do not need to reconfigure log storage for ESXi hosts that use the default configuration, which stores logs in a scratch directory on the VMFS volume. For these hosts, ESXi 5.0 autoconfigures logs to best suit your installation, and provides enough space to accommodate log messages.
- Recommended Minimum Size and Rotation Configuration for hostd, vpxa, and fdm Logs.

Log	Max log file size	Number of rotations to preserver	Min disk space required
Management Agent (hostd)	10,240KB	10	100MB
vCenter Agent (vpxa)	5,120KB	10	50MB
vSphere HA agent (fdm)	5,120KB	10	50MB

- Ports Required for Communication Between Components:

Port	Description
80	vCenter Server requires port 80 for direct HTTP connections. Port 80 redirects requests to HTTPS port 443. This redirection is useful if you accidentally use http://server instead of https://server.
389	This port must be open on the local and all remote instances of vCenter Server. This is the LDAP port number for the Directory Services for the vCenter Server group. The vCenter Server system needs to bind to port 389, even if you are not joining this vCenter Server instance to a Linked Mode group. If another service is running on this port, it might be preferable to remove it or change its port to a different port. You can run the LDAP service on any port from 1025 through 65535. If this instance is serving as the Microsoft Windows Active Directory, change the port number from 389 to an available port from 1025 through 65535.
443	The default port that the vCenter Server system uses to listen for connections from the vSphere Client. To enable the vCenter Server system to receive data from the vSphere Client, open port 443 in the firewall. The vCenter Server system also uses port 443 to monitor data transfer from SDK clients. If you use another port number for HTTPS, you must use *ip-address*:*port* when you log in to the vCenter Server system.
636	For vCenter Server Linked Mode, this is the SSL port of the local instance. If another service is running on this port, it might be preferable to remove it or change its port to a different port. You can run the SSL service on any port from 1025 through 65535.
902	The default port that the vCenter Server system uses to send data to managed hosts. Managed hosts also send a regular heartbeat over UDP port 902 to the vCenter Server system. This port must not be blocked by firewalls between the server and the hosts or between hosts.
902	Port 902 must not be blocked between the vSphere Client and the hosts. The vSphere Client uses this port to display virtual machine consoles.
8080	Web Services HTTP. Used for the VMware VirtualCenter Management Web Services.
8443	Web Services HTTPS. Used for the VMware VirtualCenter Management Web Services.
60099	Web Service change service notification port
10443	vCenter Inventory Service HTTPS
10109	vCenter Inventory Service Service Management
10111	vCenter Inventory Service Linked Mode Communication

- You use the ESXi CD/DVD or a USB flash drive to install the ESXi software onto a SAS, SATA, SCSI hard drive, or USB drive.
- Verify that the server hardware clock is set to UTC. This setting is in the system BIOS.
- When you install ESXi to a software iSCSI disk, you must configure the target iSCSI qualified name (IQN). Verify that the target IQN is configured in the iBFT BIOS target parameter setting. This setting is in the option ROM of the network interface card (NIC) to be used for the iSCSI LUN.
- The installation or upgrade script can reside in one of the following locations:
  - FTP
  - HTTP/HTTPS
  - NFS
  - USB flash drive
  - CDROM
- Scripted Installation Choices:
  - Always install on the first disk on multiple machines — Create one script.
  - Install ESXi on a different disk for each machine — Create multiple scripts.
- At boot time you might need to specify options to access the kickstart file. You can enter boot options by pressing Shift+O in the boot loader. For a PXE boot installation, you can pass options through the `kernelopts` line of the boot.cfg file.
- A ks=... option must be given, to specify the location of the installation script. Otherwise, a scripted installation or upgrade will not start. If ks=... is omitted, the text installer will proceed.
- The default ks.cfg installation script is located in the initial RAM disk at /etc/vmware/weasel/ks.cfg
- When you install ESXi using the ks.cfg script, the default root password is `mypassword`.
- The following locations are supported for the installation or upgrade script:
  - CD/DVD.
  - USB Flash drive
  - A location accessible with the following protocols: NFS, HTTP, HTTPS, FTP
- Differences Between ESXi 4.x and ESXi 5.0 Scripted Installation and Upgrade Commands — In ESXi 5.0, because the installation image is loaded directly into the host RAM when the host boots, you do not need to include the location of the installation media in the installation script.
- ESXi 5.0 supports scripted upgrades in addition to scripted installation.
- Command differences are:
  - `accepteula` or `vmaccepteula` Only in ESXi

- `autopart` Deprecated and replaced with install, upgrade, or installorupgrade.
- `auth` or `authconfig` Not supported in ESXi 5.0.
- `bootloader` Not supported in ESXi 5.0.
- `esxlocation` Deprecated and unused in ESXi.
- `firewall` Not supported in ESXi 5.0.
- `firewallport` Not supported in ESXi 5.0.
- `install`, `installorupgrade`, `upgrade` These commands replace the deprecated autopart command. Use one of these command to specify the disk to partition, and the part command to create the vmfs datastore. installorupgrade and upgrade are newly supported in ESXi5.0.
- `serialnum` or `vmserialnum` Deprecated in ESXi 5.0. You can license the host only after installation.
- `timezone` Not supported in ESXi 5.0.
- `virtualdisk` Not supported in ESXi 5.0.
- `zerombr` Not supported in ESXi 5.0.
- `%firstboot` --level option not supported in ESXi 5.0.
- `%packages` Not supported in ESXi 5.0.

- When you use a scripted upgrade to upgrade from ESX 4.x to ESXi 5.0, the MPX and VML disk names change, which might cause the upgrade to fail. To avoid this problem, use Network Address Authority Identifiers (NAA IDs) for the disk device instead of MPX and VML disk names.
- The boot loader configuration file boot.cfg specifies the kernel, the kernel options, and the boot modules that the mboot.c32 boot loader uses in an ESXi installation.
- The boot.cfg file is provided in the ESXi installer. You can modify the `kernelopt` line of the boot.cfg file to specify the location of an installation script or to pass other boot options.
- Auto Deploy enables experienced system administrators to manage large vSphere deployments efficiently. With Auto Deploy, vCenter Server loads the ESXi image directly into the host memory. Unlike the other installation options, Auto Deploy does not store ESXi state on the host disk. vCenter Server stores and manages ESXi updates and patching through an image profile, and, optionally, the host configuration through a host profile.
- You can specify the image to deploy, the hosts to deploy to, and, optionally, host profiles to apply to the hosts, and a location for each host.
- When a physical host set up for Auto Deploy is turned on, Auto Deploy uses a PXE boot infrastructure in conjunction with vSphere host profiles to provision and customize that host. No state is stored on the host itself, instead, the Auto Deploy server manages state information for each host.
- Auto Deploy Stores Information for Deployment:

Information Type	Description	Source of State Information
Image state	Executable software to run on an ESXi host.	Image profile, created with Image Builder PowerCLI.
Configuration state	Configurable settings that determine how the host is configured, for example, virtual switches and their settings, driver settings, boot parameters, and so on.	Host profile, created using the host profile UI. Often comes from a template host.
Dynamic state	Runtime state that is generated by the running software, for example, generated private keys or runtime databases.	Stored in host memory and lost during reboot.
Vitual Machine state	Virtual machines stored on a host and virtual machine auto-start information (subsequent boots only).	Managed by vCenter Server system by default. ○ If the virtual machine is in a VMware HA cluster, VMware Auto Deploy retains this information so deployment works even if the vCenter Server is unavailable. ○ If the virtual machine is not in a VMware HA cluster, vCenter Server must be available to supply information to Auto Deploy.
User input	State that is based on user input, for example, an IP address the user provides when the system starts up, cannot automatically be included in the host profile.	Custom information is stored in an answer file. You can create a host profile that requires user input for certain values. When Auto Deploy applies a host profile that requires an answer to a newly-provisioned host, the host comes up in maintenance mode. You can right-click the host and select **Update Answer File** to be prompted for the information. The answer file information is stored with the host. Each host has one answer file that can include multiple user input items.

- Auto Deploy infrastructure components:
  - *Auto Deploy server* — serves images and host profiles to ESXi hosts.
  - *Auto Deploy rules engine* — tells the Auto Deploy server which host profile to serve to which host.
  - *Image profiles* — define the set of VIBs to boot ESXi hosts with.

    - *Host profiles* — define machine-specific information such as networking or storage setup.
    - *Answer files* — store information that the user provides during the boot process. Only one answer file exists for each host.
- The rule engine maps software and configuration settings to hosts based on the attributes of the host.
- The rule engine includes rules and rule sets.
    - *Rules* **—** Rules can assign image profiles and host profiles to a set of hosts, or specify the location (folder or cluster) of a host on the target vCenter Server system. A rule can identify target hosts by boot MAC address, SMBIOS asset tag, BIOS UIID, fixed DHCP IP address, or static IP address. In most cases, rules apply to multiple hosts. You create rules using the Auto Deploy PowerCLI. After you create a rule, you must add it to a rule set. After you add a rule to a rule set, you cannot edit it.
    - *Active Rule Set* **—** When a newly started host contacts the Auto Deploy server with a request for an image, the Auto Deploy server checks the active rule set for matching rules. The image profile, host profile, and vCenter Server inventory location that are mapped by matching rules are then used to boot the host. If more than one item is mapped by the rules, the Auto Deploy server uses the item that is first in the rule set.
    - *Working Rule Set* — The working rule set allows you to test changes to rules before making them active. For example, you can use Auto Deploy PowerCLI commands for testing compliance with the working rule set to verify hosts managed by a vCenter Server system are following the rules in the working rule set. By default, commands for changing the working rule set activate the changes. Use the NoActivate parameter to add a rule only to the working rule set.
- First Boot Prerequisites:
    - Set up a DHCP server that assigns an IP address to each host upon startup and that points the host to the TFTP server to download gPXE from.
    - Identify an image profile to be used in one of the following ways.
        - Choose an ESXi image profile in a public depot.
        - Create a custom image using the Image Builder PowerCLI.
    - (Optional) define a rule that applies the host profile.
    - Specify rules for the deployment of the host and add the rules to the Auto Deploy rules engine.
- For hosts that are provisioned with Auto Deploy and managed by a vCenter Server system, subsequent boots can become completely automatic. The host is provisioned by the host's vCenter Server system, which stores the information about the image profile and host profile for each host. If the vCenter Server system is unavailable, the host contacts the Auto Deploy server for image and host profiles and the host reboots. However, Auto Deploy cannot set up vSphere distributed switches if vCenter Server is unavailable, and Auto Deploy can assign virtual machines only to hosts that participate in an HA cluster(*Forbes*: not DRS?). Until the host is reconnected to vCenter Server and the host profile is applied, the switch cannot be created and, because the host is in maintenance mode, virtual machines cannot start.
- Prerequisites to auto deploy:
    - Obtain the vCenter Server installation media, which include the Auto Deploy installer, or deploy the vCenter Server Appliance.
    - A DHCP server is included in the vCenter Server on Linux virtual appliance.
    - A TFTP server is included in the vCenter Server on Linux virtual appliance.
    - Auto Deploy transfers data over SSL to prevent casual interference and snooping. However, the authenticity of the client or the Auto Deploy server is not checked during a PXE boot.
- Install VMware Auto Deploy on a vCenter Server system or a Windows system, or deploy the vCenter Server on Linux virtual appliance to an ESXi system of your choice.
- You can set up bulk licensing using PowerCLI commands. Bulk licensing works for all ESXi hosts but is especially useful for hosts provisioned with Auto Deploy.
- The Auto Deploy PowerCLI cmdlets allow you to create rules that associate hosts with image profiles, host profiles, and location on the vCenter Server target. You can also test rule compliance and repair compliance issues.
- When you add a rule to the Auto Deploy rule set or make changes to one or more rules, unprovisioned hosts you boot are automatically provisioned according to the new rules. For all other hosts, the new rules apply only when you test their rule compliance and perform remediation.
- Provision a Host (First Boot):
    - The host contacts the DHCP server and downloads gPXE from the location the server points it to. Next, the Auto Deploy server provisions the host with the image specified by the rule engine. The Auto Deploy server might also apply a host profile to the host if one is specified in the rule set. Finally, VMware Auto Deploy adds the host to the vCenter Server system that is specified in the rule set.
    - (Optional) If a host profile is applied to the host, and if the host profile requires user input, such as an IP address, upon startup, the host is put in maintenance mode.
- A first boot using VMware Auto Deploy requires that you set up your environment and add rules to the rule set. Reprovisioning requires less preparation. Several types of reprovisioning operations are available.
    - Simple reboot.
    - Reboot of hosts for which the user answered questions during the boot operation.
    - Reprovision with a different image profile.
    - Reprovision with a different host profile.

- If a host required user input during a previous boot, the answers are saved in an answer file and work even if you change the image. You can change the host profile to require new information. If you do, the host boots into maintenance mode. You can reapply the host profile and provide answers when prompted.
- One answer file per host is available.
- In an environment where no state is stored on the host, you can create a reference host. You configure the reference host with the logging, coredump, and other settings that you want, save the host profile, and write a rule that applies the host profile to other hosts as needed. You can configure the storage, networking, and security settings on the reference host and set up services such as syslog and NTP.
- A core dump is the state of working memory in the event of host failure. By default, a core dump is saved to the local disk. You can use ESXi Dump Collector to keep core dumps on a network server for use during debugging. ESXi Dump Collector is especially useful for Auto Deploy, but is supported for any ESXi host. ESXi Dump Collector supports other customization, including sending core dumps to the local disk.
- You can use Auto Deploy with the vCenter Server Appliance in different ways.
    - Use the vCenter Server system on the appliance in conjunction with the Auto Deploy server on the appliance.
    - Use the vCenter Server system on the appliance in conjunction with an Auto Deploy server that you install separately on a Windows system.
    - Use the Auto Deploy server on the appliance in conjunction with a vCenter Server system that you install on a different vCenter Server appliance.
    - Use the Auto Deploy server on the appliance in conjunction with a vCenter Server system that you install separately on a Windows system.
- You can register only one Auto Deploy instance with a vCenter Server system, and only one vCenter Server system with an Auto Deploy server.
- Host profiles allow you to prespecify information, for example, the storage setup or Syslog setup in a reference host and apply the host profile to a set of target hosts that share the same settings. You can also use host profiles to specify that certain settings are host-dependent. If you do so, the host comes up in maintenance mode when you provision it with Auto Deploy. Apply the host profile or update the answer file to be prompted for input. The system stores your input and uses it the next time the host boots.
- The answer file is not stored in a location or format that administrators can access. Use the Host Profiles UI in the vSphere Client to manage answer files.
- For hosts that will be part of a vSphere HA cluster, include the vmware-fdm VIB in the image profile
- Using DHCP reservations is highly recommended for address allocation. Fixed IP addresses are supported by the answer file mechanism, but providing input for each host is cumbersome and not recommended.
- Using Auto Deploy in environments that do not use VLANs is highly recommended. Do not use VLAN tagged networks at the boot NIC.
- Simultaneously booting large numbers of hosts places a significant load on the Auto Deploy server. Because Auto Deploy is a web server at its core, you can use existing web server scaling technologies to help distribute the load. After a massive power outage, VMware recommends that you bring up the hosts on a per-cluster basis. If you bring up multiple clusters simultaneously, the Auto Deploy server might experience CPU bottlenecks.
- Hosts provisioned with Auto Deploy do not have a local disk to store core dumps on. Install ESXi Dump Collector and set up your first host so all core dumps are directed to ESXi Dump Collector, and apply the host profile from that host to all other hosts.
- Auto Deploy transfers data over SSL to prevent casual interference and snooping. However, the authenticity of the client or of the Auto Deploy server is not checked during a PXE boot. The administrator (root) password and user passwords that are included with host profile and answer files are MD5 encrypted. Any other passwords associated with profiles are in the clear. If you set Active Directory by using host profiles, the passwords are not protected.
- You can use software depots, image profiles, and software packages (VIBs) to specify the software you want to use during installation or upgrade of an ESXi host.
- You use Image Builder for managing the software to deploy to your ESXi hosts in several different scenarios.
    - Create image profiles for use by Auto Deploy.
    - Add custom third-party drivers to existing image profile and export to ISO.
    - Perform upgrades.
    - Create custom images with reduced footprint.
- Understanding how depots, profiles, and VIBs are structured and where you can use them is a prerequisite for in-memory installation of a custom ESXi ISO, for provisioning ESXi hosts using VMware Auto Deploy, and for some custom upgrade operations.
    - *VIB* — A VIB is an ESXi software package. VMware and its partners package solutions, drivers, CIM providers, and applications that extend the ESXi platform as VIBs. VIBs can be used to create and customize ISO images or installed asynchronously onto ESXi hosts.
    - *Image Profile* — An image profile defines an ESXi image and consists of VIBs (software packages). An image profile always includes a base VIB, and might include additional VIBs. You examine and define an image profile using the Image Builder PowerCLI.
    - *Software Depot* — A logical grouping of VIBs, bulletins, and image profiles. The software depot can be online or offline. An online depot is accessible with an HTTP URL. An offline depots is accessible as a ZIP file. Public depots are made available by VMware and by VMware partners. Companies with large VMware installations might create internal depots to provision ESXi hosts using VMware Auto Deploy, or to export an ISO for ESXi installation.
- Image Builder PowerCLI cmdlets allow you to manage image profiles and VIBs.
- Image profiles define the set of VIBs that an ESXi installation or update process uses. Image profiles apply to hosts provisioned with Auto Deploy and other ESXi 5.0 hosts. You define and manipulate image profiles using the Image Builder PowerCLI.

- You can create a custom image profile from scratch or clone an existing profile and add or remove VIBs. A profile must meet the following requirements to be valid.
  - Each image profile must have a unique name and vendor combination.
  - Each image profile has an acceptance level. When you add a VIB to an image profile using an Image Builder PowerCLI cmdlet, Image Builder checks that the VIB matches the acceptance level defined for the profile.
  - You cannot remove VIBs that are required by other VIBs.
  - You cannot include two versions of the same VIB in an image profile. When you add a new version of a VIB, the new version replaces the existing version of the VIB.
- When you make a change to an image profile, Image Builder checks that the change does not invalidate the profile.
  - Dependency Check
  - Acceptance Level Validation
  - Acceptance Level Validation
- An image profile and its VIBs must meet several criteria to be valid.
  - Image profiles must contain at least one base VIB and one bootable kernel module.
  - If any VIB in the image profile depends on another VIB, that other VIB must also be included in the image profile.
  - VIBs must not conflict with each other.
  - No VIB can replace another VIB or make another VIB obsolete. Two VIBs with the same name but two different versions cannot coexist.
  - VIBs must follow file path and resource usage rules. VMware Certified and VMware Accepted VIBs always follow those rules.
  - No acceptance level validation issues exist.
- Acceptance Levels:
  - VMwareCertified
  - VMwareAccepted
  - PartnerSupported
  - CommunitySupported
- Image Builder consists of the Image Builder server and the Image Builder PowerShell cmdlets. The Image Builder server starts when your run the first Image Builder cmdlet.
- Cloning a published profile is the easiest way to create a custom image profile. Cloning a profile is especially useful if you use hosts from different vendors and want to use the same basic profile, but want to add vendorspecific VIBs.
- You can export an image profile to an ISO image or an offline ZIP file. Use the ISO image as an ESXi installer. You can use the ZIP file, which contains metadata and the VIBs specified in the image profile, for remediation.
- Image Builder PowerCLI cmdlets allow you to examine depots and VIBs.
- You can use Image Builder cmdlets to check which depots are available, add a depot, check which channels are in the depot, display image profile information, and create a new image profile by cloning one of the available image profiles.
- Published profiles are usually read only and cannot be modified. Even if a published profile is not read only, cloning instead of modifying the profile is best practice because modifying the original profile erases the original copy. You cannot revert to the original, unmodified profile except by reconnecting to a depot.
- In most situations, you create an image profile by cloning an existing profile. Some partners might need to create an image profile from scratch. Pay careful attention to dependencies and acceptance levels if you create an image profile from scratch.
- The system expects that the acceptance level of the VIBs you add to the base image is at least as high as the level of the base image. Pass in the `-AcceptanceLevel` parameter to change the acceptance level of the image profile if you have to add a VIB with a lower acceptance level.
- When you turn on the ESXi host for the first time or after resetting the configuration defaults, the host enters an autoconfiguration phase. This phase configures system network and storage devices with default settings.
- By default, Dynamic Host Configuration Protocol (DHCP) configures IP, and all visible blank internal disks are formatted with the virtual machine file system (VMFS) so that virtual machines can be stored on the disks.(*Forbes*: It says only first disk earlier)
- Use the direct console interface for initial ESXi configuration and troubleshooting.
- DCUI — Change the user interface to high-contrast mode `F4`.
- To manage your ESXi host remotely from a serial console, you can redirect the direct console to a serial port.
- You can redirect the direct console in the following ways:
  - Edit the boot option line manually by pressing Shift+O in the host boot loader.
  - Use the vSphere Client
  - Use the vSphere Client and make the setting persist for hosts that are deployed by Auto Deploy
- Use the direct console to do simple network connectivity tests. It performs the following tests:
  - Pings the default gateway
  - Pings the primary DNS name server
  - Pings the secondary DNS name server
  - Resolves the configured host name

- Restarting the management agents restarts all management agents and services that are installed and running in /etc/init.d on the ESXi host. Typically, these agents include hostd, ntpd, sfcbd, slpd, wsman, and vobd.
- The software also restarts Fault Domain Manager (FDM) if it is installed.
- One reason to disable the management network is to isolate an ESXi host from an HA and DRS cluster, without losing your static IP and DNS configurations or rebooting the host.
- When you restore the standard switch, a new virtual adapter is created and the management network uplink that is currently connected to Distributed Switch is migrated to the new virtual switch.
- Partitioning for hosts that are upgraded to ESXi 5.0 differs significantly from partitioning for new installations of ESXi 5.0.
- ESXi overwrites any disks that appear to be blank. Disks are considered to be blank if they do not have a valid partition table or partitions. If you are using software that uses such disks, in particular if you are using logical volume manager (LVM) instead of, or in addition to, conventional partitioning schemes, ESXi might cause local LVM to be reformatted. Back up your system data before you power on ESXi for the first time.
- On the hard drive or USB device from which the ESXi host is booting, the disk-formatting software retains existing diagnostic partitions that the hardware vendor creates. In the remaining space, the software creates the partitions as follows:
- Partitions Created by ESXi on the Host Drive:
    - ESXi Installable — For fresh installations, several new partitions are created for the boot banks, the scratch partition, and the locker. Fresh ESXi installations use GUID Partition Tables (GPT) instead of MSDOS-based partitioning. The partition table itself is fixed as part of the binary image, and is written to the disk at the time the system is installed. The ESXi installer leaves the scratch and VMFS partitions blank and ESXi creates them when the host is rebooted for the first time after installation or upgrade. One 4GB VFAT scratch partition is created for system swap. The VFAT scratch partition is created only on the disk from which the ESXi host is booting. On the other disks, the software creates a VMFS5 partition on each disk, using the whole disk. When you install on a disk, the installer overwrites the entire disk. When the installer autoconfigures storage, the installer does not overwrite hardware vendor partitions. During ESXi installation, the installer creates a 110MB diagnostic partition for core dumps.
    - ESXi Embedded — One 110MB diagnostic partition for core dumps, if this partition is not present on another disk. The VFAT scratch and diagnostic partitions are created only on the disk from which the ESXi host is booting. On other disks, the software creates one VMFS5 partition per blank disk, using the whole disk. Only blank disks are formatted.
    - ESXi Installable and ESXi Embedded — One VMFS5 partition on the remaining free space.
- The scratch partition is not required. It is used to store vm-support output, which you need when you create a support bundle. If the scratch partition is not present, vm-support output is stored in a ramdisk. In low memory situations, you might want to create a scratch partition if one is not present.
- For ESXi Embedded, if a partition is not found, but an empty local disk exists, the system formats it and creates a scratch partition. If no scratch partition is created, you can configure one, but a scratch partition is not required. You can also override the default configuration. You might want to create the scratch partition on a remote NFS mounted directory.
- When you enable lockdown mode, no users other than vpxuser have authentication permissions, nor can they perform operations against the host directly. Lockdown mode forces all operations to be performed through vCenter Server.
- When a host is in lockdown mode, you cannot run vSphere CLI commands from an administration server, from a script, or from vMA against the host. External software or management tools might not be able to retrieve or modify information from the ESXi host.
- The root user is still authorized to log in to the direct console user interface when lockdown mode is enabled.
- If you enable or disable lockdown mode using the Direct Console User Interface (DCUI), permissions for users and groups on the host are discarded. To preserve these permissions, you must enable and disable lockdown mode using the vSphere Client connected to vCenter Server.
- When you reset the configuration, the software overrides all your configuration changes, deletes the password for the administrator account (root), and reboots the host.
- After you reset the configuration defaults, the virtual machines are not visible, but you can retrieve them by reconfiguring storage and reregistering the virtual machines.
- VMware recommends that you write down the license key and tape it to the server, or put the license key in a secure, easily accessible location.
- When you perform a configuration backup, the serial number is backed up with the configuration and is restored when you restore the configuration. The serial number is not preserved when you perform the repair operation. The recommended procedure is to first back up the configuration, run the repair operation, and then restore the configuration.
- Perform the backup by running the `vicfg-cfgbackup` command from the vSphere CLI.
- When you restore the configuration, the target host must be in maintenance mode, which means all virtual machines (including the vMA) must be powered off.
- Each vCenter Server instance must have its own database. vCenter Server instances cannot share the same database schema.
- vCenter Server supports IBM DB2, Oracle, and Microsoft SQL Server databases. Update Manager supports Oracle and Microsoft SQL Server databases.
- VMware recommends using separate databases for vCenter Server and Update Manager. For a small deployments, a separate database for Update Manager might not be necessary.
- vCenter Server databases require a UTF code set.

- Configuration and Patch Requirements for Databases Supported with vCenter Server:

Database Type	Patch Configuration Requirements
IBM DB2 9.5	Fix pack 5 is required. Fix pack 7 is recommended
IBM DB2 9.7	Fix pack 2 is required. Fix pack 3a is recommended.
Microsoft SQL Server 2008 R2 Express	Up to 5 hosts and 50 virtual machines.
Microsoft SQL Server 2005	32-bit and 64-bit versions. SP3 is required. SP4 is recommended.
Microsoft SQL Server 2008 & 2008 R2	32-bit and 64-bit versions. SP1 is required. SP2 is recommended.
Oracle 10g R2	Apply patch 10.2.0.4 (or later) to the client & server. Then apply patch 5699495 to the client.
Oracle 11g	Oracle 11g R1 11.1.0.7 Oracle 11g R2 11.2.0.1 with patch 5 (9966926)

- The vCenter Server system must have a 64-bit DSN. This requirement applies to all supported databases.
- vCenter Server must have a computer name that is 15 characters or fewer.
- The name-length limitation applies to the vCenter Server system. The data source name (DSN) and remote database systems can have names with more than 15 characters. The name change has no effect on communication with remote databases.
- Changing the vCenter Server computer name impacts database communication if the database server is on the same computer with vCenter Server.
- Set Database Permissions By Manually Creating Database Roles and the VMW Schema**.** By using this method, available with vCenter Server 5.0, the vCenter Server database administrator can set permissions for vCenter Server users and administrators to be granted through Microsoft SQL Server database roles.
- VMware recommends this method because it removes the requirement to set up the database role *dbo* and *db_owner* schema for vCenter Server users who install and upgrade vCenter Server.
- Alternatively, you can assign vCenter Server database permissions by creating and assigning the *dbo* role and letting the vCenter Server installer create the default schema that assigns database user permissions to that role.
- If the system that you use for your vCenter Server installation belongs to a workgroup rather than a domain, not all functionality is available to vCenter Server.
- Make sure the system on which you are installing vCenter Server is not an Active Directory domain controller.
- You can use the Microsoft Windows built-in system account or a user account to run vCenter Server. With a user account, you can enable Windows authentication for SQL Server, and it provides more security.
- The user account must be an administrator on the local machine.
- If you install vCenter Server on a system that is configured to use IPv6, vCenter Server uses IPv6
- You must enclose the IPv6 address in square brackets: [*IPv6-address*].
- The vCenter installer configures the default URLs entries as:
  - For the VirtualCenter.VimApiUrl key, the default value is http(s)://*FQDN of VC machine*/sdk.
  - For the Virtualcenter.VimWebServicesUrl key, the default value is https://*FQDN of VC machine:installed-webservices-port*/vws.
- The following components are installed when you install vCenter Server:
  - VMware vCenter Server
  - Microsoft .NET 3.5 SP1 Framework
  - Microsoft Windows Installer version 4.5
  - VMware vCenter Orchestrator (not supported on IPv6-only operating systems)
  - Microsoft SQL Server 2008 R2 Express (optional)
- Install application includes links to install the following optional components:
  - vSphere Client
  - vSphere Web Client
  - vSphere Update Manager
  - vSphere ESXi Dump Collector
  - vSphere Syslog Collector
  - vSphere Auto Deploy
  - vSphere Authentication Proxy vCenter Server — Support tool that enables ESXi hosts to join a domain without using Active Directory credentials. This tool enhances security for PXE-booted hosts and hosts that are provisioned using Auto Deploy, by removing the need to store Active Directory credentials in the host configuration.
- The vCenter Server Appliance has the default user name `root` and password `vmware`. Prerequisites:
  - Verify that vSphere Client is installed.
  - You can deploy the vCenter Server Appliance only on hosts that are running ESX version 4.x or ESXi version 4.x or later.
  - The vCenter Server Appliance requires at least 7GB of disk space, and is limited to a maximum size of 80GB.
- Before installation, create and set up an Update Manager database and 32-bit DSN, unless you are using the bundled SQL Server 2008 R2 Express.
- ESXi Dump Collector does not support vSphere distributed switches in ESXi 5.0.

- When you install the vSphere Authentication Proxy service, the installer creates a domain account with appropriate privileges to run the authentication proxy service. The account name begins with the prefix CAMand has a 32-character, randomly generated password associated with it. The password is set to never expire. Do not change the account settings.
- vCenter Server instances in a group replicate shared global data to the LDAP directory. The global data includes the following information for each vCenter Server instance:
  - Connection information (IP and ports)
  - Certificates
  - Licensing information
  - User roles
- Linked Mode groups that contain both vCenter Server 5.0 and earlier versions of vCenter Server are not supported (*Forbes*: This is correct. The current vSphere 5.0 Release Notes incorrectly state it is supported).
- The vCenter Server instances in a Linked Mode group can be in different domains if the domains have a two-way trust relationship. Each domain must trust the other domains on which vCenter Server instances are installed.
- When adding a vCenter Server instance to a Linked Mode group, the installer must be run by a domain user who is an administrator on both the machine where vCenter Server is installed and the target machine of the Linked Mode group.
- All vCenter Server instances must have network time synchronization. The vCenter Server installer validates that the machine clocks are not more than five minutes apart.
- If you are upgrading a vCenter Server that is part of a Linked Mode group, it will be removed from the group.
- You cannot join a Linked Mode group during the upgrade procedure when you are upgrading from VirtualCenter 25 to vCenter Server 5.0. You can join after the upgrade to vCenter Server is complete.
- It might take several seconds for the global data (such as user roles) that are changed on one machine to be visible on the other machines. The delay is usually 15 seconds or less. It might take a few minutes for a new vCenter Server instance to be recognized and published by the existing instances, because group members do not read the global data very often.
- Reboot the host and press “shift+R” at the beginning of the boot process to instruct the boot loader to boot off the alternate boot bank.

vSphere Upgrade (vSphere 5.0, ESXi 5.0, vCenter 5.0 & vSphere Client 5.0)

- VMware supports in-place upgrades on 64-bit systems from vCenter Server 4.x to vCenter Server 5.0.
- You can upgrade VirtualCenter 2.5 Update 6 or later and vCenter Server 4.0.x to vCenter Server 5.0 by installing vCenter Server 5.0 on a new machine and migrating the existing database. This upgrade method makes it possible to upgrade from a 32-bit system to a 64-bit system. Alternatively, if the VirtualCenter database is on a remote machine, you can simply upgrade the database.
- vCenter Server 5.0 can manage ESX 3.5.x/ESXi 3.5.x hosts in the same cluster with ESX 4.x/ESXi 4.x hosts.
- You cannot upgrade vCenter Server 4.x that is running on Windows XP Professional x64 Edition to vCenter Server 5.0, because vCenter Server 5.0 does not support Windows XP Professional x64.
- Run the vCenter Host Agent Pre-Upgrade Checker to produce a report showing known issues that might prevent a successful upgrade of the vCenter Host Agent software. Found here: `\vpx\agentupgradecheck\AgentUpgradeChecker.exe`
- Each time you run the tool, the system queries VMware.com and downloads any new updates for the tool. This action ensures that as new upgrade issues are discovered, the tool remains as useful as possible.
- Before upgrading vCenter Server to version 5.0, upgrade any ESX 3.5 hosts that are in a vSphere HA cluster to patch level 24.
- To view the database upgrade log, open %TEMP%\VCDatabaseUpgrade.log.
- In previous releases of vCenter Server, datastores and networks inherited access permissions from the datacenter. In vCenter Server 5.0, they have their own set of privileges that control access to them. This might require you to manually assign privileges, depending on the access level you require.
- In vCenter Server 5.0, users are initially granted the No Access role on all new managed objects, including datastores and networks. This means, by default, users cannot view or perform operations on them. All existing objects in vCenter Server maintain their permissions after the upgrade. To determine whether to assign permissions to existing datastores and networks, the upgrade process uses the datacenter's *Read-only* privilege.
  - If the *Read-only* privilege is nonpropagating (not inherited by child objects), VMware assumes acces privileges should not be assigned to datastores and networks. In such cases, you must update your roles to include the new datastore and network privileges desired. This is required for users to view and perform operations on these objects.
  - 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. In such cases, the default minimum privileges are automatically assigned during the upgrade process.
- You can upgrade Update Manager 1.0 Update 6 and Update Manager 4.x to Update Manager 5.0. You cannot change the installation path and patch download location (*Forbes*: without a fresh installation).
- You can install Update Manager 5.0 only on a 64-bit operating system. If you are running an earlier version of Update Manager on a 32-bit platform, you must either back up and restore your database manually, or use the data migration tool to back up the existing data on the 32-bit machine.

- Make sure that sufficient disk space is available on the host for the upgrade or migration. Migrating from ESX 4.x to ESXi 5 requires 50MB of free space on your VMFS datastore.
- If a SAN is connected to the host, detach the fibre before continuing with the upgrade or migration. Do not disable HBA cards in the BIOS. This does not apply to ESX hosts that boot from the SAN, and have the Service Console on the on the SAN LUNs. You can disconnect LUNs that contain the VMFS datastore and do not contain the Service Console.
- The migration or upgrade from ESX 4.x or ESXi 4.x to ESXi 5.0 does not migrate all host configuration files and settings.
- The upgrade process preserves as much of the ESX host configuration as possible. However, because of differences between ESX 4.x and ESXi 5.0 architecture, many configuration files cannot be migrated (*Forbes*: there are many more listed in the documentation, but I thought these were the most noteable examples):
  - /etc/logrotate.conf Not migrated. ESXi Logrotation is incompatible with prior versions.
  - /etc/localtime Not migrated. Timezones are not supported in ESXi.
  - /etc/syslog.conf Migrated for ESXi, not migrated for ESX.
  - /etc/sysconfig/network Migrated. Service Console virtual NICs (vswif) will be converted to ESXi virtual NICs (vmk).
  - /etc/sysconfig/i18n Not migrated. i18n is not supported in ESXi.
  - /etc/ssh Migrated. OpenSSH is now included in ESXi.
  - /etc/hosts.allow & hosts.deny Not migrated.
  - /etc/sudoers Not migrated. `sudo` is not supported in ESXi.
  - /etc/vmware/vmware.lic Not migrated. ESXi 5.0 upgrades are reset to evaluation mode.
  - /etc/passwd & /etc/shadow Partially migrated. Only the root user password will be saved, if possible.
- Custom ports that were opened by using the ESX/ESXi 4.1 esxcfg-firewall command do not remain open after the upgrade to ESXi 5.0. The configuration entries are ported to the esx.conf file by the upgrade, but the corresponding ports are not opened.
- The upgrade to ESXi 5.0 affects the amount of memory available to the host system. As a result, in resource pools that are set to use nearly all of the resources available, some virtual machines may not have enough resources to start after the upgrade.
- When you upgrade from ESXi 4.x to ESXi 5.x, the default maximum number of ports for a virtual switch changes from 64 to 128.
- All vswif interfaces are migrated to vmk interfaces. If a conflict is detected between two interfaces, one is left in disabled state. The upgrade disables any conflicting kernel IP addressing in favor of the management interface.
- Fresh ESXi installations use GUID Partition Tables (GPT) instead of MSDOS-based partitioning.
- The partition table itself is fixed as part of the binary image, and is written to the disk at the time the system is installed. The ESXi installer leaves the scratch and VMFS partitions blank and ESXi creates them when the host is rebooted for the first time after installation or upgrade. The scratch partition is 4GB. The rest of the disk is formatted as a VMFS5 partition.
- Upgraded systems do not use GUID Partition Tables (GPT), but retain the older MSDOS-based partition label.
- For most ESXi 4.x hosts, the partition table is not rewritten in the upgrade to ESXi 5.0. The partition table is rewritten for systems that have lopsided bootbanks. Lopsided boot banks can occur in systems that are upgraded from ESXi 3.5 to ESXi 4.x, and then directly to ESXi 5.0.
- For ESX hosts, the partitioning structure is changed to resemble that of an ESXi 4.x host. The VMFS3 partition is retained and a new MSDOS-based partition table overwrites the existing partition table.
- Upgraded hosts do not have a scratch partition. Instead, the scratch directory is created and accessed off of the VMFS volume. Each of the other partitions, such as the bootbanks, locker and vmkcore will be identical to that of any other system.
- In upgraded hosts, the VMFS partition is not upgraded from VMFS3 to VMFS5. ESXi 5.0 is compatible with VMFS3 partitions. You can upgrade the partition to VMFS5 after the host is upgraded to ESXi 5.0.
- Upgraded hosts, which retain the older MSDOS-based partitioning, do not support installing ESXi on a single physical disk or LUN larger than 2TB. To install ESXi on a disk or LUN larger than 2TB, you must do a fresh installation.
- For the VMFS partition on the disk to be preserved during an upgrade to ESXi 5.0, the partition must be physically located after the boot partition, which is partition 4, and the extended partition on the disk (8192 + 1835008 sectors). Any system that has a VMFS partition after the 1843200 sector mark can retain that VMFS partition, regardless of whether it was initially installed with ESX 3.5 or 4.x.
- For systems in which the VMFS partition is placed on a different drive from the boot drive, the entire contents of the boot drive is overwritten during the upgrade. Any extra data on the disk is erased.
- There are several ways to upgrade ESX/ESXi hosts:
  - vSphere Update Manager
  - Upgrade or migrate interactively using an ESXi installer ISO image on CD/DVD or USB flash drive
  - Perform a scripted upgrade
  - vSphere Auto Deploy — You cannot use Auto Deploy to upgrade or migrate version 4.x ESX and ESXi hosts to ESXi 5.0, because version 4.x ESX and ESXi hosts are deployed by the traditional method of installing the software on the host hard disk.
  - `esxcli` — You can upgrade and apply patches to ESXi 5.0 hosts using the `esxcli` command-line utility for ESXi. You cannot use `esxcli` to upgrade ESX/ESXi 4.x hosts to ESXi 5.
- The `esxupdate` and `vihostupdate` utilities are not supported for ESXi 5.0 upgrades.
- Hosts must have more than 350MB of free space in the /boot partition to support the Update Manager upgrade process. If the host that you are upgrading does not have more than 350MB of free space in the /boot partition, use a scripted or interactive upgrade instead.
- Once you have upgraded or migrated your host to ESXi 5.0, you cannot roll back to your version 4.x ESX or ESXi software.
- If the installer finds an existing ESX or ESXi installation and VMFS datastore you can choose from the following options:

    - Upgrade ESXi, preserve VMFS datastore
    - Install ESXi, preserve VMFS datastore
    - Install ESXi, overwrite VMFS datastore
- If you press Ctrl+C while an `esxcli` command is running, the command-line interface exits to a new prompt without displaying a message. However, the command continues to run to completion.
- You can update a host with VIBs stored in a software depot that is accessible through a URL or in an offline ZIP depot.
- Determine which VIBs are installed on the host: `esxcli --server=server_name software vib list`
- Update VIBS: `esxcli --server=server_name software vib update --depot=http://web_server/depot_name`
- You can update a host with image profiles stored in a software depot that is accessible through a URL or in an offline ZIP depot: `esxcli --server=server_name software profile update --depot=http://webserver/depot_name --profile=profile_name`
- You can update hosts with third-party VIBs or image profiles only by downloading a ZIP file of a depot prepared by the VMware partner, and downloading that ZIP file directly to the ESXi host: `esxcli --server=server_name software vib update --depot=/path_topartner_vib_ZIP/partner_ZIP_file_name.zip`
- If a third-party extension is released as a VIB package, and use the `esxcli software vib` command to add the VIB package to your system, the VIB system updates the firewall ruleset and refreshes the host daemon after you reboot your system.
- You can use the vSphere Client to export the upgrade log files (*Forbes*: the RC document said you could also find the upgrade log on the host at: /locker/db/esxupdate.log).
- After you upgrade to ESXi 5.0, reapply your host license.
- Do not use `vmware-vmupgrade.exe` to upgrade virtual machines.
- The version of VMware Tools included in vSphere 5.0 is supported on vSphere 4.x and 5.0 virtual machines. That is, you can also use this new version of VMware Tools in virtual machines on ESX/ESXi 4.x hosts.
- Virtual machines in a vSphere 5.0 environment support the versions of VMware Tools included in vSphere 4.0-5.0. That is, you are not strictly required to upgrade VMware Tools if VMware Tools was installed from an ESX/ESXi 4.x host.
- Virtual machines with virtual hardware version 8 are not supported on hosts running earlier versions of ESX/ESXi software.
- Paravirtualization (VMI) is not supported on ESXi 5.0.
- For Windows 2000 and later, VMware Tools installs a virtual machine upgrade helper tool. This tool restores the network configuration if you upgrade from virtual hardware version 4 to version 7 or higher.
- For Linux, VMware provides operating system specific packages (OSPs) as a packaging and distribution mechanism for VMware Tools. These VMware Tools OSPs are packaged using native package formats and standards such as *rpm* and *deb*.
- Use OSPs if you want to use native update mechanisms, rather than vCenter Server, to manage updates for VMware Tools. If you use an OSP, the VMware Tools status is *unmanaged* on the virtual machine *Summary* tab. The status *unmanaged* means that you cannot use vCenter Server to manage VMware Tools and you cannot use vSphere Update Manager to upgrade VMware Tools.
- When you upgrade VMware Tools on Linux guest operating systems, new network modules are available but are not used until you either reboot the guest operating system or stop networking, unload and re-load the VMware networking kernel modules, and then restart networking. This behavior means that even if VMware Tools is set to automatically upgrade, you must reboot or re-load network modules to make new features available.
- When you upgrade from virtual hardware version 4 to version 8 the upgrade is reversible if you take a virtual machine backup or snapshot before performing the upgrade.

vCenter Server and Host Management (ESXi 5.0 & vCenter 5.0)

- *Tomcat Web server* — Many vCenter Server functions are implemented as Web services that require the Tomcat Web server. The Tomcat Web server is installed on the vCenter Server machine as part of the vCenter Server installation. Features that require the Tomcat Web server to be running include: Linked Mode, CIM/Hardware Status tab, Performance charts, WebAccess, vCenter Storage Monitoring/Storage Views tab, and vCenter Service status.
- *vCenter Server agent* — On each managed host, the software that collects, communicates, and executes the actions received from vCenter Server. The vCenter Server agent is installed the first time any host is added to the vCenter Server inventory.
- *Host agent* — On each managed host, the software that collects, communicates, and executes the actions received through the vSphere Client. It is installed as part of the ESXi installation.
- *Storage DRS* — A feature that enables you to manage multiple datastores as a single compute resource, called a datastore cluster. A datastore cluster is an aggregation of multiple datastores into a single logical, load-balanced pool. You can treat the datastore cluster as a single flexible storage resource for resource management purposes. You can assign a virtual disk to a datastore cluster, & Storage DRS finds an appropriate datastore for it. The load balancer takes care of initial placement and future migrations based on workload measurements. Storage space balancing & I/O balancing minimize the risk of running out of space & the risk of I/O bottlenecks slowing the performance of virtual machines.
- *vSphere Web Client* — A Web application installed on a machine with network access to your vCenter Server installation. In this release, the vSphere Web Client includes a subset of the functionality included in the Windows-based vSphere Client, primarily related to inventory display, virtual machine deployment & configuration. Before you can connect to a vCenter Server system, you must register it with the vSphere Web Client.

- VMware recommends that you register a given vCenter Server system with only one vSphere Web Client instance, rather than using multiple vSphere Web Client instances to manage that vCenter Server system.
- The administration tool: https://localhost:port/admin-app
- The URL for the vSphere Web Client: https://server_name:9443/vsphere-client
- To configure licenses:
  - Select *Administration* > *vCenter Server Settings* to display the vCenter Server Settings dialog box.
  - Select the type of license key to assign to this vCenter Server.
  - To enter ESX 4.0/ESXi 4.0 keys, select *Home* > *Administration* > *Licensing*.
- Statistic intervals determine the frequency at which statistic queries occur, the length of time statistical data is stored in the database, and the type of statistical data collected.
- To change a collection interval attribute, select its row in the Statistics Interval section:
  - *Keep Samples for* — select an archive length. This option is configurable only for the Day and Year intervals.
  - *Statistics Interval* — select an interval duration. This option is configurable only for the Day interval.
  - *Statistics Level* — select a new level interval level. Level 4 uses the highest number of statistics counters. Use it only for debugging purposes. The statistics level must be less than or equal to the statistics level set for the preceeding statistics interval. This is a vCenter Server dependency.
- The Database Size section, estimate the effect of the statistics settings on the database.
  - Enter the number of *Physical Hosts.*
  - Enter the number of *Virtual Machines.* The estimated space required and number of database rows required are calculated and displayed.
- *vCenter Server Unique ID* — You can change this value to a number from 0 through 63 to uniquely identify each vCenter Server system running in a common environment.
- You can configure timeout intervals for vCenter Server operations for *Normal Operations* and *Long Operations.* Don't set these values to zero.
- Logging options:
  - None (Disable logging)
  - Error (Errors only)
  - Warning (Errors and warnings)
  - Info (Normal logging)
  - Verbose (Verbose)
  - Trivia (Extended verbose)
- You can configure the maximum number of database connections that can occur simultaneously. Generally, you do not need to change this value. You might want to increase this number if your vCenter Server system frequently performs many operations and performance is critical. You might want to decrease this number, if the database is shared and connections to the database are costly. VMware recommends that you not change this value unless one of these issues pertains to your system.
- In order to limit the growth of the vCenter Server database and conserve storage space, you can configure the database to discard information about tasks or events after a specified period of time. Do not use these options if you want to retain a complete history of tasks and events for your vCenter Server.
- You can use the Advanced Settings page to modify the vCenter Server configuration file, vpxd.cfg. This page can be used to add entries to the vpxd.cfg file, but not to edit or delete them.
- The following features are not supported by the vCenter Server Virtual Appliance.
  - IPv6
  - Linked Mode
  - SQL Server as a supported database
  - DB2 as a supported database
- Database Type :
  - embedded (option is available only for a small inventory size, with fewer than 100 hosts and 1000 virtual machines)
  - Oracle
- VSVA memory sizing:
  - 4 GB or higher — Memory requirement for fewer than 10 hosts and 100 virtual machines in the VSVA inventory.
  - 8 GB or higher — between 10 and 100 hosts or between 100 and 1000 virtual machines
  - 13 GB or higher — between 100 and 400 hosts or between 1000 and 4000 virtual machines
  - 17 GB or higher — over 400 hosts or 4000 virtual machines
- ESXi Dump Collector is installed and enabled by default on the vCenter Server Virtual Appliance.
- ESXi Dump Collector does not support vSphere distributed switches in ESXi 5.0.
- Enable storing vCenter Server Virtual Appliance log and core files on NFS to store vCenter Server Virtual Appliance files on an NFS datastore.
- Disable root squashing on the NFS server.
- vCenter Server Foundation and vCenter Server Essentials editions do not support Linked Mode.
- When you join a vCenter Server system to a linked mode group, the roles defined on each vCenter Server system in the group are replicated to the other systems in the group.

- If the roles defined on each vCenter Server system are different, the roles lists of the systems are combined into a single common list. For example, if vCenter Server 1 has a role named Role A and vCenter Server 2 has a role named Role B, then both servers will have both Role A and Role B after they are joined in a linked mode group.
- If two vCenter Server systems have roles with the same name, the roles are combined into a single role if they contain the same privileges on each vCenter Server system. If two vCenter Server systems have roles with the same name that contain different privileges, this conflict must be resolved by renaming at least one of the roles. You can choose to resolve the conflicting roles either automatically or manually.
- If you choose to reconcile the roles automatically, the role on the joining system is renamed to *vcenter_namerole_name* where *vcenter_name* is the name of the vCenter Server system that is joining the Linked Mode group and *role_name* is the name of the original role.
- To change the domain of a vCenter Server system in a Linked Mode group, isolate the vCenter Server system from the Linked Mode group first.
- If you connect a vCenter Server system to a Linked Mode group and the vCenter Server system has a machine name that does not match the domain name, several connectivity problems arise. Correct this situation by changing the URLs.
    - For the Virtualcenter.VimApiUrl key, the default value is http(s)://Fully qualified domain name (FQDN) of vCenter Server machine/sdkvCenter Server.
    - For the Virtualcenter.VimWebServicesUrl key, the default value is https://FQDN of vCenter Server machine:installed-webservices-port/vwsvCenter Server.
- vRAM is a licensing-specific measure for the memory that is configured for powered-on virtual machines that run on ESXi 5.0 hosts. vSphere 5.0 license keys have per-processor capacity with pooled vRAM entitlements. When you assign a vSphere 5.0 license key of a certain edition to an ESXi 5.0 host, the key adds a certain amount of vRAM to a vRAM pool that is created for the corresponding license edition.
- The amount of vRAM that a vSphere 5.0 key adds to a vRAM pool is equal to the processor capacity of the license key multiplied by the vRAM entitlement for license edition.
- When you power on a virtual machine on an ESXi 5.0 host, the amount of memory that is configured to the virtual machine counts against the same amount of vRAM in the vRAM pool that is available for the ESXi 5.0 host where the machine runs.
- Certain products have licenses that implement hard enforcement of the license compliance. Hard-enforced licenses prevent operations that result in exceeding the capacity of the license key. If you try to perform an operation that results in exceeding the capacity of a hard-enforced license, the operation is unsuccessful and you receive an error message. For example, vCenter Server Essentials that is included in the Essentials Kits, implements a hard enforcement on the maximum amount of pooled vRAM that you can use for powered-on virtual machines. When the maximum amount of the pooled vRAM is reached, you cannot power on virtual machines.
- Products that have licenses implementing soft enforcement rely on vCenter Server alerts and license reporting to notify when the license use for a product exceeds the available license capacity. When the license use for a product that has soft-encored licenses exceeds the available license capacity for the product, an alarm triggers on vCenter Server or on all vCenter Server instances in a Linked Mode group.
- The policy of VMware is that license capacity must be purchased in advance of use. You should plan accordingly to avoid exceeding the license capacity for products with soft-enforced licenses.
- The consumed vRAM is equal to the total amount of memory that is provisioned on powered-on virtual machines. The amount of consumed vRAM is calculated for every vRAM pool that is available.
- You can extend a vRAM pool of a certain license edition in three ways.
    - Assign new license keys of the same edition to unlicensed ESXi 5.0 hosts.
    - Assign a license key of the same edition and a larger processor capacity to ESXi 5.0 hosts.
    - Upgrade the editions of the license keys that are assigned to ESXi 5.0 hosts.
- vCenter Server systems in Linked Mode share a common license inventory.
- If you have license keys for vSphere 5.0 or later, but you need to license VMware Infrastructure 3 or vSphere 4.0 assets, use the license portal to downgrade the license keys. When you downgrade the license keys, the vSphere 5.0 license keys remain valid. When you are ready to upgrade your assets, you can stop using the VMware Infrastructure or vSphere licenses and start using the vSphere 5.0 license keys by adding them to the vCenter Server license inventory and assigning them to your upgraded assets.
- If a product is marked with one asterisk (*) symbol, the license has a strong (*Forbes*: strong means hard?) limit.
- If a product is marked with two asterisks (**) symbol, the license usage for the product includes data from vCenter Server instances that do not participate in the current selection.
- The Licensing Accounting Module available on each vCenter Server system collects the license usage snapshots every hour.
- A temper-detection feature in vCenter Server protects the license usage information. If the licensing data in the vCenter Server database has been edited, you cannot export a license usage report.
- The license usage information that you view in the *Reporting* tab is aggregated data from the license usage snapshots available in the vCenter Server database over the time period you select.
- The exported report contains raw data on the license usage for a single product or a product group over the time period you select. The exported report includes all the snapshots collected over the selected time period. The exported report can contain host, vCenter Server, or solution asset IDs, but does not include user-defined names. The exported report is free of user and company-specific information.
- If a vCenter Server system manages an ESXi host, changes made to the host license through direct connection to the host do not persist, because the license key assigned through vCenter Server overwrites the changes.
- You can view system logs from the Home page of a vSphere Client.
- You can use the direct console interface to view the system logs on an ESXi host.

- On Windows systems, several log files are stored in the Local Settings directory.
- The ESXi system logs can be found in /var/run/log.
- Client Installation log: Temp directory on the vSphere Client machine.
- Client Service log \vpx directory in the Application Data directory on the vSphere Client machine.
- All ESXi hosts run a syslog service (vmsyslogd), which logs messages from the VMkernel and other system components to log files.
- You can use the vSphere Client or the `esxcli system syslog` vCLI command to configure the syslog service.
- You can collect and package all relevant ESXi system and configuration information, a well as ESXi log files, by running the following script on the ESXi Shell: `/usr/bin/vm-support`
- Do not schedule multiple tasks to be performed at the same time on the same object. The results are unpredictable.
- Canceling a task stops a running task from occurring, regardless of whether the task was a real-time task or a scheduled task. The operation cancels only the running task. If the task being canceled is a scheduled task, subsequent runs are not canceled.
- Tasks that aren't running can be cleared when they are in a queued or scheduled state. In such cases, because the cancel operation is not available, either remove the task or reschedule it to run at a different time. Removing a scheduled task requires that you recreate it to run it in the future, rescheduling does not.
- The user performing the task in the vSphere Client must have the correct permissions on the relevant objects. After a scheduled task is created, it will be performed even if the user no longer has permission to perform the task.
- Some errors that appear for tasks, events, and in error dialogs can be submitted to VMware for further investigation. In some cases, more information or links to knowledge base articles are provided.
- Disconnecting a managed host does not remove it from vCenter Server; it temporarily suspends all monitoring activities performed by vCenter Server.
- Storage vMotion requires that the 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. If you convert the mapping file, a new virtual disk is created and the contents of the mapped LUN are copied to this disk. For physical compatibility mode RDMs, you can migrate the mapping file only.
- EVC masks only those processor features that affect vMotion compatibility. Enabling EVC does not prevent a virtual machine from taking advantage of faster processor speeds, increased numbers of CPU cores, or hardware virtualization support that might be available on newer hosts.
- When a virtual machine is powered on, it determines the EVC mode of the cluster in which it is running. If the EVC mode of the cluster is subsequently raised, the virtual machine does not change its EVC mode until it is powered off and powered on again.
- You can run the `svmotion` command in either interactive or noninteractive mode. In interactive mode using the `--interactive` switch), you are prompted for all the information necessary to complete the storage migration.
- vCenter places limits on the number of simultaneous VM migrations and provisioning operations. Each operation, such as a migration with vMotion or cloning a virtual machine, is assigned a resource cost. Each type of resource, such as host, datastore, or network, has a maximum cost that it can support at any one time. Any new migration or provisioning operation that would cause a resource to exceed its maximum cost does not proceed immediately, but is queued until other operations complete and release resources. Each of the network, datastore, and host limits must be satisfied in order for the operation to proceed.
- Network limits apply to migrations with vMotion only.
- Maximum Cost of vMotion:
  - 1GigE **4**
  - 10GigE **8**
  - All migrations with vMotion have a network resource cost of 1.
- Datastore limits apply to migrations with vMotion and with Storage vMotion. 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.
  - vMotion/Storage vMotion **128**
- Datastore Resource Costs for vMotion and Storage vMotion
  - Storage vMotion **16**
- Host limits apply to migrations with vMotion, Storage vMotion, and other provisioning operations such as cloning, deployment, and cold migration.
  - vMotion **8**
  - Storage vMotion **2**
  - Other provisioning operations **8**
- Host Resource Costs for vMotion, Storage vMotion, and Provisioning Operations
  - vMotion **1**
  - Storage vMotion **4**
  - Other provisioning operations **1**
- Resource Map Icons show a host that is compatible for vMotion migration in a green circle. The color of the circle varies in intensity based on the load of the current host. Heavily used hosts are pale; low-load hosts are saturated green.

What's New in VMware vSphere 5.0 Platform Whitepaper

- Up to 32 virtual CPUs (vCPUs) and up to 1TB of RAM
- vSphere 5.0 provides new forms of SSD handling and optimization. The VMkernel automatically recognizes and tags SSD devices that are local to an ESXi host or are on the network. In addition, the VMkernel scheduler is modified to allow ESXi swap to extend to local or network SSD devices, which enables memory over commitment and minimizes performance impact.
- vSphere 5.0 marks the beginning of efforts by VMware to standardize on a single CLI for both local and remote administration, as well as to help reduce the overall number of CLI tools. The `esxcli` command is available on each VMware ESXi host via the VMware ESXi shell. It is also available as part of the optional vCLI package that can be installed on any supported Windows or Linux server, or through the vSphere Management Assistant (vMA).
- In vSphere 5.0, the new `esxcli` command replaces the deprecated `esxcfg-*` style commands. However, it doesn't yet provide a comprehensive set of command-line capabilities. The `esxcli` command will continue to be enhanced in future releases and will eventually replace the non-`esxcli` commands. Until that time, you will continue to augment the `esxcli` command with the `vicfg-*` commands and other familiar CLI tools such as `vmware-cmd` and `vmkfstools` to troubleshoot and administer your VMware ESXi hosts. Of course, you can also continue to use the vSphere PowerCLI.
- In addition to the new `esxcli` command, a new `localcli` command has been added in vSphere 5.0. The `localcli` command is largely equivalent to the `esxcli` command, with the notable exception that it bypasses the local `hostd` process on the server. The `localcli` command is intended for situations where the VMware ESXi host's `hostd` daemon becomes unresponsive. It is recommended that you do not use the `localcli` command outside of the direction of VMware global services because it can result in host instability.
- In VMware ESXi, the access control capability is provided through a vmknic (VMkernel network adaptor)-level firewall module. This module sits between a vmknic and a virtual switch. It inspects packets against firewall rules. Based on the results, it determines whether to drop or pass packets.
- Host Profiles have been extended to include additional configuration settings not in earlier versions, such as support for iSCSI, FCoE, storage multipathing, individual device settings and kernel module settings. In addition, for host-specific configuration attributes Host Profiles now enables creating a per-host answer file.
- Update Manager now monitors the cluster's available capacity and uses the information on spare capacity, to optimize the number of hosts that can be patched simultaneously.

What's new in vSphere 5.0 Release Notes

- 32 way virtual SMP
- 1TB virtual machine RAM
- 3D graphics for Windows Aero support
- USB 3.0 device support
- UEFI virtual BIOS
- Support for up to 512 VM per host
- Max 2048 virtual CPUs per host
- Up to 160 logical CPUs and up to 2TB RAM (per host)
- Memory fault isolation — on supported platforms, ESXi 5.0 detects and quarantines physical memory regions that exhibit frequent correctable errors. This preemptive action reduces the risk of uncorrectable errors that result in VM or host downtime. Should an uncorrectable memory error occur, ESXi 5.0 quarantines the failed memory region and restarts the affected virtual machines. ESXi halts with a purple diagnostic screen only if the memory error affects the hypervisor itself.

vSphere Virtual Machine Administration (ESXi 5.0 & vCenter 5.0)

- Virtual Machine files:

File	Usage	Description
.vmx	*vmname*.vmx	Virtual machine configuration file
.vmxf	*vmname*.vmxf	Additional virtual machine configuration file
.vmdk	*vmname*.vmdk	Virtual disk characteristics
-flat.vmdk	*vmname*-flat.vmdk	Preallocated virtual disk
.nvram	*vmname*.nvram or nvram	Virtual Machine BIOS or EFI configuration
.vmsd	*vmname*.vmsd	Virtual machine snapshots
.vmsn	*vmname*.vmsn	Virtual machine snapshot data file
.vswp	*vmname*.vswp	Virtual machine swap file
.vmss	*vmname*.vmss	Virtual machine suspend file
.log	*vmname*.log	Current virtual machine log file
-#.log	*vmname*-#.log (where # is a number starting with 1)	Old virtual machine log entries

- The Mac OS X Server must run on Apple hardware. You cannot power on a Mac OS X Server if it is running on other hardware.

- A VM's name can be up to 80 characters long. If you are connected to vCenter Server and have folders in your inventory, names must be unique within the folder. Names are not case-sensitive, so the name my_vm is identical to My_Vm.
- Minimum memory size is 4MB for virtual machines that use BIOS firmware. Virtual machines that use EFI firmware require at least 96MB of RAM or they cannot power on.
- Maximum VM memory is:

VM version	Maximum memory
4	64GB
7	255 GB
8	1,011 GB

- To use vMotion for virtual machines with enabled NPIV, make sure that the RDM files of the virtual machines are located on the same datastore. You cannot perform Storage vMotion or vMotion between datastores (*Forbes*: I assume this should read "or cold migrate between datastores") when NPIV is enabled.
- A virtual machine with a physical compatibility RDM cannot be cloned, made into a template, or migrated if the migration involves copying the disk.
- A virtual RDM behaves as if it were a virtual disk, so you can use such features as taking a snapshot, cloning, and so on. When you clone the disk or make a template from it, the contents of the LUN are copied into a .vmdk virtual disk file. When you migrate a virtual compatibility mode RDM, you can migrate the mapping file or copy the contents of the LUN into a virtual disk.
- *Show all storage recommendations* — this option appears only when the virtual machine disks are stored on a datastore cluster and Storage DRS is enabled. When you select this option, the Virtual Machine Storage Placement Recommendations dialog box appears when you click *Continue*. The dialog box lists the datastores in the datastore cluster that are recommended for virtual machine placement.
- *Edit Storage DRS rules* — this option appears only when the virtual machine disks are stored on a datastore cluster. This option is selected when you select *Edit virtual hardware*. You can edit Storage DRS rules on the *Options* tab of the Virtual Machine Properties dialog box. When you select the *Edit Storage DRS rules* check box, the Storage DRS rules dialog box appears when you click *Continue*.
- Select the datastore location where you want to store the virtual machine files.
    - (Optional) Apply a virtual machine storage profile for the virtual machine home files and the virtual disks from the *VM Storage Profile* drop-down menu. The list of datastores shows which datastores are compatible and which are incompatible with the selected virtual machine storage profile.
    - (Optional) If you selected a datastore cluster and do not want to use Storage DRS with this virtual machine, select *Disable Storage DRS for this virtual machine* and select a datastore within the datastore cluster.
- The guest operating system being customized must be installed on a disk attached as SCSI node 0:0 in the virtual machine configuration.
- Customization of Linux guest operating systems requires that Perl is installed in the Linux guest operating system.
- As an alternative to specifying the computer name or IP addresses for virtual NICs during customization, you can create a script to generate these items.
- The script is executed by the Windows command-line script host (`cscript.exe`). The script can be written in any compatible scripting language, including JScript or VBScript. See the Microsoft documentation for cscript.exe for more information on writing scripts.
- Duplicate SIDs do not cause problems when the computers are part of a domain and only domain user accounts are used. However, if the computers are part of a Workgroup or local user accounts are used, duplicate SIDs can compromise file access controls
- If the new virtual machine encounters customization errors while it is booting, the errors are logged to %WINDIR%\temp\vmware-imc.
- View Virtual machine customization errors by opening /var/log/vmwareimc/toolsDeployPkg.log.
- You can change the administrator password only if the administrator password on the source Windows virtual machine is blank. If the source Windows virtual machine or template already has a password, the administrator password does not change.
- A custom sysprep answer file is a file that stores a number of customization settings such as computer name, licensing information, and workgroup or domain settings. You can supply a custom sysprep answer file as an alternative to specifying many of the settings in the Guest Customization wizard.
- Windows 2000, Windows Server 2003, and Windows XP use a text file called sysprep.inf. Windows Server 2008, Windows Vista, and Windows 7 use an XML file called sysprep.xml. You can create these files using a text editor, or use the Microsoft Setup Manager utility to generate them. For more information about how to create a custom sysprep answer file, see the documentation for the relevant operating system.
- The OVF format offers the following advantages:
    - OVF files are compressed, allowing for faster downloads.
    - The vSphere Client validates an OVF file before importing it, and ensures that it is compatible with the intended destination server. If the appliance is incompatible with the selected host, it cannot be imported and an error message appears.
    - OVF can encapsulate multi-tiered applications and more than one virtual machine.
- OVF template as a set of files (.ovf, .vmdk, and .mf) This format is optimal if you plan to publish the OVF files on a web server or image library.
- OVF template into a single .ova file. This might be convenient to distribute the OVF package as a single file if it needs to be explicitly downloaded from a web site or moved around using a USB key.
- The Sysprep tool is built in to the Windows Vista and Windows 2008 operating systems.
  C:\ALLUSERSPROFILE\Application Data\Vmware\VMware VirtualCenter\sysprep
  ...\1.1\

...\2k\
...\xp\
...\svr2003\
...\xp-64\
...\svr2003-64\

- In the vSphere Web Client, you can limit the number of simultaneous connections to a virtual machine and lock the guest operating system when the last remote user disconnects from the virtual machine console.
- You might change the guest operating system, for example, if you are upgrading the guest operating system installed in the virtual machine.
- When you set the guest operating system type for a new virtual machine, vCenter Server chooses configuration defaults based on the guest type. Changing the guest operating system type after the virtual machine is created does not retroactively change those settings. It affects the recommendations and setting ranges offered after the change.
- VMware multicore virtual CPU support lets you control the number of cores per virtual socket in a virtual machine. This capability lets operating systems with socket restrictions use more of the host CPU's cores, which increases overall performance.
- Using multicore virtual CPUs can be useful when you run operating systems or applications that can take advantage of only a limited number of CPU sockets. Previously, each virtual CPU was, by default, assigned to a single-core socket, so that the virtual machine would have as many sockets as virtual CPUs.
- Adding CPU resources to a running virtual machine with CPU hot plug enabled disconnects and reconnects all USB passthrough devices connected to that virtual machine.
    - For best results, use hardware version 8 virtual machines.
    - Hot-adding multicore virtual CPUs is supported only with hardware version 8 virtual machines.
    - To use the CPU hot-add feature with hardware version 7 virtual machines, set *Number of cores per socket* to 1.
- With CPU hot plug enabled and the virtual machine running, you can select the number of sockets to add from the *Number of virtual sockets* drop-down menu. The *Number of cores per socket* drop-down menu is unavailable and retains a value of 2.
- HT Sharing sharing mode:
    - *Any* (default) — The virtual CPUs of this virtual machine can share cores with other virtual CPUs of this or other virtual machines.
    - *None* — The virtual CPUs of this virtual machine have exclusive use of a processor core whenever they are scheduled to it. The other hyperthread of the core is halted while this virtual machine is using the core.
    - *Internal* — On a virtual machine with exactly two virtual processors, the two virtual processors are allowed to share one physical core (at the discretion of the host scheduler), but this virtual machine never shares a core with any other virtual machine. If this virtual machine has any other number of processors other than two, this setting is the same as the None setting.
- The *Scheduling Affinity* option gives you detailed control over how virtual machine CPUs are distributed across the host's physical cores (and hyperthreads if hyperthreading is enabled). This panel does not appear for virtual machines in a DRS cluster or when the host has only one processor core and no hyperthreading.
- Using CPU affinity, you can assign a virtual machine to a specific processor. This assignment allows you to restrict the assignment of virtual machines to a specific available processor in multiprocessor systems.
- Memory hot add requires that the virtual machine is using hardware version 7 or later.
- You can specify that all future memory allocations on a virtual machine use pages associated with a single NUMA node. The NUMA code is also known as manual memory affinity. When the virtual machine uses local memory, the performance improves on that virtual machine.
- The following conditions apply to memory optimization with NUMA:
    - The NUMA option is available on the Advanced Memory Resources page only if the host uses NUMA memory architecture.
    - Affinity settings are meaningful only when used to modify the performance of a specific set of virtual machines on one host. This option is not available when the virtual machine resides on a DRS cluster. All affinity values are cleared when you move the virtual machine to a new host.
    - You can specify nodes to use for future memory allocations only if you also specified CPU affinity. If you make manual changes only to the memory affinity settings, automatic NUMA rebalancing does not work properly.
    - Checking all the boxes is the same as applying no affinity.
- The following NIC types are supported:
    - *E1000* — Emulated version of the Intel 82545EM Gigabit Ethernet NIC, with drivers available in most newer guest operating systems, including Windows XP and later and Linux versions 2.4.19 and later.
    - *Flexible* — Identifies itself as a Vlance adapter when a virtual machine boots, but initializes itself and functions as either a Vlance or a VMXNET adapter, depending on which driver initializes it. With VMware Tools installed, the VMXNET driver changes the Vlance adapter to the higher performance VMXNET adapter.
    - *Vlance* — Emulated version of the AMD 79C970 PCnet32 LANCE NIC, an older 10 Mbps NIC with drivers available in most 32bit guest operating systems except Windows Vista and later. A virtual machine configured with this network adapter can use its network immediately.
    - *VMXNET* — Optimized for performance in a virtual machine and has no physical counterpart. Because operating system vendors do not provide built-in drivers for this card, you must install VMware Tools to have a driver for the VMXNET network adapter available.

    - *VMXNET 2 (Enhanced)* — Based on the VMXNET adapter but provides high-performance features commonly used on modern networks, such as jumbo frames and hardware offloads. VMXNET 2 (Enhanced) is available only for some guest operating systems on ESX/ESXi 3.5 and later.
    - *VMXNET 3* — Next generation of a paravirtualized NIC designed for performance. VMXNET 3 offers all the features available in VMXNET 2 and adds several new features, such as multiqueue support (also known as Receive Side Scaling in Windows), IPv6 offloads, and MSI/MSI-X interrupt delivery. VMXNET 3 is not related to VMXNET or VMXNET 2.
- Manually assigned MAC addresses for virtual machines on ESXi hosts must begin with the OUI 00:50:56. The address must have the form 00:50:56:XX:YY:ZZ, where XX is a hexadecimal number between 00 and 3F, and YY and ZZ are hexadecimal numbers between 00 and FF.
- You can set up virtual serial ports to send data in the following ways.
    - Physical serial port on the host — Sets the virtual machine to use a physical serial port on the host computer
    - Output to file — Sends output from the virtual serial port to a file on the host computer.
    - Connect to a named pipe — With this method, two virtual machines or a virtual machine and a process on the host can communicate as if they were physical machines connected by a serial cable
    - Connect over the network — Enables a serial connection to and from a virtual machine's serial port over the network. The Virtual Serial Port Concentrator (vSPC) aggregates traffic from multiple serial ports onto one management console. vSPC behavior is similar to physical serial port concentrators. Using a vSPC also allows network connections to a virtual machine's serial ports to migrate seamlessly when you use vMotion to migrate the virtual machine.
- When you use a physical serial port for serial port passthrough from an ESXi host to a virtual machine, the following conditions apply. The following are not supported:
    - Migration with vMotion
    - Serial ports present on add-on expansion cards might be supported by PCI DirectPath I/O.
    - Serial ports connected through USB are not supported for serial port passthrough. They might be supported by USB passthrough from an ESXi host to a virtual machine.
- A virtual machine can use up to four virtual serial ports.
- N-port ID virtualization (NPIV) provides the ability to share a single physical Fibre Channel HBA port among multiple virtual ports, each with unique identifiers. This capability lets you control virtual machine access to LUNs on a per-virtual machine basis.
- Each virtual port is identified by a pair of world wide names (WWNs): a world wide port name (WWPN) and a world wide node name (WWNN). These WWNs are assigned by vCenter Server.
- NPIV support is subject to the following limitations:
    - NPIV must be enabled on the SAN switch. Contact the switch vendor for information about enabling NPIV on their devices.
    - NPIV is supported only for virtual machines with RDM disks. Virtual machines with regular virtual disks continue to use the WWNs of the host's physical HBAs.
    - The physical HBAs on the ESXi host must have access to a LUN using its WWNs in order for any virtual machines on that host to have access to that LUN using their NPIV WWNs. Ensure that access is provided to both the host and the virtual machines.
    - The physical HBAs on the ESXi host must support NPIV. If the physical HBAs do not support NPIV, the virtual machines running on that host will fall back to using the WWNs of the host's physical HBAs for LUN access.
    - Each virtual machine can have up to 4 virtual ports. NPIV-enabled virtual machines are assigned exactly 4 NPIV-related WWNs, which are used to communicate with physical HBAs through virtual ports. Therefore, virtual machines can utilize up to 4 physical HBAs for NPIV purposes.
- To edit the virtual machine's WWNs, power off the virtual machine.
- NFS datastores with Hardware Acceleration and VMFS datastores support the following disk provisioning policies. On NFS datastores that do not support Hardware Acceleration, only thin format is available.
- Thick Provisioned Lazy Zeroed — Creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine. Using the default flat virtual disk format does not zero out or eliminate the possibility of recovering deleted files or restoring old data that might be present on this allocated space. You cannot convert a flat disk to a thin disk.
- Thick Provision Eager Zeroed — A type of thick virtual disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create other types of disks.
- Thin Provision — Use this format to save storage space. For the thin disk, you provision as much datastore space as the disk would require based on the value that you enter for the disk size. However, the thin disk starts small and at first, uses only as much datastore space as the disk needs for its initial operations.
- Disk mode:
    - *Dependent* — Dependent disks are included in snapshots.
    - *Independent-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.
    - *Independent-Nonpersistent* — Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset.

- Virtual machine storage profiles list the storage capabilities that virtual machine home files and virtual disks require to run the applications within the virtual machine.
- You can create a list of virtual machine storage profiles to define different levels of storage requirements.
- The virtual machine home files (.vmx, .vmsd, .nvram, .log, and so on) and the virtual disks (.vmdk) can have separate virtual machine storage profiles.
- The choice of SCSI controller does not affect whether your virtual disk is an IDE or SCSI disk. The IDE adapter is always ATAPI.
- PVSCSI controllers are available for virtual machines running hardware version 7 and later.
- PVSCSI controllers have the following limitations:
    - Hot add or remove requires a bus rescan from within the guest operating system.
    - Disks on PVSCSI controllers might not experience performance gains if they have snapshots or if memory on the ESXi host is over committed.
    - If you upgrade your Linux virtual machine to an unsupported kernel, you might not be able to access data on the disks attached to a PVSCSI controller.
    - MSCS clusters are not supported.
    - PVSCSI controllers do not support boot disks, the disk that contains the system software, on Red Hat Linux 5 virtual machines.
- vSphere DirectPath I/O allows a guest operating system on a virtual machine to directly access physical PCI and PCIe devices connected to a host. Each virtual machine can be connected to up to six PCI devices.
- Snapshots are not supported with PCI vSphere DirectPath I/O devices:
    - To use DirectPath, verify that the host has Intel® Virtualization Technology for Directed I/O (VT-d) or AMD I/O Virtualization Technology (IOMMU) enabled in the BIOS.
    - Verify that the PCI devices are connected to the host and marked as available for passthrough.
    - Verify that the virtual machine is using hardware version 7 or later.
- You can change the number of displays for a virtual machine, allocate memory for the displays, and enable 3D support.
- Some 3D applications require a minimum video memory of 64MB. Keep this in mind when you assign video memory.
- Use the Video Memory Calculator to calculate the required video memory based on the maximum number of displays, resolution, and color depth that the guest operating system must support, and click *OK*.
- A vService specifies a particular service on which vApps and virtual machines can depend.
- The vService configuration tab monitors and manages vService dependencies. This tab displays all the dependencies that a virtual machine or vApp has and each of their states. Each dependency shows the dependency name, description, requirement, bound status, and provider name.
- When you attach a USB device to a physical host, the device is available only to virtual machines that reside on that host. The device cannot connect to virtual machines that reside on another host in the datacenter.
- A USB device is available to only one virtual machine at a time. When a device is connected to a powered-on virtual machine, it is not available to connect to other virtual machines that run on the host. When you remove the active connection of a USB device from a virtual machine, it becomes available to connect to other virtual machines that run on the host.
- The USB passthrough autoconnect feature identifies the device by using the USB path of the device on the host.
- It uses the physical topology and port location, rather than the device identity. This feature can seem confusing if you expect the autoconnect feature to match the connection target by device ID.
- If the same device is plugged back in to the host through a different USB port, it cannot re-establish connection with the virtual machine. If you unplug the device from the host and plug in a different device to the same USB path, the new device appears and is connected to the virtual machine by the autoconnect feature that the previous device connection enabled.
- Autoconnect is useful in cases where devices mutate during usage. For example, for iPhones and other such devices, the device VID:PID changes during software or firmware upgrades. The upgrade process disconnects and reconnects the devices to the USB port.
- vSphere features available while using USB passthrough:

Feature	USB device passthrough support
DPM	No
DRS	Yes
FT	No
vMotion	Yes

- If a host with connected USB devices resides in a DRS cluster with DPM enabled, you must disable DPM for that host. Otherwise DPM might turn off the host with the device, which disconnects the device from the virtual machine.
- With USB passthrough from a host to a virtual machine, you can migrate a virtual machine to another ESXi host in the same datacenter and maintain the USB passthrough device connections to the original host.
- When you migrate a virtual machine with attached USB devices away from the host to which the devices are connected, the devices remain connected to the virtual machine. However, if you suspend or power off the virtual machine, the USB devices are disconnected and cannot reconnect when the virtual machine is resumed. The device connections can be restored only if you move the virtual machine back to the host to which the devices are attached.

- If you resume a suspended virtual machine that has a Linux guest operating system, the resume process might mount the USB devices at a different location on the file system.
- You can add two USB controllers to a virtual machine. The xHCI controller, available for Linux guest operating systems only, supports USB 3.0 superspeed, 2.0, and 1.1 devices. The EHCI+UHCI controller supports USB 2.0 and 1.1 devices.
- USB Controller support:

Controller type	Passthrough from ESXi to VM	Passthrough from Client to VM
EHCI & UHCI	Yes	Yes
xHCI (includes USB 3.0)	No	Yes (Linux guests only)

  - Drivers are not available for the xHCI controller on Windows guest operating systems.
- The following NICs support Wake on LAN:
  - Flexible (VMware Tools required).
  - vmxnet
  - Enhanced vmxnet
  - vmxnet 3
- The vApp metadata resides in the vCenter Server's database, so a vApp can be distributed across multiple ESXi hosts. This information can be lost if the vCenter Server database is cleared or if a standalone ESXi host that contains a vApp is removed from vCenter Server. You should back up vApps to an OVF package to avoid losing any metadata.
- vApp metadata for virtual machines within vApps do not follow the snapshots semantics for virtual machine configuration. So, vApp properties that are deleted, modified, or defined after a snapshot is taken remain intact (deleted, modified, or defined) after the virtual machine reverts to that snapshot or any prior snapshots.
- vApp IP Allocation Policy:
  - *Fixed* — IP addresses are manually configured. No automatic allocation is performed.
  - *Transient* — IP addresses are automatically allocated using IP pools from a specified range when the vApp is powered on. The IP addresses are released when the appliance is powered off.
  - *DHCP* — A DHCP server is used to allocate the IP addresses. The addresses assigned by the DHCP server are visible in the OVF environments of virtual machines started in the 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 virtual machines.
- A vSphere administrator uses the vCenter Solutions Manager to view the installalled solutions, view detailed information about the solutions, and monitor the solution health status.
- A solution is an extension of the vCenter Server that adds new functions to a vCenter Server instance. For example, vSphere ESX Agent Manager is a standard vCenter solution provided by VMware that allows you to manage ESX host agents that add new capabilities to ESX hosts. Another standard solution that vSphere provides is vService Manager. VMware products that integrate with vCenter Sever are also considered solutions. You can install a solution to add functionality from third-party technologies to the standard functions of vCenter Server. Solutions typically are delivered as OVF packages. You can install and deploy solutions from vSphere Client. Solutions can be integrated into the vCenter Solutions Manager.
- If a virtual machine or vApp is running a solution, an icon appears next to it in the inventory view of the vSphere Client. When you power on or power off a virtual machine or vApp, you are notified that you are performing this operation on an entity that is managed by the solution manager.
- The agent's health status is indicated by a specific color:
  - *Red* — The solution must intervene for the ESX Agent Manager to proceed.
  - *Yellow* — The ESX Agent Manager is actively working to reach a goal state.
  - *Green* — A solution and all its agents reached the goal state.
- A vService is a service or function that a solution provides to virtual machines and vApps. A solution can provide one or more vService. These vServices integrate with the platform and are able to change the environment in which the vApp or virtual machine runs.
- To access a virtual machine's console in the vSphere Web Client, you must install the Client Integration Plugin.
  - Windows 32 bit and 64 bit Internet Explorer 7 & 8 and Firefox 3.5 & 3.6
  - Linux 32 bit Firefox 3.5 & 3.6
- Snapshots are comprised of the following files:
  - vm-number.vmdk and vm-number-delta.vmdk. A collection of .vmdk and -delta.vmdk files for each virtual disk is connected to the virtual machine at the time of the snapshot. These files are referred to as child disks and delta disks. The child disks can later be considered parent disks for future child disks. From the original parent disk, each child constitutes a delta disk pointing back from the present state of the virtual disk, one step at a time, to the original. Delta disks are stored with the parent disk. The number value might not be consistent across all child disks from the same snapshot. The file names are chosen based on filename availability.
  - vm.vmsd. The .vmsd file is a database of the virtual machine's snapshot information and the primary source of information for the snapshot manager. The file contains line entries, which define the relationships between snapshots as well as the child disks for each snapshot.
  - vmSnapshotnumber.vmsn. These files are the active state at the time of the snapshot.

- The time it takes to commit or delete snapshots depends on how much data the guest operating system has written to the virtual disks since the last snapshot was taken. The required time is directly proportional to the amount of data (committed or deleted) and the amount of RAM allocated to the virtual machine.

vSphere Host Profiles (ESXi 5.0 & vCenter 5.0)

- The host profiles feature creates a profile that encapsulates the host configuration and helps to manage the host configuration.
- Host profiles eliminates per-host, manual, or UI-based host configuration and maintain configuration consistency and correctness across the datacenter by using host profile policies. These policies capture the blueprint of a known, validated reference host configuration and use this to configure networking, storage, security, and other settings on multiple hosts or clusters. You can then check a host or cluster against a profile's configuration for any deviations.
- A reference host is the host from which the profile is created.
- Host profiles is only supported for VMware vSphere 4.0 hosts or later. This feature is not supported for VMware Infrasturcture 3.5 or earlier hosts.
- While you can attach a host profile to a mixed cluster that contains VMware Infrastructure 3.5 or earlier hosts, the compliance check for those hosts fails.
- For hosts provisioned with Auto Deploy, vCenter Server owns the entire host configuration, which is captured in a host profile. In most cases, the host profile information is sufficient to store all configuration information. Sometimes the user is prompted for input when the host provisioned with Auto Deploy boots. The Auto Deploy answer file mechanism manages those cases.
- You can export a profile to a file that is in the VMware profile format (.vpf).
- When a host profile is exported, administrator passwords are not exported.
- Host Profile Sub-profile Configurations:

Sub-Profile Configuration	Example Policies and Compliance Checks
Memory reservation	Set memory reservation to a fixed value.
Storage	Configure NFS storage.
Networking	Configure virtual switch, port groups, physical NIC speed, security and NIC teaming policies, vNetwork Distributed Switch, and vNetwork Distributed Switch uplink port.
Date and Time	Configure time settings, timezone of server.
Firewall	Enable or disable a ruleset.
Security	Add a user or a usergroup, set root password.
Service	Configure settings for a service.
Advanced	Modify advanced options.

- When you need to create or edit a storage subprofile, use the vSphere CLI to configure or modify the NMP and PSA policies on a reference host first, and then extract the host profile from that host.
- Other profile configuration categories include: user group, authentication, kernel module, DCUI keyboard, host cache settings, SFCB, resource pools, login banner, SNMP agent, power system, and CIM indication subscriptions.
- When a host profile is attached to a cluster, the host or hosts within that cluster are also attached to the host profile. However, when the host profile is detached from the cluster, the association between the host or host within the cluster and that host profile remains.
- Once you create a host profile, you might need to make incremental updates to the profile. You can do this using two methods:
  - Make the configuration changes to the reference host in the vSphere Client, then update the profile from the reference host. The settings within the existing profile are updated to match those of the reference host.
  - Update the profile directly using the Profile Editor.
- For hosts provisioned with Auto Deploy, the answer file contains the user input policies for a host profile. The file is created when the profile is initially applied to a particular host.
- To apply a host profile to a host, the host must be placed into maintenance mode. During this process, the user is prompted to enter "answers" for policies that are specified during the host profiles creation.
- Placing the host into maintenance mode each time you apply a profile to the host can be costly and time consuming. A host provisioned with Auto Deploy can be rebooted while the host profile is attached to the host. After rebooting values stored in the answer file help the host provisioned with Auto Deploy to apply the profile. An answer file is created that contains a series of key value pairs for the user input options.
- The Answer File Status:
  - Incomplete
  - Complete
  - Unknown
- Hosts provisioned with Auto Deploy usually do not have sufficient local storage to save system logs. You can specify a remote syslog server for those hosts by setting up a reference host, saving the host profile, and applying that host profile to other hosts as needed. Best practice is to set up the syslog server on the reference host with the vSphere Client or the `esxcli system syslog` command and save the host profile.

vSphere Networking (ESXi & vCenter 5.0)

- 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 default number of logical ports for a standard switch is 120.
- For a port group to reach port groups located on other VLANs, the VLAN ID must be set to 4095.
- You can view IP and IPv6 routing information, such as network, prefix, and gateway, for a VMkernel network interface on a vSphere standard switch.
- Network resource pools allow you to manage network traffic by type of network traffic.
- Set the maximum number of ports on a host to limit the number of distributed ports that can exist on one or more hosts associated with a vSphere distributed switch.
- If you are changing the maximum number of ports for a host after the host is added to the distributed switch, you must restart the host before the new maximum takes effect.
- Port binding — Choose when ports are assigned to virtual machines connected to this distributed port group.
    - *Static binding* — assigns a port to a virtual machine when the virtual machine connects to the distributed port group. This option is not available when the vSphere Client is connected directly to ESXi.
    - *Dynamic binding* — assigns a port to a virtual machine the first time the virtual machine powers on after it is connected to the distributed port group. Dynamic binding is depricated in ESXi 5.0.
    - *Ephemeral* — for no port binding. This option is not available when the vSphere Client is connected directly to ESXi.
- Private VLANs are used to solve VLAN ID limitations and waste of IP addresses for certain network setups.
- 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, or *Community*, communicating with both promiscuous ports and other ports on the same secondary VLAN.
- To use private VLANs between a host and the rest of the physical network, the physical switch connected to the host needs to be private VLAN-capable and configured with the VLAN IDs being used by ESXi for the private VLAN functionality. For physical switches using dynamic MAC+VLAN ID based learning, all corresponding private VLAN IDs must be first entered into the switch's VLAN database.
- For each host associated with a vSphere distributed switch, you must assign at least one physical network adapter, or uplink, to the vSphere distributed switch.
- When network I/O control is enabled, distributed switch traffic is divided into the following predefined network resource pools: Fault Tolerance traffic, iSCSI traffic, vMotion traffic, management traffic, vSphere Replication (VR) traffic, NFS traffic, and virtual machine traffic.
- You can also create custom network resource pools for virtual machine traffic. You can control the priority that the traffic from each network resource pool is given by setting the physical adapter shares and host limit for each network resource pool.
- The iSCSI traffic resource pool shares do not apply to iSCSI traffic on a dependent hardware iSCSI adapter.
- Assigning a QoS priority tag to a network resource pool applies an 802.1p tag to all outgoing packets associated with that network resource pool.
- Network resource pools:
    - Custom shares — from 1 to 100
    - High — 100
    - Normal — 50
    - Low — 25
- You can change network resource pool settings such as allocated shares and limits for each network resource pool to change the priority network traffic for that network resource pool is given.
- Enable jumbo frames on a vSphere distributed switch or vSphere standard switch by changing the maximum transmission units (MTU). TCP Segmentation Offload (TSO) is enabled on the VMkernel interface by default, but must be enabled at the virtual machine level.
- 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 replacement might result in a change in the MAC address of the virtual network adapter.
- Enabling jumbo frame support on a virtual machine requires an enhanced vmxnet adapter for that virtual machine.
- NetQueue takes advantage of the ability of some network adapters to deliver network traffic to the system in multiple receive queues that can be processed separately, allowing processing to be scaled to multiple CPUs, improving receive-side networking performance.
- NetQueue is enabled by default.
- To disable NetQueue on a Host: `vicfg-advcfg --set false VMkernel.Boot.netNetQueueEnable`
- DirectPath I/O allows virtual machine access to physical PCI functions on platforms with an I/O Memory Management Unit.
- The following features are unavailable for virtual machines configured with DirectPath:
    - Hot adding and removing of virtual devices
    - Suspend and resume
    - Record and replay
    - Fault tolerance
    - High availability

    - DRS (limited availability — the virtual machine can be part of a cluster, but cannot migrate across hosts)
    - Snapshots
- vMotion is available for virtual machines configured with DirectPath only on Cisco UCS systems through supported Cisco distributed switches.
- You can configure passthrough networking devices on a host.
- The Passthrough Configuration page appears, listing all available passthrough devices. A green icon indicates that a device is enabled and active. An orange icon indicates that the state of the device has changed and the host must be rebooted before the device can be used.
- Adding a DirectPath device to a virtual machine sets memory reservation to the memory size of the virtual machine.
- Policies set at the standard switch or distributed port group level apply to all of the port groups on the standard switch or to ports in the distributed port group. The exceptions are the configuration options that are overridden at the standard port group or distributed port level.
- Load balancing and failover policies:
    - *Load Balancing policy* — determines how outgoing traffic is distributed among the network adapters associated with a switch or port group. Incoming traffic is controlled by the load balancing policy on the physical switch.
    - *Failover Detection* — controls the link status and beacon probing. Beaconing is not supported with guest VLAN tagging.
    - *Network Adapter Order* — can be active or standby.
- Load Balancing options for selecting an uplink:
    - Route based on the originating port ID
    - Route based on ip hash
    - Route based on source MAC hash
    - Use explicit failover order
    - Route based on physical NIC load (DVS Port Group only — choose an uplink based on the current loads of physical NICs)
- Do not use beacon probing with IP-hash load balancing.
- Do not use notify switches when the virtual machines using the port group are using Microsoft Network Load Balancing (NLB) in unicast mode. No such issue exists with NLB running in multicast mode.
- If you are using iSCSI Multipathing, your VMkernel interface must be configured to have one active adapter and no standby adapters.
- When using IP-hash load balancing, do not configure standby uplinks.
- VLAN policies: set the VLAN to:
    - *None* – don't use VLANing
    - *VLAN* – set a VLAN ID
    - *VLAN trunking* – set a VLAN trunk range
    - *Private VLAN*
- Secutity Policies:
    - *Promiscuous mode* — In nonpromiscuous 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 nonpromiscuous mode.
    - *MAC Address Changes* — If you set the MAC Address Changes to Reject and the guest operating system changes the MAC address of the adapter to anything other than what is in the .vmx configuration file, all inbound frames are dropped.
    - *Forged Transmits* — Reject means any outbound frame with a source MAC address that is different from the one currently set on the adapter are dropped.
- Traffic shaping policy is defined by average bandwidth, peak bandwidth, and burst size. You can establish a traffic shaping policy for each port group and each distributed port or distributed port group. ESXi shapes outbound network traffic on standard switches and inbound and outbound traffic on distributed switches. If you disable the policy, services have a free and clear connection to the physical network.
    - *Average Bandwidth* — The number of bits per second to allow across a port, averaged over Time.—the allowed average load.
    - *Peak Bandwidth* — Maximum number of bits per second to allow across a port when it is sending or receiving a burst of traffic. It can never be smaller than the average bandwidth.
    - *Burst Size* — Maximum number of bytes to allow in a burst (KB).
- The monitoring policy enables or disables NetFlow monitoring on a distributed port or port group.
- Port blocking policies allow you to selectively block ports from sending or receiving data.
- VDS Policies:
    - *Security* — Set MAC address changes, forged transmits, and promiscuous mode for the selected port groups.
    - *Traffic Shaping* — Set the average bandwidth, peak bandwidth, and burst size for inbound and outband traffic on the selected port groups.
    - *VLAN* — Configure how the selected port groups connect to physical VLANs.
    - *Teaming and Failover* — Set load balancing, failover detection, switch notification, and failover order for the selected port groups.
    - *Resource Allocation* — Set network resource pool association for the selected port groups. This option is available for vSphere distributed switch versions 5.0.0 and later only.
    - *Monitoring* — Enable or disable NetFlow on the selected port groups. This option is available for vSphere distributed switch versions 5.0.0 and later only.
    - *Miscellaneous* — Enable or disable port blocking on the selected port groups.

- IPv6 uses 128-bit addresses rather than the 32-bit addresses used by IPv4. This combats the problem of address exhaustion that is present with IPv4 and eliminates the need for network address translation. Other notable differences include link-local addresses that appear as the interface is initialized, addresses that are set by router advertisements, and the ability to have multiple IPv6 addresses on an interface.
- IPv6 is disabled by default.
- Configuring ESXi with VLANs is recommended for the following reasons.
    - It integrates the host into a pre-existing environment.
    - It secures network traffic.
    - It reduces network traffic congestion.
    - iSCSI traffic requires an isolated network.
- You can configure VLANs in ESXi using three methods: External Switch Tagging (EST), Virtual Switch Tagging (VST), and Virtual Guest Tagging (VGT).
    - *EST* — VLAN ID set to 0
    - *VST* — VLAN ID specified
    - *VGT* — all VLAN tagging is performed by the virtual machine.
- When using VGT, you must have an 802.1Q VLAN trunking driver installed on the virtual machine.
- Port mirroring allows you to mirror a distributed port's traffic to other distributed ports or specific physical switch ports. Port mirroring settings:
    - Select *Allow normal IO on destination ports* to allow normal IO traffic on destination ports. If you do not select this option, mirrored traffic is allowed out on destination ports, but no traffic is allowed in. So, do you want the destination port you define to handle normal I/O as well as receive mirror traffic?
    - Select *Encapsulation VLAN* to create a VLAN ID that encapsulates all frames at the destination ports. Useful if you are sending it outside your virtual environment.
    - If the original frames have a VLAN and *Preserve original VLAN* is not selected, the encapsulation VLAN replaces the original VLAN..
    - Select *Preserve original VLAN* to keep the original VLAN in the mirrored frames. This option is only available if you select *Encapsulation VLAN*.
    - Select *Mirrored packet length* to put a limit on the size of mirrored frames. If this option is selected, all mirrored frames are truncated to the specified length.
    - Select whether to use this source for *Ingress* or *Egress* traffic, or select *Ingress/Egress* to use this source for both types of traffic.
    - Type one or more port IDs or ranges of port IDs to add as source for the port mirroring session.
- NetFlow is a network analysis tool that you can use to monitor network monitoring and virtual machine traffic.
- With an IP address to the vSphere distributed switch, the NetFlow collector can interact with the vSphere distributed switch as a single switch, rather than interacting with a separate, unrelated switch for each associated host.
- Netflow settings:
    - *Active flow export timeout.*
    - *Idle flow export timeout*.
    - *Sampling Rate* — The sampling rate determines what portion of data NetFlow collects, with the sampling rate number determining how often NetFlow collects the packets. A collector with a sampling rate of 2 collects data from every other packet. A collector with a sampling rate of 5 collects data from every fifth packet.
    - *Process internal flows only* — collects data only on network activity between virtual machines on the same host.
- vSphere 5.0 supports Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP). CDP is available for vSphere standard switches and vSphere distributed switches connected to Cisco physical switches. LLDP is available for vSphere distributed switches version 5.0.0 and later.
- To circumvent the limit of 256 virtual network adapters per physical machine and possible MAC address conflicts between virtual machines, system administrators can manually assign MAC addresses. By default, VMware uses the Organizationally Unique Identifier (OUI) 00:50:56 for manually generated addresses, but all unique manually generated addresses are supported.
- You can set the addresses by adding the following line to a virtual machine‘s configuration file: ethernetnumber.address = 00:50:56:XX:YY:ZZ
- The value for XX must not be greater than 3F to avoid conflict with MAC addresses that are generated by the VMware Workstation and VMware Server products
- You must also set the option in a virtual machine’s configuration file: ethernetnumber.addressType="static"
- After the MAC address has been generated, it does not change unless the virtual machine is moved to a different location, for example, to a different path on the same server. The MAC address in the configuration file of the virtual machine is saved. All MAC addresses that have been assigned to network adapters of running and suspended virtual machines on a given physical machine are tracked.
- The MAC address of a powered off virtual machine is not checked against those of running or suspended virtual machines. It is possible that when a virtual machine is powered on again, it can acquire a different MAC address. This acquisition is caused by a conflict with a virtual machine that was powered on when this virtual machine was powered off.
- When using passthrough devices with a Linux kernel version 2.6.20 or earlier, avoid MSI and MSI-X modes because these modes have significant performance impact.
- The iSCSI initiator is enabled by Default in vSphere 5.

- Other blogger's points to highlight new features at the networking stack include (*Forbes*: Sorry, I can't find the reference for this):
  - LLDP (802.1AB) for dvSwitch (similar to Cisco CDP but non-proprietary)
  - NetFlow for dvSwitch (monitor application flow and measure application performance overtime)
  - DVMirror for dvSwitch (Port Mirroring capability to send copy of packets to different ports)
  - DVMirrror Encapsulation VLAN Session when destination is an uplink port
  - NIOC VM Traffic – per VM management

What's New in VMware vSphere 5.0 Networking Whitepaper

- All the networking capabilities discussed in this document are available only with the VMware vSphere Distributed Switch (Distributed Switch).
- There are two broad types of networking capabilities that are new or enhanced in the VMware vSphere 5.0 release. The first type improves the network administrator's ability to monitor and troubleshoot virtual infrastructure traffic:
  - NetFlow
  - Port mirror
- The second type focuses on enhancements to the network I/O control (NIOC):
  - User-defined resource pool
  - vSphere replication traffic type
  - IEEE 802.1p tagging
- NetFlow is a networking protocol that collects IP traffic information as records and sends them to a collector such as CA NetQoS for traffic flow analysis.
  - Intrahost virtual machine traffic (virtual machine–to–virtual machine traffic on the same host)
  - Interhost virtual machine traffic (virtual machine–to–virtual machine traffic on different hosts)
  - Virtual machine–physical infrastructure traffic
- Port mirroring is also referred to as Switch Port Analyzer (SPAN) on Cisco switches.
- Port mirroring provides visibility into:
  - Intrahost virtual machine traffic (virtual machine–to–virtual machine traffic on the same host)
  - Interhost virtual machine traffic (virtual machine–to–virtual machine traffic on different hosts)
- In the VMware vSphere 5.0 platform, NIOC supports traffic management capabilities for the following traffic types:
  - Virtual machine traffic
  - Management traffic
  - iSCSI traffic
  - NFS traffic
  - Fault-tolerant traffic
  - VMware vMotion™ traffic
  - *New* — User-defined traffic
  - *New* — vSphere replication traffic
- vSphere replication is a new system traffic type that carries replication traffic from one host to another.
- IEEE 802.1p is a standard for enabling QoS at MAC level. The IEEE 802.1p tag provides a 3-bit field for prioritization, which allows packets to be grouped into seven different traffic classes. Higher-number tags typically indicate critical traffic that has higher priority.

vSphere Storage(ESXi 5.0 & vCenter 5.0)

- You cannot use IDE/ATA drives to store virtual machines.
- ESXi does not support the delegate user functionality that enables access to NFS volumes using non-root credentials.
- Shared Serial Attached SCSI (SAS) stores virtual machines on direct-attached SAS storage systems that offer shared access to multiple hosts. This type of access permits multiple hosts to access the same VMFS datastore on a LUN.
- In the ESXi context, the term *target* identifies a single storage unit that the host can access.
- Different storage vendors present the storage systems to ESXi hosts in different ways. Some vendors present a single target with multiple storage devices or LUNs on it, while others present multiple targets with one LUN each.
- The iSCSI targets use iSCSI names, while Fibre Channel targets use World Wide Names (WWNs).
- A device, or LUN, is identified by its UUID name. If a LUN is shared by multiple hosts, it must be presented to all host with the same UUID.
- Depending on the type of storage, the ESXi host uses different algorithms and conventions to generate an identifier for each storage device:
  - *SCSI INQUIRY identifiers* — The host uses the SCSI INQUIRY command to query a storage device and uses the resulting data to generate a unique identifier. Device identifiers that are based on this information are unique across all hosts, do not change, and have one of the following formats:
    - naa.*number*
    - t10.*number*
    - eui.*number*

    These formats follow the T10 committee standards.

- *Path-based identifier* — When the device does not provide the information required above, the host generates an mpx.*path* name, where *path* represents the path to the device, for example, mpx.vmhba1:C0:T1:L3. This identifier can be used in the same way as the SCSI INQUIRY identifies. The mpx. identifier is created for local devices on the assumption that their path names are unique. However, this identifier is neither unique nor persistent and could change after every boot.
- *Legacy Identifier* — In addition to the SCSI INQUIRY or mpx. identifiers, for each device, ESXi generates an alternative legacy name. The identifier has the following format: *vml.number*. The legacy identifier includes a series of digits that are unique to the device and can be derived in part from the Page 83 information (*Forbes*: not sure what this is meant to refer), if it is available. For nonlocal devices that do not support Page 83 information, the vml. name is used as the only available unique identifier. You can use the `esxcli --server=server_name storage core device list` command to display all device names in the vSphere CLI.
- *Runtime Name* — In the vSphere Client, you can see the device identifier and a runtime name. The runtime name is generated by the host and represents the name of the first path to the device. It is not a reliable identifier for the device, and is not persistent. Typically, the path to the device has the following format: vmhba*Adapter*:C*Channel*:T*Target*:L*LUN*
  - vmhba*Adapter* is the name of the storage adapter. The name refers to the physical adapter on the host, not to the SCSI controller used by the virtual machines.
  - C*Channel* is the storage channel number. Software iSCSI initiators use the channel number to show multiple paths to the same target.
  - T*Target* is the target number. Target numbering is determined by the host and might change if the mappings of targets visible to the host change. Targets that are shared by different hosts might not have the same target number.
  - L*LUN* is the LUN number that shows the position of the LUN within the target. The LUN number is provided by the storage system. If a target has only one LUN, the LUN number is always zero (0).

- Supported network storage:

Technology	Protocols	Transfers	Interface
Fibre Channel	FC/SCSCI	Block access of data/LUN	FC HBA
Fibre Channel over Ethernet (FCoE)	FCoE/SCSI	Block access of data/LUN	○ Converged Network Adapter (hardware FCoE) ○ NIC with FCoE support (software FCoE)
iSCSI	IP/SCSI	Block access of data/LUN	○ iSCSI HBA or iSCSI-enabled NIC (hardware iSCSI) ○ Network adapter (software iSCSI)
NAS	IP/NFS	File (no direct LUN access)	Network adapter

- vSphere Features Supported by Storage:

Storage type	Boot VM	vMotion	Datastore	RDM	VM Cluster	HA and DRS	Data Protection APIs
Local Storage	Yes	No	VMFS	No	No	No	Yes
Fibre Channel	Yes	Yes	VMFS	Yes	Yes	Yes	Yes
iSCSI	Yes	Yes	VMFS	Yes	No	Yes	Yes
NAS over NFS	Yes	Yes	NFS	No	No	Yes	Yes

- *WWPN (World Wide Port Name)* — A globally unique identifier for a port that allows certain applications to access the port. The FC switches discover the WWPN of a device or host and assign a port address to the device.
- *Port_ID (or port address)* — Within a SAN, each port has a unique port ID that serves as the FC address for the port. This unique ID enables routing of data through the SAN to that port. The FC switches assign the port ID when the device logs in to the fabric. The port ID is valid only while the device is logged on.
- When N-Port ID Virtualization (NPIV) is used, a single FC HBA port (N-port) can register with the fabric by using several WWPNs. This method allows an N-port to claim multiple fabric addresses, each of which appears as a unique entity. When ESXi hosts use a SAN, these multiple, unique identifiers allow the assignment of WWNs to individual virtual machines as part of their configuration.
- Zoning defines which HBAs can connect to which targets.
- With ESXi hosts, use a single-initiator zoning or a single-initiator-single-target zoning.
- You should not mix FC HBAs from different vendors in a single server. Having different models of the same HBA is supported, but a single LUN cannot be accessed through two different HBA types, only through the same type. Ensure that the firmware level on each HBA is the same.
- A software FCoE adapter is a software code that performs some of the FCoE processing. This adapter is used with a NIC that supports partial FCoE offload.
- For the software FCoE adapter, you must properly configure networking and then activate the adapter.
- Follow these guidelines when you configure a network switch for software FCoE environment:
  - On the ports that communicate with your ESXi host, disable the Spanning Tree Protocol (STP). Having the STP enabled might delay the FCoE Initialization Protocol (FIP) response at the switch and cause the all paths down (APD) condition. The FIP is a protocol that FCoE uses to discover and initialize FCoE entities on the Ethernet.
  - Turn on Priority-based Flow Control (PFC) and set it to AUTO.
- Before you activate the software FCoE adapters, you need to connect the VMkernel to physical FCoE NICs installed on your host.
- ESXi 5.0 supports a maximum of four network adapter ports used for software FCoE.
- N-Port ID Virtualization (NPIV) — When a virtual machine has a WWN assigned to it, the virtual machine's configuration file (.vmx) is updated to include a WWN pair (consisting of a World Wide Port Name, WWPN, and a World Wide Node Name, WWNN). As that virtual machine is

powered on, the VMkernel instantiates a virtual port (VPORT) on the physical HBA which is used to access the LUN. The VPORT is a virtual HBA that appears to the FC fabric as a physical HBA, that is, it has its own unique identifier, the WWN pair that was assigned to the virtual machine. Each VPORT is specific to the virtual machine, and the VPORT is destroyed on the host and it no longer appears to the FC fabric when the virtual machine is powered off. When a virtual machine is migrated from one host to another, the VPORT is closed on the first host and opened on the destination host.

- If NPIV is enabled, WWN pairs (WWPN & WWNN) are specified for each virtual machine at creation time. When a virtual machine using NPIV is powered on, it uses each of these WWN pairs in sequence to try to discover an access path to the storage. The number of VPORTs that are instantiated equals the number of physical HBAs present on the host. A VPORT is created on each physical HBA that a physical path is found on. Each physical path is used to determine the virtual path that will be used to access the LUN. Note that HBAs that are not NPIV-aware are skipped in this discovery process because VPORTs cannot be instantiated on them.
- NPIV can be used only for virtual machines with RDM disks.
- HBAs on your host must support NPIV.
- The switches in the fabric must be NPIV-aware.
- NPIV supports vMotion.
- If your FC SAN environment supports concurrent I/O on the disks from an active-active array, the concurrent I/O to two different NPIV ports is also supported.
- NPIV does not support Storage vMotion.
- If you want to use vMotion for a virtual machine with enabled NPIV, make sure that the RDM file is located on the same datastore where the virtual machine configuration file resides.
- Path thrashing only occurs on active-passive-arrays. Configure the path to use the Most Recently Used PSP (the default).
- iSCSI ports are end-points of an iSCSI session. Each node can be identified in a number of ways.
    - *IP Address* — Each iSCSI node can have an IP address associated with it
    - *iSCSI Name* — A worldwide unique name for identifying the node. iSCSI uses the iSCSI Qualified Name (IQN), Extended Unique Identifier (EUI), and Network Address Authority (NAA) formats. By default, ESXi generates unique iSCSI names for your iSCSI initiators
    - *iSCSI Alias* — A more manageable name for an iSCSI device or port used instead of the iSCSI name. iSCSI aliases are not unique and are intended to be just a friendly name to associate with a port.
- iSCSI names are formatted in two different ways. The most common is the IQN format.
    - The IQN format takes the form iqn.yyyy-mm.naming-authority:unique name
    - The EUI format takes the form eui.16 hex digits. The top 24 bits are a company ID that IEEE registers with a particular company. The lower 40 bits are assigned by the entity holding that company ID and must be unique.
- Dependent Hardware iSCSI Adapter — Depends on VMware networking, and iSCSI configuration and management interfaces provided by VMware. This type of adapter can be a card that presents a standard network adapter and iSCSI offload functionality for the same port. The iSCSI offload functionality depends on the host's network configuration to obtain the IP, MAC, and other parameters used for iSCSI sessions.
- Independent Hardware iSCSI Adapter — Implements its own networking and iSCSI configuration and management interfaces.
- *Active-active storage system* — Allows access to the LUNs simultaneously through all the storage ports that are available without significant performance degradation. All the paths are active at all times, unless a path fails.
- *Active-passive storage system* — A system in which one storage processor is actively providing access to a given LUN. The other processors act as backup for the LUN and can be actively providing access to other LUN I/O. I/O can be successfully sent only to an active port for a given LUN. If access through the active storage port fails, one of the passive storage processors can be activated by the servers accessing it.
- *Asymmetrical storage system* — Supports Asymmetric Logical Unit Access (ALUA). ALUA-complaint storage systems provide different levels of access per port. ALUA allows hosts to determine the states of target ports and prioritize paths.
- *Virtual port storage system (iSCSI)* — Allows access to all available LUNs through a single virtual port. These are active-active storage devices, but hide their multiple connections though a single port. The ESXi multipathing cannot detect the multiple connections to the storage. These storage systems handle port failover and connection balancing transparently. This is often referred to as transparent failover.
- The two types of discovery available on ESXi are dynamic and static. 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.
- Access control is a policy set up on the iSCSI storage system. Three types:
    - Initiator name
    - IP address
    - CHAP protocol
- You can activate only one software iSCSI adapter.
- If you use a vSphere distributed switch with multiple uplink ports, for port binding, create a separate distributed port group per each physical NIC. Then set the team policy so that each distributed port group has only one active uplink port.
- For each virtual adapter on a vSphere standard switch, all network adapters appear as active. You must override this setup, so that each VMkernel interface maps to only one corresponding active NIC. For example, vmk1 maps to vmnic1, vmk2 maps to vmnic2, and so on. After you perform this task, bind the virtual VMkernel adapters to the software iSCSI or dependent hardware iSCSI adapters.
- Set up target discovery addresses so that the iSCSI adapter can determine which storage resource on the network is available for access.

    - *Dynamic Discovery* — Also known as SendTargets discovery. Each time the initiator contacts a specified iSCSI server, the initiator sends the SendTargets request to the server. The server 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. The initiator has a list of targets it can contact and uses their IP addresses and target names to communicate with them.
- ESXi supports the following CHAP authentication methods:
    - *One-way CHAP* — In one-way CHAP authentication, also called unidirectional, the target authenticates the initiator, but the initiator does not authenticate the target.
    - *Mutual CHAP* — In mutual CHAP authentication, also called bidirectional, an additional level of security enables the initiator to authenticate the target. VMware supports this method for software and dependent hardware iSCSI adapters only.
- The CHAP name should not exceed 511 alphanumeric characters and the CHAP secret should not exceed 255 alphanumeric characters
- With independent hardware iSCSI only, you can place the diagnostic partition on the boot LUN. If you configure the diagnostic partition in the boot LUN, this LUN cannot be shared across multiple hosts. If a separate LUN is used for the diagnostic partition, it can be shared by multiple hosts.
- If your ESXi host uses an independent hardware iSCSI adapter, such as QLogic HBA, you need to configure the adapter to boot from the SAN.
- ESXi hosts can boot from an iSCSI SAN using the software or dependent hardware iSCSI adapters and network adapters.
- To deploy ESXi and boot from the iSCSI SAN, the host must have an iSCSI boot capable network adapter that supports the iSCSI Boot Firmware Table (iBFT) format. The iBFT is a method of communicating parameters about the iSCSI boot device to an operating system.
- With VMFS5, you can have up to 256 VMFS datastores per host, with the minimum size of 1.3GB and the maximum of 64TB.
- Up to 128 hosts to a single VMFS5 volume.
- VMFS supports SCSI reservations and atomic test and set (ATS) locking.
- VMFS uses SCSI reservations on storage devices that do not support hardware acceleration. SCSI reservations lock an entire storage device while an operation that requires metadata protection is performed. Because this lock is exclusive, excessive SCSI reservations by a host can cause performance degradation on other hosts that are accessing the same VMFS.
- For storage devices that support hardware acceleration, VMFS uses the ATS algorithm, also called hardware assisted locking. In contrast with SCSI reservations, ATS supports discrete locking per disk sector.
- When a storage device contains a VMFS datastore copy, you can mount the datastore with the existing signature or assign a new signature. Each VMFS datastore created in a storage disk has a unique UUID that is stored in the file system superblock. When the storage disk is replicated or snapshotted, the resulting disk copy is identical, byte-for-byte, with the original disk.
- If you do not need to resignature a VMFS datastore copy, you can mount it without changing its signature. You can keep the signature if, for example, you maintain synchronized copies of virtual machines at a secondary site as part of a disaster recovery plan.
- You can mount a VMFS datastore copy only if it does not collide with the original VMFS datastore that has the same UUID. To mount the copy, the original VMFS datastore has to be offline.
- Use datastore resignaturing if you want to retain the data stored on the VMFS datastore copy. When resignaturing a VMFS copy, ESXi assigns a new UUID and a new label to the copy, and mounts the copy as a datastore distinct from the original.
- When you perform VMFS datastore management operations, such as creating a VMFS datastore or RDM, adding an extent, and increasing or deleting a VMFS datastore, your host or the vCenter Server automatically rescans and updates your storage.
- Perform the manual rescan each time you make one of the following changes.
    - Zone a new disk array on a SAN.
    - Create new LUNs on a SAN.
    - Change the path masking on a host.
    - Reconnect a cable.
    - Change CHAP settings (iSCSI only).
    - Add or remove discovery or static addresses (iSCSI only).
    - Add a single host to the vCenter Server after you have edited or removed from the vCenter Server a datastore shared by the vCenter Server hosts and the single host.
- By default, the VMkernel scans for LUN 0 to LUN 255 for every target (a total of 256 LUNs). You can modify the *Disk.MaxLUN* parameter to improve LUN discovery speed.
- You can disable the default sparse LUN support to decrease the time ESXi needs to scan for LUNs.
- If all LUNs that your storage system presents are sequential, you can disable the sparse LUN support.
- When you perform VMFS datastore management operations, vCenter Server uses default storage protection filters. The filters help you to avoid storage corruption by retrieving only the storage devices that can be used for a particular operation. Unsuitable devices are not displayed for selection. You can turn off the filters to view all devices.
- Your host can determine whether a device disconnection is a temporary, all-paths-down event or whether a permanent device disconnection occurred. You can perform a storage reconfiguration in which you detach a datastore and later reattach it. If your host detects an unplanned device loss, ESXi marks the storage device as permanently unavailable to conserve resources and memory.
- Although the ESXi host cannot determine the reason for a device loss, the host supports the loss detection. When the device becomes permanently unavailable, ESXi receives sense codes from storage arrays and recognizes that the device is permanently lost, not just

temporarily unavailable. ESXi marks the device as not connected and a warning about the device being permanently unavailable appears in the VMkernel log file. Typically, unplanned device loss is unintentional and can occur when a storage device is unmapped, removed, or its unique ID changes, or when there is an unrecoverable hardware error.

- To run successfully, your host must have a diagnostic partition or a dump partition to store core dumps for debugging and technical support.
- Typically, a local diagnostic partition is created during ESXi installation. You can override this default behavior if, for example, you use shared storage devices instead of local storage. To prevent automatic disk formatting, detach the local storage devices from the host before you install ESXi and power on the host for the first time.
- You can later create a diagnostic partition on a local disk or on a private or shared SAN LUN using the vSphere Client.
- A diagnostic partition cannot be located on an iSCSI LUN accessed through the software iSCSI or dependent hardware iSCSI adapter.
- Unless you are using diskless servers, set up a diagnostic partition on a local storage.
- Each host must have a diagnostic partition of 110MB. If multiple hosts share a diagnostic partition on a SAN LUN, the partition should be large enough to accommodate core dumps of all hosts.
- If a host that uses a shared diagnostic partition fails, reboot the host and extract log files immediately after the failure. Otherwise, the second host that fails before you collect the diagnostic data of the first host might not be able to save the core dump.
    - Click *Datastores* and click *Add Storage*.
    - Select *Diagnostic* and click *Next*.
    - If you do not see *Diagnostic* as an option, the host already has a diagnostic partition.
        - *Private Local* — Creates the diagnostic partition on a local disk
        - *Private SAN Storage* — Creates the diagnostic partition on a non-shared SAN LUN
        - *Shared SAN Storage* — Creates the diagnostic partition on a shared SAN LUN
- When you unmount a datastore, it remains intact, but can no longer be seen from the hosts that you specify. The datastore continues to appear on other hosts, where it remains mounted.
- When you unmount a datastore, it remains intact, but can no longer be seen from the hosts that you specify.
- The datastore continues to appear on other hosts, where it remains mounted. You cannot unmount an active mounted datastore.
- To unmount a datastore, verify that all of the following conditions are met:
    - No virtual machines reside on the datastore.
    - The datastore is not part of a datastore cluster.
    - The datastore is not managed by Storage DRS.
    - Storage I/O control is disabled for this datastore.
    - The datastore is not used for vSphere HA heartbeating.
- If you are using the RDM in physical compatibility mode, you cannot use a snapshot with the disk.
- Feature available with VMDKs and RDMs:

Feature	VMDK	Virtual Mode RDM	Physical Mode RDM
SCSI commands passed through	No	No	Yes (REPORT LUNs is not passed through)
vCenter support	Yes	Yes	Yes
Snapshots	Yes	Yes	No
Distributed locking	Yes	Yes	Yes
Clustering (MSCS)	Cluster in a box (CIB)	Cluster in a box (CIB) & Cluster across boxes (CAB)	Physcial to Virtual (n+1) Cluster across boxes (CAB)
SCSI target based software	No	No	Yes

- You can use PSA SATP claim rules to tag SSD devices that are not detected automatically.
- In addition to path failover, multipathing provides load balancing.
- To manage storage multipathing, ESXi uses a collection of Storage APIs, also called the Pluggable Storage Architecture (PSA). The PSA is an open, modular framework that coordinates the simultaneous operation of multiple multipathing plug-ins (MPPs). The PSA allows 3rd party software developers to design their own load balancing techniques and failover mechanisms for particular storage array, and insert their code directly into the ESXi storage I/O path.
- The VMkernel multipathing plug-in that ESXi provides by default is the VMware Native Multipathing Plug-In (NMP). The NMP is an extensible module that manages sub plug-ins. There are two types of NMP sub plugins, Storage Array Type Plug-Ins (SATPs), and Path Selection Plug-Ins (PSPs). SATPs and PSPs can be built-in and provided by VMware, or can be provided by a third party.
- By default, the VMware NMP supports the following PSPs:
    - *VMW_PSP_MRU* — MRU is the default policy for most active-passive storage devices.
    - *VMW_PSP_FIXED* — Fixed is the default policy for most active-active storage devices.
    - *VMW_PSP_RR* — The host uses an automatic path selection algorithm rotating through all active paths when connecting to active-passive arrays, or through all available paths when connecting to active-active arrays. RR is the default for a number of arrays and can be used with both active-active and active-passive arrays to implement load balancing across paths for different LUNs.
- Based on a set of claim rules defined in the /etc/vmware/esx.conf file, the host determines which multipathing plug-in (MPP) should claim the paths to a particular device and become responsible for managing the multipathing support for the device.
- By default, the host performs a periodic path evaluation every 5 minutes causing any unclaimed paths to be claimed by the appropriate MPP.

- For the paths managed by the NMP module, a second set of claim rules is applied. These rules determine which Storage Array Type Plug-In (SATP) should be used to manage the paths for a specific array type, & which Path Selection Plug-In (PSP) is to be used for each storage device.
- If you are using the Fixed path policy, you can see which path is the preferred path. The preferred path is marked with an asterisk (*) in the Preferred column.
- The following considerations help with multipathing:
  - If no SATP is assigned to the device by the claim rules, the default SATP for iSCSI or FC devices is VMW_SATP_DEFAULT_AA. The default PSP is VMW_PSP_FIXED.
  - When you search the SATP rules to locate a SATP for a given device, the NMP searches the driver rules first. If there is no match, the vendor/model rules are searched, and finally the transport rules are searched. If no match occurs, NMP selects a default SATP for the device.
  - If VMW_SATP_ALUA is assigned to a specific storage device, but the device is not ALUA-aware, no claim rule match occurs for this device. The device is claimed by the default SATP based on the device's transport type.
  - The default PSP for all devices claimed by VMW_SATP_ALUA is VMW_PSP_MRU. The VMW_PSP_MRU selects an active/optimized path as reported by the VMW_SATP_ALUA, or an active/unoptimized path if there is no active/optimized path. This path is used until a better path is available (MRU). For example, if the VMW_PSP_MRU is currently using an active/unoptimized path and an active/optimized path becomes available, the VMW_PSP_MRU will switch the current path to the active/optimized one.
  - If you enable VMW_PSP_FIXED with VMW_SATP_ALUA, the host initially makes an arbitrary selection of the preferred path, regardless of whether the ALUA state is reported as optimized or unoptimized. As a result, VMware does not recommend to enable VMW_PSP_FIXED when VMW_SATP_ALUA is used for an ALUA-compliant storage array. The exception is when you assign the preferred path to be to one of the redundant storage processor (SP) nodes within an active-active storage array. The ALUA state is irrelevant.
- The hardware acceleration functionality enables the ESXi host to integrate with compliant storage arrays and offload specific virtual machine and storage management operations to storage hardware. With the storage hardware assistance, your host performs these operations faster and consumes less CPU, memory, and storage fabric bandwidth.
- Hardware Acceleration Storage Requirements:

ESXi Version	Block Storage Devices	NAS Devices
ESX/ESXi version 4.1	Support block storage plug-ins for array integration (VAAI)	Not supported
ESXi version 5.0	Support T10 SCSI standard or block storage plug-ins for array integration (VAAI)	Support NAS plug-ins for array Integration

- Block Storage — ESXi hardware acceleration supports the following array operations:
  - Full copy, also called clone blocks or copy offload. Enables the storage arrays to make full copies of data within the array without having the host read and write the data.
  - Block zeroing, also called write same. Enables storage arrays to zero out a large number of blocks to provide newly allocated storage, free of previously written data
  - Hardware assisted locking, also called Atomic Test and Set (ATS). Supports discrete virtual machine locking without use of SCSI reservations. This operation allows disk locking per sector, instead of the entire LUN as with SCSI reservations.
- On your host, the hardware acceleration for block storage devices is enabled by default. You can use the vSphere Client advanced settings to disable the hardware acceleration operations.
- NAS Devices — Hardware Acceleration operations:
  - File clone. This operation is similar to the VMFS block cloning except that NAS devices clone entire files instead of file segments.
  - Reserve space. Enables storage arrays to allocate space for a virtual disk file in thick format.
  - Extended file statistics. Enables storage arrays to accurately report space utilization for virtual machines.
- SSD enablement results in several benefits:
  - It enables usage of SSD as swap space for improved system performance.
  - It increases virtual machine consolidation ratio as SSDs can provide very high I/O throughput.
  - It supports identification of virtual SSD device by the guest operating system.
- You can use PSA SATP claim rules to tag devices that cannot be auto-detected.
- With ESXi, you can use two models of thin provisioning, array-level and virtual disk-level.
- NFS datastores with Hardware Acceleration and VMFS datastores support the following disk provisioning policies. On NFS datastores that do not support Hardware Acceleration, only thin format is available.
- When allocating storage for your VMs, the following space allocation information is available in the Resources section:
  - *Provisioned Storage* – Shows datastore space guaranteed to the virtual machine. The entire space might not be used by the virtual machine if it has disks in thin provisioned format. Other virtual machines can occupy any unused space.
  - *Not-shared Storage* – Shows datastore space occupied by the virtual machine and not shared with any other virtual machines.
  - *Used Storage* – Shows datastore space actually occupied by virtual machine files, including configuration and log files, snapshots, virtual disks, and so on. When the virtual machine is running, the used storage space also includes swap files.
- When you use the Storage APIs — Array Integration, the host can integrate with physical storage and become aware of underlying thin-provisioned LUNs and their space usage.

- Using thin provision integration, your host can perform these tasks:
  - Monitor the use of space on thin-provisioned LUNs to avoid running out of physical space. As your datastore grows or if you use Storage vMotion to migrate virtual machines to a thin-provisioned LUN, the host communicates with the LUN and warns you about breaches in physical space and about out-ofspace conditions.
  - Inform the array about the datastore space that is freed when files and RDMs are deleted or removed from the datastore by Storage vMotion. The array can then reclaim the freed blocks of space.
- The vendor provider is a software plug-in developed by a third party through the Storage APIs – Storage Awareness. The vendor provider component is typically installed on the storage array side and acts as a server in the vSphere environment. The vCenter Server uses vendor providers to retrieve information about:
  - Storage topology. (it includes such data as storage vendors, array IDs, processors, ports, and so on.)
  - Storage capabilities.
  - Storage status
- A storage capability outlines the quality of service that a storage system can deliver. It is a guarantee that the storage system can provide a specific set of characteristics for capacity, performance, availability, redundancy, and so on.
- If a storage system uses Storage APIs — Storage Awareness, it informs vCenter Server that it can guarantee a specific set of storage features by presenting them as a storage capability. vCenter Server recognizes the capability and adds it to the list of storage capabilities in the Manage Storage Capabilities dialog box. Such storage capabilities are system-defined. vCenter Server assigns the system-defined storage capability to each datastore that you create from that storage system.
- `vmkfstools` is one of the ESXi Shell commands for managing VMFS volumes and virtual disks. You can perform many storage operations using this command.
- The following was taken from Duncan Epping's VMFS5 post: http://www.yellow-bricks.com/2011/07/13/vsphere-5-0-what-has-changed-for-vmfs/
  - VMFS-5 uses GPT instead of MBR
  - VMFS-5 supports volumes up to 64TB
  - This includes Pass-through RDMs!
  - VMFS-5 uses a Unified Blocksize –> 1MB
  - VMFS-5 uses smaller Sub-Blocks
  - ~30.000 8KB blocks versus ~3000 64KB blocks with VMFS-3
  - VMFS-5 has support for very small files (1KB)
  - Non-disruptive upgrade from VMFS-3 to VMFS-5
  - ATS locking enhancements (as part of VAAI)

What's New in VMware vSphere 5.0 Storage Whitepaper

- 64TB device support.
- Unified block size.
- Improved subblock mechanism.
- vSphere 5.0 facilitates a nondisruptive upgrade from VMFS-3 to VMFS-5.
- Storage DRS provides initial placement and ongoing balancing recommendations.
- I/O load is evaluated by default every 8 hours.
- I/O latency threshold is 15ms by default.
- Storage DRS offers three types of affinity rules:
  - *VMDK Anti-Affinity* — Virtual disks of a virtual machine with multiple virtual disks are placed on different datastores.
  - *VMDK Affinity* — Virtual disks are kept together on the same datastore.
  - *VM Anti-Affinity* — Two specified virtual machines, including associated disks, are placed on different datastores.
- Storage DRS offers Datastore Maintenance Mode, which automatically evacuates all virtual machines and virtual disk drives from the selected datastore to the remaining datastores in the datastore cluster.
- Storage DRS works with both VMFS- and NFS-based datastores. But mixing VMFS and NFS datastores in a single datastore cluster is currently not supported.
- With vSphere 5.0, support for the VAAI primitives has been enhanced and additional primitives have been introduced:
  - Thin Provisioning
  - Hardware acceleration for NAS
- VAAI Thin Provisioning introduces the following:
  - Dead Space Reclamation informs the array about the datastore space that is freed when files are deleted or removed from the datastore by Storage vMotion. The array can then reclaim the freed blocks of space.
  - Out-of-Space Conditions monitors the space usage on thin-provisioned LUNs to prevent running out of physical space.
- Hardware Acceleration for NAS will enable faster provisioning and the use of thick virtual disks through two newly introduced VAAI primitives:
  - Full File Clone – Similar to Full Copy. Enables virtual disks to be cloned by the NAS device
  - Reserve Space – Enables creation of thick virtual disk files on NAS

- Storage vMotion in vSphere 5.0 now also supports the migration of virtual machines with a vSphere snapshot and the migration of linked clones.
- Mirror Mode enables a single-pass block copy of the source disk to the destination disk by mirroring I/Os of copied blocks.
- When the guest OS of the virtual machine that is undergoing the process using Storage vMotion initiates a write to an already copied block, the mirror driver will synchronously mirror this write and wait for both acknowledgements before communicating this to the guest OS.

VMware vSphere Blog post (Storage features part 1 — VMFS-5)

http://blogs.vmware.com/vsphere/2011/07/new-vsphere-50-storage-features-part-1-vmfs-5.html

- VMFS-5 Enhancements:
  - Unified 1MB File Block Size.
  - Large Single Extent Volumes limit has been increased to ~ 60TB.
  - Smaller Sub-Block. now 8KB rather than the 64KB.
  - Small File Support For files less than or equal to 1KB, VMFS-5 uses the file descriptor location in the metadata for storage rather than file blocks. When they grow above 1KB, these files will then start to use the new 8KB sub blocks.
  - Increased File Count. VMFS-5 introduces support for greater than 100,000 files.
  - ATS Enhancement. This Hardware Acceleration primitive, Atomic Test & Set (ATS), is now used throughout VMFS-5 for file locking. ATS is part of the VAAI (vSphere Storage APIs for Array Integration).
- VMFS-3 to VMFS-5 Upgrades:
  - Upgrading from VMFS-3 to VMFS-5 is an online & non-disruptive upgrade operation, i.e. VMs can continue to run on the datastore.
  - Upgraded VMFS-5 can use the new 1KB small-files feature.
  - Upgraded VMFS-5 can be grown to ~ 60TB, same as a newly created VMFS-5.
  - Upgraded VMFS-5 has all the VAAI ATS improvements that a newly created VMFS-5 has.
- Differences between newly created and upgraded VMFS-5 datastores:
  - VMFS-5 upgraded from VMFS-3 continues to use the previous file block size which may be larger than the unified 1MB file block size.
  - VMFS-5 upgraded from VMFS-3 continues to use 64KB sub-blocks and not new 8K sub-blocks.
  - VMFS-5 upgraded from VMFS-3 continues to have a file limit of 30720 rather than new file limit of > 100000 for newly created VMFS-5.
  - VMFS-5 upgraded from VMFS-3 continues to use MBR (Master Boot Record) partition type; when the VMFS-5 volume is grown above 2TB, it automatically & seamlessly switches from MBR to GPT (GUID Partition Table) with no impact to the running VMs.
  - VMFS-5 upgraded from VMFS-3 continue to have its partition starting on sector 128; newly created VMFS5 partitions will have their partition starting at sector 2048.
- RDM — Raw Device Mappings:
  - Support for passthru RDMs to be ~ 60TB in size.
  - Non-passthru RDMs are still limited to 2TB — 512 bytes.
  - Both upgraded VMFS-5 & newly created VMFS-5 support the larger passthru RDM.
- Misc:
  - Maximum size of a VMDK on VMFS-5 is still 2TB -512 bytes.
  - Maximum size of a non-passthru (virtual) RDM on VMFS-5 is still 2TB -512 bytes.
  - Maximum number of LUNs that are supported on an ESXi 5.0 host is still 256.

vSphere Security (ESXi & vCenter 5.0)

- From a security perspective, ESXi consists of three major components: the virtualization layer, the virtual machines, and the virtual networking layer.
- The VMkernel is fully dedicated to supporting virtual machines and is not used for other purposes, the interface to the VMkernel is strictly limited to the API required to manage virtual machines. ESXi provides additional VMkernel protection with the following features:
  - Memory Hardening — The ESXi kernel, user-mode applications, and executable components such as drivers and libraries are located at random, non-predictable memory addresses. This combines with non-executable memory protections in CPUs.
  - Kernel Module Integrity — Digital signing ensures the integrity and authenticity of modules, drivers and applications as they are loaded by the VMkernel.
  - Trusted Platform Module (TPM)- This module is a hardware element that represents the core of trust for a hardware platform and enables attestation of the boot process, as well as cryptographic key storage and protection. Each time ESXi boots, TPM measures the VMkernel with which ESXi booted in one of its Platform Configuration Registers (PCRs). TPM measurements are propagated to vCenter Server when the host is added to the vCenter Server system.
    You can use TPM with third-party solutions to provide policy-based protection against the following threats against the ESXi image:
    - Corruption of the stored image
    - Certain kinds of tampering
    - Unexpected or unauthorized updates or other types of changes

Enable the dynamic launch of the VMkernel using TPM with an advanced configuration option, enableTboot, in the vSphere Client. This is referred to as Dynamic Root of Trust for Measurement (DRTM). By default, the use of DRTM for measuring VMkernel is disabled.

- The ESXi firewall in ESXi 5.0 does not allow per-network filtering of vMotion traffic. Therefore, you must install rules on your external firewall to ensure that no incoming connections can be made to the vMotion socket.
- If you have a firewall between two ESXi hosts:
  - 443 (server-to-server migration and provisioning traffic)
  - 902 (server-to-server migration and provisioning traffic)
  - 2050–2250 (for HA traffic)
  - 8000 (for vMotion)
  - 8042–8045 (for HA traffic)

Port	Purpose	Traffic type
22	SSH Server	Incoming TCP
53	DNS client	Incoming and outgoing UDP
68	DHCP Client	Incoming and outgoing UDP
80	• vSphere Fault Tolerance (FT) (outgoing TCP, UDP) • HTTP access • The default non-secure TCP Web port typically used in conjunction with port 443 as a front end for access to ESXi networks from the Web. Port 80 redirects traffic to an HTTPS landing page (port 443). • WS-Management	Incoming TCP
123	NTP Client	Outgoing UDP
161	SNMP Server	Incoming UDP
427	The CIM client uses the Service Location Protocol, version 2 (SLPv2) to find CIM servers.	Incoming and outgoing UDP
443	• HTTPS access • vCenter Server access to ESXi hosts • Default SSL Web port • vSphere Client access to vCenter Server • vSphere Client access to ESXi hosts • WS-Management • vSphere Client access to vSphere Update Manager • Third-party network management client connections to vCenter • Third-party network management clients access to hosts	Incoming TCP
902	• Host access to other hosts for migration and provisioning • Authentication traffic for ESXi and remote console traffic (xinetd/vmware-authd) • vSphere Client access to virtual machine consoles • (UDP) Status update (heartbeat) connection from ESXi to vCenter	Incoming and outgoing TCP, outgoing UDP
903	• Remote console traffic generated by user access to virtual machines on a specific host. • vSphere Client access to virtual machine consoles • MKS transactions (xinetd/vmware-authd-mks)	Incoming TCP
1234, 1235	Host Based Replication (HBR)	Outgoing TCP
2049	Transactions from NFS storage devices, the port is used on the VMkernel interface.	Incoming and outgoing TCP
2050-2250	Between ESX hosts for HA and EMC Autostart Manager	Outgoing TCP, incoming and outgoing UDP
3260	Transactions to iSCSI storage devices	Outgoing UDP
5900-5964	RFB protocol, which is used by management tools such as VNC	Incoming and outgoing TCP
5988	CIM transactions over HTTP	Incoming TCP
5989	CIM XML transactions over HTTPS	Incoming and outgoing TCP
8000	vMotion requests	Incoming and outgoing TCP
8042-8045	Between hosts for HA and EMC Autostart Manager	Outgoing TCP, incoming and outgoing UDP
8100, 8200	Between hosts for Fault Tolerance (FT)	Incoming and outgoing UDP & TCP

- Virtual switches and VLANs can protect against the following types of attacks.
  - MAC flooding — Floods a switch with packets that contain MAC addresses tagged as having come from different sources. Many switches use a content-addressable memory (CAM) table to learn and store the source address for each packet. When the table is full, the switch can enter a fully open state in which every incoming packet is broadcast on all ports, letting the attacker see all of the switch's traffic. This state might result in packet leakage across VLANs. Although VMware virtual switches store a MAC address table, they do not get the MAC addresses from observable traffic and are not vulnerable to this type of attack.
  - 802.1q and ISL tagging attacks — Force a switch to redirect frames from one VLAN to another by tricking the switch into acting as a trunk and broadcasting the traffic to other VLANs. VMware virtual switches do not perform the dynamic trunking required for this type of attack and, therefore, are not vulnerable.
  - Double-encapsulation attacks — Occur when an attacker creates a double-encapsulated packet in which the VLAN identifier in the inner tag is different from the VLAN identifier in the outer tag. For backward compatibility, native VLANs strip the outer tag from transmitted packets unless configured to do otherwise. When a native VLAN switch strips the outer tag, only the inner tag is left, and that inner tag routes the packet to a different VLAN than the one identified in the now-missing outer tag. VMware virtual switches drop any double-encapsulated frames that a virtual machine attempts to send on a port configured for a specific VLAN. Therefore, they are not vulnerable to this type of attack.
  - Multicast brute-force attacks — Involve sending large numbers of multicast frames to a known VLAN almost simultaneously to overload the switch so that it mistakenly allows some of the frames to broadcast to other VLANs. VMware virtual switches do not allow frames to leave their correct broadcast domain (VLAN) and are not vulnerable to this type of attack.
  - Spanning-tree attacks — Target Spanning-Tree Protocol (STP), which is used to control bridging between parts of the LAN. The attacker sends Bridge Protocol Data Unit (BPDU) packets that attempt to change the network topology, establishing themselves as the root bridge. As the root bridge, the attacker can sniff the contents of transmitted frames. VMware virtual switches do not support STP and are not vulnerable to this type of attack.
  - Random frame attacks — Involve sending large numbers of packets in which the source and destination addresses stay the same, but in which fields are randomly changed in length, type, or content. The goal of this attack is to force packets to be mistakenly rerouted to a different VLAN. VMware virtual switches are not vulnerable to this type of attack.
- The iSCSI initiator relies on being able to get MAC address changes from certain types of storage. If you are using ESXi iSCSI and have iSCSI storage, set the *MAC Address Changes* option to *Accept*.
- In some situations, you might have a legitimate need for more than one adapter to have the same MAC address on a network—for example, if you are using Microsoft Network Load Balancing in unicast mode. When Microsoft Network Load Balancing is used in the standard multicast mode, adapters do not share MAC addresses.
- You might have a legitimate reason to configure a virtual switch to operate in promiscuous mode—for example, if you are running network intrusion detection software or a packet sniffer.
- ESXi hosts support IPsec using IPv6. When you set up IPsec on a host, you enable authentication and encryption of incoming and outgoing packets.
- Add a security association to specify encryption parameters for associated IP traffic. You can add a security association using the `vicfg-ipsec` vSphere CLI command.
- List of all security associations available for use by security policies. The list includes both user created security associations and any security associations the VMkernel installed using Internet Key Exchange.
  - `vicfg-ipsec --server=server_name -l`
- ESXi can provide a list of all security policies on the host.
  - `vicfg-ipsec --server=server_name -L`
- ESXi does not support Kerberos, Secure Remote Protocol (SRP), or public-key authentication methods for iSCSI. Additionally, it does not support IPsec authentication and encryption.
- A primary security risk in iSCSI SANs is that an attacker might sniff transmitted storage data. Take additional measures to prevent attackers from easily seeing iSCSI data. Neither the hardware iSCSI adapter nor ESXi iSCSI initiator encrypts the data that they transmit to and from the targets, making the data more vulnerable to sniffing attacks.
- To ensure the protection of the data transmitted to and from external network connections, ESXi uses one of the strongest block ciphers available—256-bit AES block encryption. ESXi also uses 1024-bit RSA for key exchange. These encryption algorithms are the default for the following connections.
  - vSphere Client connections to vCenter Server and to ESXi through the management interface.
  - SDK connections to vCenter Server and to ESXi.
  - Management interface connections to virtual machines through the VMkernel.
  - SSH connections to ESXi through the management interface.
- SSH configuration in ESXi is enhanced to provide a high security level.
  - Version 1 SSH protocol disabled
  - Improved cipher strength — SSH supports only 256-bit and 128-bit AES ciphers for your connections.
- Security of the ESXi management interface is critical to protect against unauthorized intrusion and misuse.
- Consider the following recommendations when evaluating host security and administration.
  - Limit user access.

    - Use the vSphere Client to administer your ESXi hosts.
    - Use only VMware sources to upgrade ESXi components.
- You can add supported services and management agents that are required to operate the host by adding rule set configuration files to the ESXi firewall directory /etc/vmware/firewall/. The default rule set configuration file is service.xml. The file contains firewall rules and describes each rule's relationship with ports and protocols.
- The behavior of the NFS Client rule set (nfsClient) is different from other rule sets. When the NFS Client rule set is enabled, all outbound TCP ports are open for the destination hosts in the list of allowed IP addresses.
- To add a service to the host security profile, you define the port rules for the service in a configuration file. Name the configuration file service_name.xml.
- After you add or edit a firewall rule set, you must refresh the firewall to load the new rule.
    - `esxcli --server=server_name network firewall refresh`
- ESXi automatically configures NFS Client settings when you mount or unmount an NFS datastore. When you add or mount an NFS datastore, ESXi checks the state of the NFS Client (nfsClient) firewall rule set.
    - If the NFS Client rule set is disabled, ESXi enables the rule set and disables the Allow All IP Addresses policy by setting the *allowedAll* flag to FALSE. The IP address of the NFS server is added to the allowed list of outgoing IP addresses.
    - If the NFS Client rule set is enabled, the state of the rule set and the allowed IP address policy are not changed. The IP address of the NFS server is added to the allowed list of outgoing IP addresses.
- If you manually enable the NFS Client rule set or manually set the Allow All IP Addresses policy, either before or after adding an NFS datastore to the system, your settings are overridden when the last NFS datastore is unmounted. The NFS Client rule set is automatically disabled when all NFS datastores are unmounted.
- ESXi can automate whether services start based on the status of firewall ports.
- `esxcli network firewall get` Returns the enabled or disabled status of the firewall and lists default actions.
- `esxcli network firewall ruleset list` List rule sets information.
- The PAM configuration for VMware services is located in /etc/pam.d/vmware-authd, which stores paths to authentication modules.
- The host checks for password compliance using the default authentication plug-in, pam_passwdqc.so.
- Your user password must meet the following length requirements:
    - Passwords containing characters from one or two character classes must be at least eight characters long.
    - Passwords containing characters from three character classes must be at least seven characters long.
    - Passwords containing characters from all four character classes must be at least six characters long.
- To be granted shell access, users must also have an Administrator role for an inventory object on the host.
- Users who are logged in and are removed from the domain keep their host permissions until you restart the host.
- An uppercase character that begins a password does not count toward the number of character classes used. A number that ends a password does not count toward the number of character classes used.
- For ESXi and vCenter Server, permissions are defined as access roles that consist of a user and the user's assigned role for an object such as a virtual machine or ESXi host. 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 a user assigned to the Administrator role on all objects. When connecting directly to the host, the root and vpxuser user accounts have the same access rights as any user assigned the Administrator role on all objects.
- All other users initially have no permissions on any objects.
- Moving an object in the inventory hierarchy requires appropriate privileges on the object itself, the source parent object (such as a folder or cluster), and the destination parent object.
- When you assign a permission to an object, you can choose whether the permission propagates down the object hierarchy. You set propagation for each permission. Propagation is not universally applied. Permissions defined for a child object always override the permissions that are propagated from parent objects.
- Most inventory objects inherit permissions from a single parent object in the hierarchy.
- Virtual machines inherit permissions from both the parent virtual machine folder and the parent host, cluster, or resource pool simultaneously.
- To set permissions for a distributed switch and its associated distributed port groups, set permissions on a parent object, such a folder or datacenter.
- Objects might have multiple permissions, but only one permission for each user or group.
- Permissions applied on a child object always override permissions that are applied on a parent object. Virtual machine folders and resource pools are equivalent levels in the hierarchy.
- If multiple group permissions are defined on the same object and the user belongs to two or more of those groups, two situations are possible:
    - If no permission is defined for the user on that object, the user is assigned the set 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.
- `vicfg` commands do not perform an access check. Therefore, even if you limit the root user's privileges, it does not affect what that user can do using the command-line interface commands.

- Use the *No Access* role to masks specific areas of the hierarchy that you don't want particular users to have access to.
- vCenter Server and ESXi grant access to objects only to users who are assigned permissions for the object. When you assign a user or group permissions for the object, you do so by pairing the user or group with a role.
- A role is a predefined set of privileges.
- ESXi hosts provide three default roles, and you cannot change the privileges associated with these roles.
- Privileges define individual rights that a user requires 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.
- 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.
- 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.
- Changes to permissions and roles take effect immediately, even if the users involved are logged in. The exception is searches, where permission changes take effect after the user has logged out and logged back in.
- When you use the vSphere Authentication Proxy, you do not need to store Active Directory credentials on the host. Users supply the domain name of the Active Directory server and the IP address of the authentication proxy server when they join a host to a domain.
- When you install the vSphere Authentication Proxy service, the installer creates a domain account with appropriate privileges to run the authentication proxy service. The account name begins with the prefix CAMand has a 32-character, randomly generated password associated with it.
- Before you use the vSphere Authentication Proxy to connect ESXi to a domain, you must authenticate the vSphere Authentication Proxy server to ESXi. If you use Host Profiles to connect a domain with the vSphere Authentication Proxy server, you do not need to authenticate the server.
- To authenticate ESXi to use the vSphere Authentication Proxy, export the server certificate from the vSphere Authentication Proxy system and import it to ESXi. You need only authenticate the server once.
- ESXi supports synchronizing time with an external NTPv3 or NTPv4 server that is compliant with RFC 5905 and RFC 1305. The Microsoft Windows W32Time service does not meet these requirements.
- Host certificate checking is enabled by default and SSL certificates are used to encrypt network traffic.
- ESXi supports only X.509 certificates.
- When you enable lockdown mode, no users other than vpxuser have authentication permissions, nor can they perform operations against the host directly. Lockdown mode forces all operations to be performed through vCenter Server.
- When a host is in lockdown mode, you cannot run vSphere CLI commands from an administration server, from a script, or from vMA against the host. External software or management tools might not be able to retrieve or modify information from the ESXi host.
- The root user is still authorized to log in to the direct console user interface when lockdown mode is enabled.
- Lockdown mode is only available on ESXi hosts that have been added to vCenter Server.
- Lockdown Mode Behavior:

Service	Normal Mode	Lockdown Mode
vSphere WebServices API	All users, based on ESXi permissions	vCenter only (vpxuser)
CIM Providers	Root users and users with Admin role on the host	vCenter only (ticket)
Direct Console UI (DCUI)	Root users and users with Admin role on the host	Root users
ESXi Shell	Root users and users with Admin role on the host	No users
SSH	Root users and users with Admin role on the host	No users

- You can enable or disable remote and local access to the ESXi Shell to create different lockdown mode configurations.
- If you lose access to vCenter Server while running in Total Lockdown Mode, you must reinstall ESXi to gain access to the host.
- By default, ESXi imposes no restrictions on the root password.

vSphere Resource Management (ESXi & vCenter 5.0)

- Resources include CPU, memory, power, storage, and network resources.
- Specifying shares makes sense only with regard to sibling virtual machines or resource pools.
- Share values:

Setting	CPU share values	Memory share values
High	2000 shares per virtual CPU	20 shares per megabyte of configured virtual machine memory
Normal	1000 shares per virtual CPU	10 shares per megabyte of configured virtual machine memory
Low	500 shares per virtual CPU	5 shares per megabyte of configured virtual machine memory

- A reservation specifies the guaranteed minimum allocation
- A limit specifies an upper bound.
- A server can allocate more than the reservation to a virtual machine, but never allocates more than the limit.
- If there are frequent changes to the total available resources, use Shares.

- Hyperthreading options:
  - *Any* — The default can freely share cores with other virtual CPUs from this or any other virtual machine at any time.
  - *None* — Should not share cores with each other or with virtual CPUs from other virtual machines.
  - *Internal* — Virtual CPUs from this virtual machine cannot share cores with virtual CPUs from other virtual machines. Only for SMP virtual machines
- CPU power management policies:
  - *Not supported* — The host does not support any power management features or power management is not enabled in the BIOS.
  - *High Performance* — The VMkernel detected certain power management features, but will not use them unless the BIOS requests them for power capping or thermal events.
  - *Balanced* (Default) — The VMkernel uses the available power management features conservatively to reduce host energy consumption with minimal compromise to performance.
  - *Low Power* — The VMkernel aggressively uses available power management features to reduce host energy consumption at the risk of lower performance.
  - *Custom* — The VMkernel bases its power management policy on the values of several advanced configuration parameters. You can set these parameters in the vSphere Client Advanced Settings dialog box.
- A host determines allocations for each virtual machine based on the number of shares allocated to it and an estimate of its recent working set size.
- Working set size — ESXi hosts estimate the working set for a virtual machine by monitoring memory activity over successive periods of virtual machine execution time. Estimates are smoothed over several time periods using techniques that respond rapidly to increases in working set size and more slowly to decreases in working set size. This approach ensures that a virtual machine from which idle memory is reclaimed can ramp up quickly to its full share-based allocation when it starts using its memory more actively. Memory activity is monitored to estimate the working set sizes for a default period of 60 seconds.
- Virtual machine executable (VMX) swap files allow the host to greatly reduce the amount of overhead memory reserved for the VMX process.
  - VMX swap files are not related to the swap to host cache feature or to regular host-level swap files.
  - Memory for certain components, such as the virtual machine monitor (VMM) and virtual devices, is fully reserved when a virtual machine is powered on. However, some of the overhead memory that is reserved for the VMX process can be swapped. The VMX swap feature reduces the VMX memory reservation significantly (for example, from about 50MB or more per virtual machine to about 10MB per virtual machine). This allows the remaining memory to be swapped out when host memory is overcommitted, reducing overhead memory reservation for each virtual machine.
  - The host creates VMX swap files automatically, provided there is sufficient free disk space at the time a virtual machine is powered on.
- If a virtual machine is not actively using all of its currently allocated memory, 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.
- The idle memory tax is applied in a progressive fashion. The effective tax rate increases as the ratio of idle memory to active memory for the virtual machine rises.
- Guest-level swap space must be greater than or equal to the difference between the virtual machine's configured memory size and its Reservation.
- Datastores that are created on solid state drives (SSD) can be used to allocate space for host cache. The host reserves a certain amount of space for swapping to SSD.
- Only SSD-backed datastores appear in the list of datastores on the Host Cache Configuration page.
- You can change the percentage of space allocated or disable the host's ability to swap to SSD.
- To determine the effectiveness of memory sharing for a given workload, try running the workload, and use `resxtop` or `esxtop` to observe the actual savings. Find the information in the *PSHARE* field of the interactive mode in the Memory page.
- Pages that can be compressed to 2 KB or smaller are stored in the virtual machine's compression cache, increasing the capacity of the host.
- If you do not set the size of the compression cache, ESXi uses the default value of 10 percent.
- Memory Granted is the amount of guest physical memory that is mapped to machine memory.
- Memory Consumed is the amount of machine memory allocated to the virtual machine, accounting for savings from shared memory.
- A similar result is obtained when determining Memory Shared and Memory Shared Common for the host.
- Memory reliability, also known as error insolation, allows ESXi to stop using parts of memory when it determines that a failure might occur, as well as when a failure did occur. When enough corrected errors are reported at a particular address, ESXi stops using this address to prevent the corrected error from becoming an uncorrected error.
- Storage I/O Control extends the constructs of shares and limits to handle storage I/O resources. By default, all virtual machine shares are set to Normal (1000) with unlimited IOPS.
- Storage I/O Control is enabled by default on Storage DRS-enabled datastore clusters.
- SIOC has several requirements and limitations:
  - Datastores that are Storage I/O Control-enabled must be managed by a single vCenter Server system.
  - Storage I/O Control is supported on Fibre Channel-connected, iSCSI-connected, and NFS-connected storage. Raw Device Mapping (RDM) is not supported.
  - Storage I/O Control does not support datastores with multiple extents.

    - Verify whether your automated tiered storage array has been certified to be compatible with Storage I/O Control.
- If a virtual machine has more than one virtual disk, you must set the limit on all of its virtual disks. Otherwise, the limit will not be enforced for the virtual machine.
- The congestion threshold value for a datastore is the upper limit of latency that is allowed for a datastore before Storage I/O Control begins to assign importance to the virtual machine workloads according to their shares.
- Storage I/O Control will not function correctly unless all datatores that share the same spindles on the array have the same congestion threshold.
- Set the values based on:
    - A higher value typically results in higher aggregate throughput and weaker isolation.
    - If throughput is more critical than latency, do not set the value too low.
    - A lower value will result in lower device latency and stronger virtual machine I/O performance isolation.
    - The value must be between 10 ms and 100 ms.
    - The default value is 30 ms.
- When you power on a virtual machine in a resource pool, or try to create a child resource pool, the system performs additional admission control to ensure the resource pool's restrictions are not violated.
- How available CPU and memory resources are computed and whether actions are performed depends on the Reservation Type.
- Reservation types:
    - *Fixed* — The system checks whether the selected resource pool has sufficient unreserved resources. If it does, the action can be performed. If it does not, a message appears and the action cannot be performed.
    - *Expandable* (default) — The system considers the resources available in the selected resource pool and its direct parent resource pool. If the parent resource pool also has the *Expandable Reservation* option selected, it can borrow resources from its parent resource pool. Borrowing resources occurs recursively from the ancestors of the current resource pool as long as the *Expandable Reservation* option is selected. Leaving this option selected offers more flexibility, but, at the same time provides less protection. A child resource pool owner might reserve more resources than you anticipate.
- Expandable reservations cause a loss of strict isolation.
- The DRS migration threshold is a measure of how much cluster imbalance across host (CPU and memory) loads is acceptable.
- The Conservative setting generates only priority-one recommendations (mandatory recommendations), the Aggressive level which generates priority-five recommendations and higher (that is, all recommendations.)
- A priority level for each migration recommendation is computed using the load imbalance metric of the cluster. This metric is displayed as Current host load standard deviation. A higher load imbalance leads to higher-priority migration recommendations.
- When DRS is disabled, the cluster's resource pool hierarchy and affinity rules are not reestablished when DRS is turned back on. So if you disable DRS, the resource pools are removed from the cluster. To avoid losing the resource pools, instead of disabling DRS, you should suspend it by changing the DRS automation level to manual (and disabling any virtual machine overrides).
- A DRS cluster can be valid, overcommitted (yellow), or invalid (red).
    - Overcommitted if a host fails.
    - Invalid if vCenter Server is unavailable and you power on virtual machines using a vSphere Client connected directly to a host.
    - Invalid if the user reduces the reservation on a parent resource pool while a virtual machine is in the process of failing over.
    - If changes are made to hosts or virtual machines using a vSphere Client connected to a host while vCenter Server is unavailable. When vCenter Server becomes available again, you might find that clusters have turned red or yellow because cluster requirements are no longer met.
- If two VM-VM affinity rules are in conflict, you cannot enable both. For example, if one rule keeps two virtual machines together and another rule keeps the same two virtual machines apart, you cannot enable both rules.
- When two VM-VM affinity rules conflict, the older one takes precedence and the newer rule is disabled. DRS only tries to satisfy enabled rules and disabled rules are ignored. DRS gives higher precedence to preventing violations of anti-affinity rules than violations of affinity rules.
- A datastore cluster is a collection of datastores with shared resources and a shared management interface.
- Resource management capabilities are:
    - Space utilization load balancing
    - I/O latency load balancing
    - Anti-affinity rules
- Initial placement recommendations are made in accordance with space constraints and with respect to the goals of space and I/O load balancing.
- Possible reasons for Storage DRS recommendations are:
    - Balance datastore space use
    - Balance datastore I/O load
- Storage DRS makes mandatory recommendations for migration in the following situations:
    - The datastore is out of space.
    - Anti-affinity or affinity rules are being violated.
    - The datastore is entering maintenance mode and must be evacuated.
- Guidelines when you create a datastore cluster:

    - Datastore clusters must contain similar or interchangeable datastores. A datastore cluster can contain a mix of datastores with different sizes and I/O capacities, and can be from different arrays and vendors. However, the following types of datastores cannot coexist in a datastore cluster.
        - NFS and VMFS datastores cannot be combined in the same datastore cluster.
        - Replicated datastores cannot be combined with non-replicated datastores in the same Storage-DRS enabled datastore cluster.
    - All hosts attached to the datastores in a datastore cluster must be ESXi 5.0 and later. If datastores in the datastore cluster are connected to ESX/ESXi 4.x and earlier hosts, Storage DRS does not run.
    - Datastores shared across multiple datacenters cannot be included in a datastore cluster.
    - As a best practice, do not include datastores that have hardware acceleration enabled in the same datastore cluster as datastores that do not have hardware acceleration enabled. Datastores in a datastore cluster must be homogeneous to guarantee hardware acceleration-supported behavior.
- When you enable Storage DRS, you enable the following functions:
    - Space load balancing among datastores within a datastore cluster.
    - I/O load balancing among datastores within a datastore cluster.
    - Initial placement for virtual disks based on space and I/O workload.
- The Enable Storage DRS check box in the Datastore Cluster Settings dialog box enables or disables all of these components at once. If necessary, you can disable I/O-related functions of Storage DRS independently of space balancing functions.
- When you disable Storage DRS on a datastore cluster, Storage DRS settings are preserved. When you enable Storage DRS, the settings for the datastore cluster are restored to the point where Storage DRS was disabled.
- Storage DRS migration recommendations appear on the *Storage DRS* tab in the vSphere Client. You can refresh these recommendations by running Storage DRS.
- SDRS Automation Levels:
    - Default (Manual)
    - Fully Automated
    - Disabled
- Thresholds to set the aggressiveness level for Storage DRS:
    - *Space Utilization* — Storage DRS generates recommendations or performs migrations when the percentatge of space used on the datastore is greater than the threshold you set in the vSphere Client.
    - *I/O Latency* — Storage DRS generates recommendations or performs migrations when the 90th percentile I/O latency measured over a day for the datastore is greater than the threshold.
- You can also set advanced options:
    - *Space utilization difference* — This threshold ensures that there is some minimum difference between the space utilization of the source and the destination. For example, if the space used on datastore A is 82% and datastore B is 79%, the difference is 3. If the threshold is 5, Storage DRS will not make migration recommendations from datastore A to datastore B.
    - *I/O load balancing invocation intercal* — After this interval, Storage DRS runs to balance I/O load.
    - ***I****/O imbalance threshold* — Lowering this value makes I/O load balancing less aggressive. Storage DRS computes an I/O fairness metric between 0 and 1, which 1 being the fairest distribution. I/O load balancing runs only if the computed metric is less than 1 — (I/O imbalance threshold / 100).
- You can create a scheduled task to change Storage DRS settings for a datastore cluster so that migrations for fully automated datastore clusters are more likely to occur during off-peak hours.
- You can create a scheduled task to change the automation level and aggressiveness level for a datastore cluster. For example, you might configure Storage DRS to run less aggressively during peak hours, when performance is a priority, to minimize the occurrence of storage migrations. During non-peak hours, Storage DRS can run in a more aggressive mode and be invoked more frequently.
- By default, a virtual machine's virtual disks are kept together on the same datastore.
- When you create an anti-affinity rule, it applies to the relevant virtual disks in the datastore cluster. Anti-affinity rules are enforced during initial placement and Storage DRS-recommendation migrations, but are not enforced when a migration is initiated by a user.
- Anti-affinity rules do not apply to CD-ROM ISO image files that are stored on a datastore in a datastore cluster, nor do they apply to swapfiles that are stored in user-defined locations.
- *Inter-VM Anti-Affinity Rules* — Specify which virtual machines should never be kept on the same datastore.
- *Intra-VM Anti-Affinity Rules* — Specify which virtual disks associated with a particular virtual machine must be kept on different datastores.
- You can create an anti-affinity rule to indicate that all virtual disks of certain virtual machines must be kept on different datastores. The rule applies to individual datastore clusters.
- Virtual machines that participate in an inter-VM anti-affinity rule in a datastore cluster must be associated with an intra-VM affinity rule in the datastore cluster. The virtual machines must also comply with the intra-VM affinity rule.
- VMDK affinity rules indicate that all virtual disks in a datastore cluster that are associated with a particular virtual machine are located on the same datastore in the datastore cluster. The rules apply to individual datastore clusters.
- VMDK affinity rules are enabled by default for all virtual machines that are in a datastore cluster. You can override the default setting for the datastore cluster or for individual virtual machines.

- When you enable the option to clear Storage DRS statistics, statistics are cleared every time Storage DRS runs until you disable the option. Always disable the option after you diagnose the Storage DRS problem.
- Storage DRS and SRM are not aware of each other.
- SRM with Array replication:
  - On the protection site, datastore clusters should have only one consistency group.
  - Use manual mode for load balancing recommendations.
  - If load balancing is manually invoked, or if datastore maintenance mode is invoked, SRM might not be able to recover the affected virtual machine. However, you can manually recover the virtual machine.
- SRM with Host Based replication:
  - Disable I/O load balancing for datastore clusters. After a Storage vMotion migration, Host Based Replication (HBR) might generate a large amount of storage I/O activity that can affect Storage DRS.
  - HBR guarantees crash consistency among a virtual machine's disks. Using VSS might allow a higher level of consistency.
  - Storage vMotion might impact recovery point objectives (RPO) if the virtual machine home directory is moved. While RPO might be impacted, you can recover SRM-protected virtual machines on the recovery site.
- ESXi uses a sophisticated NUMA scheduler to dynamically balance processor load and memory locality or processor load balance.
  a. Each virtual machine managed by the NUMA scheduler is assigned a home node. A home node is one of the system's NUMA nodes containing processors and local memory, as indicated by the System Resource Allocation Table (SRAT).
  b. When memory is allocated to a virtual machine, the ESXi host preferentially allocates it from the home node. The virtual CPUs of the virtual machine are constrained to run on the home node to maximize memory locality.
  c. The NUMA scheduler can dynamically change a virtual machine's home node to respond to changes in system load. The scheduler might migrate a virtual machine to a new home node to reduce processor load imbalance. Because this might cause more of its memory to be remote, the scheduler might migrate the virtual machine's memory dynamically to its new home node to improve memory locality. The NUMA scheduler might also swap virtual machines between nodes when this improves overall memory locality.

vSphere Availability (ESXi and vCenter 5.0)

- When you create a vSphere HA cluster, you choose a single host as the master host to communicate with vCenter Server and to monitor the state of the other, slave, hosts and their virtual machines.
- The master host must distinguish between a failed host and one that is in a network partition or that has become network isolated. The master host uses datastore heartbeating to determine the type of failure.
- All active hosts (those not in standby or maintenance mode, or not disconnected) participate in an election to choose the cluster's master host. The host that mounts the greatest number of datastores has an advantage in the election. Only one master host exists per cluster and all other hosts are slave hosts. If the master host fails, is shut down, or is removed from the cluster a new election is held.
- Master host responsibilities:
  - Monitoring the state of slave hosts. The master host identifies which virtual machines need to be restarted.
  - Monitoring the power state of all protected virtual machines. If one virtual machine fails, the master host ensures that it is restarted. Using a local placement engine, the master host also determines where the restart should be done.
  - Managing the lists of cluster hosts and protected virtual machines.
  - Acting as vCenter Server management interface to the cluster and reporting the cluster health state.
- The slave hosts monitoring their runtime states, and reporting state updates to the master host.
- Both slave hosts and master hosts implement the VM and Application Monitoring features.
- HA guarantees that it attempts to power it back on after a failure. A master host commits to protecting a virtual machine when it observes that the power state of the virtual machine changes from powered off to powered on in response to a user action.
- The master host that has exclusively locked a system-defined file on the datastore that contains a virtual machine's configuration file.
- This state is reported on the host's *Summary* tab in the vSphere Client and in the Host List view for a cluster or datacenter, if the HA State column has been enabled. An HA state of "Running (Master)" indicates the host is serving as a vSphere HA master host. A state of "Connected (Slave)" indicates the host is serving as a vSphere HA slave host. Several other states are provided to indicate when an election is underway or an error condition has occurred.
- If you disconnect a host from a cluster, all of the virtual machines registered to that host are unprotected by vSphere HA.
- Three types of host failure are detected:
  - A host stops functioning (fails).
  - A host becomes network isolated.
  - A host loses network connectivity with the master host.
- Network heartbeats every second. When the master host stops receiving these heartbeats from a slave host, it checks for host liveness before declaring the host to have failed. The liveness check that the master host performs is to determine whether the slave host is exchanging heartbeats with one of the datastores. Also, the master host checks whether the host responds to ICMP pings sent to its management IP addresses.
- If a master host is unable to communicate directly with the agent on a slave host, the slave host does not respond to ICMP pings, and the agent is not issuing heartbeats it is considered to have failed.

- If such a slave host is exchanging heartbeats with a datastore, the master host assumes that it is in a network partition or network isolated and so continues to monitor the host and its virtual machines.
- Host network isolation occurs when a host is still running, but it can no longer observe traffic from vSphere HA agents on the management network. If a host stops observing this traffic, it attempts to ping the cluster isolation addresses. If this also fails, the host declares itself as isolated from the network.
- The master host monitors the virtual machines that are running on an isolated host and if it observes that they power off, and the master host is responsible for the virtual machines, it restarts them.
- A partitioned cluster leads to degraded virtual machine protection and cluster management functionality.
- Virtual machine protection. vCenter Server allows a virtual machine to be powered on, but it is protected only if it is running in the same partition as the master host that is responsible for it. The master host must be communicating with vCenter Server. A master host is responsible for a virtual machine if it has exclusively locked a system-defined file on the datastore that contains the virtual machine's configuration file.
- If vCenter Server can communicate with only some of the hosts in the cluster, and it can connect to only one master host, changes in configuration that affect vSphere HA might not take effect until after the partition is resolved. This failure could result in one of the partitions operating under the old configuration, while another uses the new settings.
- If a vSphere HA cluster contains pre-ESXi 5.0 hosts and a partition occurs, vSphere HA might incorrectly power on a virtual machine that was powered off by the user or it might fail to restart a virtual machine that failed.
- When a vSphere HA cluster is partitioned, you can add a host only to the partition that vCenter Server is communicating with.
- vCenter Server selects a preferred set of datastores for heartbeating. To maximize the number of hosts that have access to a heartbeating datastore and minimize the likelihood that the datastores are backed by the same storage array or NFS server.
- The *Datastore Heartbeating* tab lets you specify alternative datastores. Only datastores mounted by at least two hosts are available.
- You can use the advanced attribute *das.heartbeatdsperhost* to change the number of heartbeat datastores selected by vCenter Server for each host. The default is two and the maximum valid value is five.
- If you deploy a converged network environment, where storage and management network traffic travel over the same physical NICs, disable datastore heartbeating. It does not provide any benefit in this type of network.
- vSphere HA creates a directory at the root of each datastore that is used for both datastore heartbeating and for persisting the set of protected virtual machines. The name of the directory is .vSphere-HA.
- Because more than one cluster might use a datastore, subdirectories for this directory are created for each cluster. Root owns these directories and files and only root can read and write to them.
- With vmfs3, the maximum usage is approximately 2GB and the typical usage (*Forbes*: minimum usage) is approximately 3MB. With vmfs5 the maximum and typical usage is approximately 3MB. vSphere HA use of the datastores adds negligible overhead and has no performance impact on other datastore operations.
- vSphere HA uses TCP and UDP port 8182 for agent-to-agent communication.
- For ESXi 5.x hosts, vSphere HA writes to syslog only by default.
- vSphere HA logs onto the vSphere HA agents using a user account, vpxuser.
- All communication between vCenter Server and the vSphere HA agent is done over SSL. Agent-to-agent communication also uses SSL except for election messages, which occur over UDP. Election messages are verified over SSL.
- Each host generates a self-signed certificate when it is booted for the first time.
- If the certificate is replaced, vSphere HA needs to be reconfigured on the host.
- HA might not be able to fail over virtual machines because of resource constraints.
    - HA admission control is disabled and Distributed Power Management (DPM) is enabled. This can result in DPM consolidating virtual machines onto fewer hosts and placing the empty hosts in standby mode leaving insufficient powered-on capacity to perform a failover.
    - VM-Host affinity (required) rules might limit the hosts on which certain virtual machines can be placed.
    - There might be sufficient aggregate resources but these can be fragmented across multiple hosts so that they can not be used by virtual machines for failover.
- Slot size is comprised of two components, CPU and memory.
    - vSphere HA calculates the CPU component by obtaining the CPU reservation of each powered-on virtual machine and selecting the largest value. If you have not specified a CPU reservation for a virtual machine, it is assigned a default value of 32 MHz. You can change this value by using the *das.vmcpuminmhz* advanced attribute.)
    - vSphere HA calculates the memory component by obtaining the memory reservation, plus memory overhead, of each powered-on virtual machine and selecting the largest value. There is no default value for the memory reservation.
- 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.
- New Advanced Runtime Info — *Failover slots* — The total number of slots not counting the used slots or the available slots.
- vSphere HA uses the actual reservations of the virtual machines. If a virtual machine does not have reservations, meaning that the reservation is 0, a default of 0MB memory and 256MHz (*Forbes*: not 32?) CPU is applied.

- The total host resources available for virtual machines is calculated by adding the host's CPU and memory resources. These amounts are those contained in the host's root resource pool, not the total physical resources of the host.
- With the Specify Failover Hosts admission control policy, when a host fails, vSphere HA attempts to restart its virtual machines on one of the specified failover hosts. If this is not possible, for example the failover hosts have failed or have insufficient resources, then vSphere HA attempts to restart those virtual machines on other hosts in the cluster.
- To ensure that spare capacity is available on a failover host, you are prevented from powering on virtual machines or using vMotion to migrate virtual machines to a failover host. Also, DRS does not use a failover host for load balancing.
- DRS does not load balance failover hosts and VM-VM affinity rules are not supported.
- vSphere HA includes the resource usage of Fault Tolerance Secondary VMs when it performs admission control calculations. For the Host Failures Cluster Tolerates policy, a Secondary VM is assigned a slot, and for the Percentage of Cluster Resources policy, the Secondary VM's resource usage is accounted for when computing the usable capacity of the cluster.
- There should be at least one management network in common among all hosts and best practice is to have at least two.
- Host isolation response — To use the Shut down VM setting, you must install VMware Tools in the guest operating system of the virtual machine. Shutting down the virtual machine provides the advantage of preserving its state. Shutting down is better than powering off the virtual machine, which does not flush most recent changes to disk or commit transactions. Virtual machines that are in the process of shutting down will take longer to fail over while the shutdown completes. Virtual Machines that have not shut down in 300 seconds, or the time specified in the advanced attribute das.isolationshutdowntimeout seconds, are powered off.

- VM and Application monitoring — the default settings for monitoring sensitivity:

Setting	Failure Interval	Reset Period
High	30	1 Hr
Medium	60	24 Hrs
Low	120	7 Days

- In clusters where Storage vMotion is used extensively or where Storage DRS is enabled, VMware recommends that you do not deploy vSphere HA.
- It is considered a Best Practice to use the Select the Percentage of Cluster Resources Reserved admission control policy. This policy offers the most flexibility in terms of host and virtual machine sizing. In most cases, a calculation of 1/N, where N is the number of total nodes in the cluster, yields adequate sparing.
- On ESXi hosts in the cluster, vSphere HA communications, by default, travel over VMkernel networks, except those marked for use with vMotion. If there is only one VMkernel network, vSphere HA shares it with vMotion, if necessary.
- The anti-affinity check is performed when the Primary VM is powered on. It is possible that the Primary and Secondary VMs can be on the same host when they are both in a powered-off state. This is normal behavior and when the Primary VM is powered on, the Secondary VM is started on a different host at that time.
- You can use vSphere Fault Tolerance with vSphere Distributed Resource Scheduler (DRS) when the Enhanced vMotion Compatibility (EVC) feature is enabled.
- A VM-VM affinity rule applies to the Primary VM only, while a VM-Host affinity rule applies to both the Primary VM and its Secondary VM. If a VM-VM affinity rule is set for a Primary VM, DRS attempts to correct any violations that occur after a failover (that is, after the Primary VM effectively moves to a new host).
- Use the VMware SiteSurvey utility to better understand the configuration issues associated with the cluster, host, and virtual machines being used for vSphere FT.
- The failover of fault tolerant virtual machines is independent of vCenter Server, but you must use vCenter Server to set up your Fault Tolerance clusters.
- Cluster requirements before you use Fault Tolerance:
    - Host certificate checking enabled.
    - At least two FT-certified hosts running the same Fault Tolerance version or host build number.
    - ESXi hosts have access to the same virtual machine datastores and networks.
    - Fault Tolerance logging and VMotion networking configured.
    - vSphere HA cluster created and enabled.
- Host requirements for FT:
    - Hosts must have processors from the FT-compatible processor group.
    - Hosts must be licensed for Fault Tolerance.
    - Hosts must be certified for Fault Tolerance.
    - The configuration for each host must have Hardware Virtualization (HV) enabled in the BIOS.
- VM requirements for FT:
    - No unsupported devices attached to the virtual machine.
    - Virtual machines must be stored in virtual RDM or virtual machine disk (VMDK) files that are thick provisioned.
    - Shared storage

    - Only virtual machines with a single vCPU are compatible with Fault Tolerance.
    - Supported guest operating systems.
- Not supported for fault tolerant virtual machines.
    - Snapshots.
    - Storage vMotion
    - Linked clones.
    - Virtual Machine Backups.
- The vMotion and FT logging NICs must be on different subnets and IPv6 is not supported on the FT logging NIC.
- If you configure networking to support FT but subsequently disable the Fault Tolerance logging port, pairs of fault tolerant virtual machines that are already powered on remain powered on. However, if a failover situation occurs, when the Primary VM is replaced by its Secondary VM a new Secondary VM is not started, causing the new Primary VM to run in a Not Protected state.
- The option to turn on Fault Tolerance is unavailable (dimmed):
    - The virtual machine resides on a host that does not have a license for the feature.
    - The virtual machine resides on a host that is in maintenance mode or standby mode.
    - The virtual machine is disconnected or orphaned (its .vmx file cannot be accessed).
    - The user does not have permission to turn the feature on.
- Several validation checks are performed on a virtual machine before Fault Tolerance can be turned on.
    - SSL certificate checking must be enabled in the vCenter Server settings.
    - The host must be in a vSphere HA cluster or a mixed vSphere HA and DRS cluster.
    - ESX/ESXi 4.0 or greater installed.
    - The virtual machine must not have multiple vCPUs.
    - The virtual machine must not have snapshots.
    - The virtual machine must not be a template.
    - The virtual machine must not have vSphere HA disabled.
    - The virtual machine must not have a video device with 3D enabled.
- Several additional validation checks are performed for powered-on virtual machines (or those that are in the process of being powered on).
    - Hardware Virtualization (HV) enabled.
    - The host that supports the Primary VM must have a processor that supports Fault Tolerance.
    - The host that supports the Secondary VM must have a processor that supports Fault Tolerance and is the same CPU family or model as the host that supports the Primary VM.
    - Your hardware should be certified as compatible with Fault Tolerance.
    - The combination of the virtual machine's guest operating system and processor must be supported by Fault Tolerance
- You cannot disable Fault Tolerance from the Secondary VM.
- Fault Tolerance status:
    - Protected — Indicates that the Primary and Secondary VMs are powered on and running as expected.
    - Not Protected
    - Starting
    - Need Secondary VM — This generally occurs when there is no compatible host in the cluster
    - Disabled
    - VM not Running
- vLockstep Interval — Typically, this interval is less than one-half of one second.
- Hosts running the Primary and Secondary VMs should operate at approximately the same processor frequencies, otherwise the Secondary VM might be restarted more frequently. Platform power management features that do not adjust based on workload can cause processor frequencies to vary greatly.
- Apply the same instruction set extension configuration to all hosts.
- vSphere Fault Tolerance can function in clusters with nonuniform hosts, but it works best in clusters with compatible nodes.
    - Processors from the same compatible processor group.
    - Common access to datastores used by the virtual machines.
    - The same virtual machine network configuration.
    - The same ESXi version.
    - The same Fault Tolerance version number.
    - The same BIOS settings for all hosts.
- Run the *Check Compliance* to identify FT incompatibilities and to correct them.
- VMware recommends:
    - No more than four fault tolerant virtual machines (primaries or secondaries) on any single host
    - Ensure that a resource pool containing fault tolerant virtual machines has excess memory above the memory size of the virtual machines.
    - Use a maximum of 16 virtual disks per fault tolerant virtual machine.
    - Minimum of three hosts in the cluster.

What's New in VMware vSphere 5.0 Availability Whitepaper

- VMware vCenter Server Heartbeat 6.4 provides an enhanced architecture that allows both the active and passive servers to be represented as unique entities within Microsoft Active Directory. This allows for the assignment of a unique IP address to both servers, making them accessible through the network at all times.
- Now the VMware vCenter Server instance can be associated with a virtual IP address. When a failover is required, this virtual IP address follows the active VMware vCenter Server instance.
- VMware vCenter Server Heartbeat now also provides availability of VMware View Composer and Microsoft SQL Server 2008 R2.

vSphere Monitoring and Performance (vSphere, vCenter Server & ESXi 5.0)

- There are several tools to help you monitor your virtual environment and to locate the source of potential issues and current problems:
  - Performance charts in the vSphere Client
  - Performance monitoring command-line utilities
  - Host health
  - Storage maps and charts
  - Events, alerts, and alarms in the vSphere Client
- Collection levels — determine the number of counters for which data is gathered during each collection interval. Collection levels are also referred to as statistics levels. Levels 1-4, 4 being the most stats.
- Collection intervals — determine the time period during which statistics are aggregated, calculated, rolled up, and archived in the vCenter Server database.

Collection frequency	Retention
5 Minutes	1 Day
30 Minutes	1 Week
2 Hours	1 Month
1 Day	1 Year

- Real-time data appears in the performance charts only for hosts and virtual machines that are powered on.
- Historical data appears for all supported inventory objects.
- Use advanced charts, or create your own custom charts, to see more performance data.
- If Virtual machine CPU usage is above 90% and the CPU ready value is above 20%, application performance is impacted.
- Memory Performance issues experienced if:
  - Memory usage is constantly high (94% or greater) or constantly low (24% or less).
  - Free memory consistently is 6% or less and swapping frequently occurs.
- Disk Performance issues if:
  - *kernelLatency* data counter is greater than 4ms.
  - *deviceLatency* data counter is greater than 15ms indicates there are probably problems with the storage array.
  - *queueLatency* data counter measures above zero.
- VMware-specific performance objects are loaded into Microsoft Windows Perfmon and enabled when VMware Tools is installed.
- The host health monitoring tool presents data gathered using Systems Management Architecture for Server Hardware (SMASH) profiles. The information displayed depends on the sensors available on your server hardware.
- Some host hardware sensors display data that is cumulative over time. You can reset these sensors to clear the data in them and begin collecting new data.
- Maps are automatically updated every 30 minutes.
- You can move individual items on the storage map to make the map visually clearer.
- You can export maps to various graphic file types, including jpeg, tiff, and gif.
- Events are records of user actions or system actions that occur on objects in vCenter Server or on a host.
- Three types of events:
  - Information
  - Warning
  - Error
- Alarms are notifications that are activated in response to an event, a set of conditions, or the state of an inventory object.
  - *Triggers* — Defines the event, condition, or state that will trigger the alarm and defines the notification severity.
  - *Tolerance thresholds (Reporting)* — Provides additional restrictions on condition and state triggers thresholds that must be exceeded before the alarm is triggered.
  - *Actions* — Defines operations that occur in response to triggered alarms.
- Alarms have the following severity levels:
  - *Normal* – green
  - *Warning* – yellow

    - *Alert* – red
- You can enable, disable, and modify alarms only from the object in which the alarm is defined. For example, if you defined an alarm in a cluster to monitor virtual machines, you can only enable, disable, or modify that alarm through the cluster; you can not make changes to the alarm at the individual virtual machine level.
- You cannot save an alarm without triggers defined for it.
- Reporting settings include Range and Frequency. Range is the threshold the monitored condition or state must exceed the specified trigger limit for the alarm to trigger. Frequency is the length of time between each retriggering for as long as the condition or state exists.
- Disabling alarm actions is not the same as disabling an alarm, nor is it the same as acknowledging an alarm.
- Acknowledging an alarm lets other users know that you are taking ownership of the issue. After an alarm is acknowledged, its alarm actions are discontinued. The alarm, however, is still visible in the system. Alarms are neither cleared, nor reset when acknowledged.
- The `resxtop` and `esxtop` command-line utilities provide a detailed look at how ESXi uses resources in real time.
- You can use `resxtop` remotely, whereas you can start `esxtop` only through the ESXi Shell of a local ESXi host. You must have root user privileges.
- The `esxtop` utility reads its default configuration from .esxtop41rc (*Forbes*: I suspect this is a typo and should be .esxtop50rc).This configuration file consists of nine lines.
- The first eight lines contain lowercase and uppercase letters to specify which fields appear in which order on the CPU, memory, storage adapter, storage device, virtual machine storage, network, interrupt, and CPU power panels. The letters correspond to the letters in the Fields or Order panels for the respective `esxtop` panel.
- The ninth line contains information on the other options. Most important, if you saved a configuration in secure mode, you do not get an insecure `esxtop` without removing the s from the seventh line of your .esxtop50rc file. A number specifies the delay time between updates. As in interactive mode, typing c, m, d, u, v, n, I, or p determines the panel with which `esxtop` starts.
- Do not edit the .esxtop41rc file (*Forbes*: again, this should probably be .esxtop50rc). Instead, select the fields and the order in a running `esxtop` process, make changes, and save this file using the `W` interactive command.
- The `resxtop` utility is a vSphere CLI command.
- A load average of 1.00 means that there is full utilization of all physical CPUs.
- Memory overcommitment of 1.00 means a memory overcommitment of 100 percent.
- Managed systems run SNMP agents, which can provide information to a management program in at least one of the following ways:
    - In response to a GET operation, which is a specific request for information from the management system.
    - By sending a trap, which is an alert sent by the SNMP agent to notify the management system of a particular event or condition.
- Optionally, you can configure ESXi hosts to convert CIM indications to SNMP traps, allowing this information to be received by SNMP monitoring systems.
- Management Information Base (MIB) files define the information that can be provided by managed devices. The MIB files contain object identifiers (OIDs) and variables arranged in a hierarchy.
- vCenter Server and ESXi have SNMP agents. The agent provided with each product has differing capabilities.
- The vCenter Server SNMP agent functions only as a trap emitter and does not support other SNMP operations, such as GET.
- The traps sent by vCenter Server are defined in VMWARE-VC-EVENT-MIB.mib.
- 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.
- By default, the embedded SNMP agent is disabled. To enable it, you must configure it using the vSphere CLI command `vicfg-snmp`.
- IPMI sensors were used for hardware monitoring in ESX/ESXi 4.x and earlier. The conversion of CIM indications to SNMP notifications is newly available in ESXi 5.0.

What's New in Performance in VMware vSphere 5.0 Whitepaper

- Storage I/O Control (Storage I/O Control) now supports NFS.
- Virtual NUMA (vNUMA) exposes host NUMA topology to the guest operating system.
- vSphere 5.0 now enables users to choose to configure a swap cache on the SSD. VMware ESXi 5.0 will then use this swap cache to store the swapped-out pages instead of sending them to the regular and slower hypervisor swap file on the disk.
- Host Power Management (VMware HPM) in vSphere 5.0 provides power savings working at the host level. With vSphere 5.0, the default power management policy is "balanced." The balanced policy uses an algorithm that exploits a processor's P-states.
- A multi-network adaptor feature enables users to employ multiple network adaptors for vMotion. The VMkernel will transparently load-balance the vMotion traffic over all the vMotion-enabled vmknics in an effort to saturate all of the connections. In fact, even when there is a single vMotion, VMkernel uses all the available network adaptors to spread the vMotion traffic.
- vSphere 5.0 introduces a new latency-aware "Metro vMotion" feature that provides better performance over long latency networks and also increases the round-trip latency limit for vMotion networks from 5 milliseconds to 10 milliseconds. Previously, vMotion was supported only on networks with round-trip latencies of up to 5 milliseconds.

vSphere Troubleshooting (ESXi & vCenter 5.0)

- You must enable Hardware Virtualization (HV) before you use vSphere Fault Tolerance.

- You can only enable Fault Tolerance on a virtual machine with a maximum of 64GB of memory.
- Migrating a running fault tolerant virtual machine using vMotion also can fail if its memory is greater than 15GB or if memory is changing at a rate faster than vMotion can copy over the network. The default timeout window (8 seconds).
- Before you enable Fault Tolerance, power off the virtual machine and increase its timeout window with *ft.maxSwitchoverSeconds = "30"*.
- If you increase the timeout to 30 seconds, the fault tolerant virtual machine might become unresponsive for a longer period of time (up to 30 seconds) when enabling FT or when a new Secondary VM is created after a failover.
- Orphaned virtual machines exist in the vCenter Server database, but the ESXi host no longer recognizes them.
- Regenerate Certificates for an ESXi Host — back up any existing certificates in /etc/vmware/ssl. rui.crt and rui.key. Run the command `/sbin/generate-certificates` to generate new certificates.
- vCenter Server plug-ins that run on the Tomcat server have extension.xml files, which contain the URL where the corresponding Web application can be accessed. These files are located in C:\Program Files\VMware\Infrastructure\VirtualCenter Server\extensions. Extension installers populate these XML files using the DNS name for the machine.
- Reregister the plug-in after you edit its extension.xml file.
- When you have multiple vCenter Server instances, each instance must have a working relationship with the domain controller and not conflict with another machine that is in the domain.
- If the domain controller is unreachable, vCenter Server might be unable to start.
- If resolving the problem with the domain controller is impossible, you can restart vCenter Server by removing the vCenter Server system from the domain and isolating the system from its current Linked Mode group.
- vCenter Server uses Microsoft ADAM/AD LDS to enable Linked Mode, which uses the Windows RPC port mapper to open RPC ports for replication.
- vSphere HA or vSphere FT error messages: http://kb.vmware.com/kb/1033634.
- When you use the *Host Failures Cluster Tolerates* admission control policy, vSphere HA clusters might become invalid (red) due to insufficient failover resources.
- You might get a *not enough failover resources* fault when trying to power on a virtual machine in a vSphere HA cluster.
    - Hosts in the cluster are disconnected, in maintenance mode, not responding, or have a vSphere HA error.
    - Cluster contains virtual machines that have much larger memory or CPU reservations than the others.
    - No free slots in the cluster.
- If Storage DRS rules are preventing Storage DRS from making migration recommendations, you can set the Storage DRS advanced option *IgnoreAffinityRulesForMaintenance* to 1.
- When a datastore is shared across multiple datacenters, Storage DRS I/O load balancing is disabled for the entire datastore cluster. However, Storage DRS space balancing remains active for all datastores in the datastore cluster that are not shared across datacenters.
- ESXi hosts use the SCSI reservations mechanism only when storage devices do not support the hardware acceleration. For storage devices that support the hardware acceleration, the hosts use the atomic test and set (ATS) algorithm to lock the LUN.
- SCSI device drivers have a configurable parameter called the LUN queue depth that determines how many commands to a given LUN can be active at one time. If the host generates more commands to a LUN, the excess commands are queued in the VMkernel.
- If the sum of active commands from all virtual machines consistently exceeds the LUN depth, increase the queue depth. The procedure that you use to increase the queue depth depends on the type of storage adapter the host uses.
- Which HBA module is currently loaded by entering one of the following commands:
    - For QLogic: `esxcli --server=server_name system module list |grep qla`
    - For Emulex: `esxcli --server=server_name system module list |grep lpfc`
    - The software iSCSI module is: *iscsi_vmk*.
- The *iscsivmk_LunQDepth* parameter sets the maximum number of outstanding commands, or queue depth, for each LUN accessed through the software iSCSI adapter. The default value is 128.
- Setting the queue depth to a value higher than the default can decrease the total number of LUNs supported.
- If you adjusted the LUN queue depth, change the *Disk.SchedNumReqOutstanding* parameter, so that its value matches the queue depth. The parameter controls the maximum number of outstanding requests that each virtual machine can issue to the LUN. Change this parameter only when you have multiple virtual machines active on a LUN. The parameter does not apply when only one virtual machine is active on a LUN. In that case, the bandwidth is limited by the queue depth of the storage adapter.
- When your host uses a network adapter with iBFT, the software iSCSI adapter is always enabled by default.
- Attempts to mount NFS datastores with names in international languages (non-ASCII characters) can result in failures.

Solutions and Examples for VMware vSphere 5 (ESXi & vCenter 5.0)

- The administrative user name for the ESXi host is root. By default, the administrative password is not set.
- If DHCP is not available, the host assigns the link local IP address, which is in the subnet 169.254.x.x/16.
- When you replace default vCenter Server certificates, the certificates you obtain for your servers must be signed and conform to the Privacy Enhanced Mail (PEM), a key format that stores data in a Base-64 encoded Distinguished Encoding Rules (DER) format. The key used to sign certificates must be a standard RSA key with an encryption length ranging from 512 to 2048 bits. The recommended length is 1024 bits.

- You can use host profiles to specify how hosts join a directory service domain: with user-supplied Active Directory credentials or using the vSphere Authentication Proxy server (CAM server).
- When you join a host to an Active Directory domain, you must define roles on the host for a user or group in that domain. If you do not define such roles, the host is not accessible to Active Directory users or groups. You can use host profiles to set a required role for a user or group and to apply the the change to one or more hosts.
- Change the network policy for all VMkernel adapters, so that it is compatible with the network binding requirements. You can then bind the iSCSI VMkernel adapters to the software iSCSI or dependent hardware iSCSI adapters.
- If the NIC you use with your iSCSI adapter, either software or dependent hardware, is not in the same subnet as your iSCSI target, your host cannot establish sessions from this network adapter to the target.
- Duplicate SIDs do not cause problems when the computers are part of a domain and only domain user accounts are used. However, if the computers are part of a Workgroup or local user accounts are used, duplicate SIDs can compromise file access controls.
- If a new Windows virtual machine encounters customization errors while it is booting, the errors are logged to %WINDIR%\temp\vmware-imc.
- If a new Linux virtual machine encounters customization errors while it is booting, the errors are reported using the guest's system logging mechanism. View the errors by opening /var/log/vmware-imc/toolsDeployPkg.log.

VMware vSphere 5.0 — Licensing, Pricing and Packaging Whitepaper

- The vSphere 5.0 licensing model is per processor (CPU) with pooled vRAM entitlements. Benefits:
  - Simplicity – removes two physical constraints (core and physical RAM)
  - Flexibility – resource pooling
  - Fairness – aligns cost with actual use
  - Evolution – a cloud-like "pay for consumption" model
- vSphere 5.0 licensing removes all restrictions on physical cores and physical RAM.
- The vRAM entitlements of vSphere CPU licenses are pooled—that is, aggregated—across all CPU licenses managed by a VMware vCenter instance (or multiple linked VMware vCenter instances).
- The easiest way to expand pooled vRAM capacity is to add more vSphere CPU licenses of the same edition to the vRAM pool. Alternatively, customers can upgrade all CPU licenses in the vRAM pool to a vSphere edition with a higher base vRAM entitlement.
- Available and consumed vRAM capacity can be monitored and managed using the licensing-management module of VMware vCenter Server. Customers can create reports and set up alerts to obtain automated notification of when the level of vRAM consumption surpasses a specified level of the available pooled capacity
- When a license key is assigned by vCenter Server, it is copied to the host and saved in a persistent format. If the host becomes disconnected from vCenter Server, the license key remains active on the host indefinitely.
- You also have the option to assign their license keys directly to individual hosts.
- vCenter includes Orchestrator and Linked Mode in vCenter Standard.

♦ New in vSphere 5.0	Essentials	Essentials Plus	Standard	Enterprise	Enterprise Plus
Entitlements per CPU license					
• vRAM Entitlement	32 GB	32 GB	32 GB	64 GB	96 GB
• vCPU	8 way	8 way	8 way	8 way	32 way
Features					
Hypervisor					
High Availability					
Data Recovery					
vMotion					
Virtual Serial Port Concentrator					
Hot Add					
vShield Zones					
Fault Tolerance					
Storage APIs for Array Integration					
Storage vMotion					
Distribute Resource Scheduler & Distributed Power Management					
Distributed Switch					
I/O Controls (Network and Storage)					
Host Profiles					
Auto deploy ♦					
Profile-Driven Storage ♦					
Storage DRS ♦					

- Customers who purchased vSphere 4.x Standard with vMotion and Storage vMotion are entitled to vSphere 5.0 Enterprise Edition.
- vRAM is the total memory configured to a virtual machine (*Forbes*: Only powered on VMs).
- Only vSphere Essentials and Essentials Plus implement hard enforcement of vRAM capacity.
- Additional notes from www.virtual-al.net:
  - vRAM is pooled per license type
  - Make sure all license keys you have purchased are entered into vCenter
  - All licenses (whether deployed or not) add pooled vRAM capacity
  - Make sure existing vCenter servers are linked — vRAM licenses are pooled across linked mode vCenters giving a total vRAM per license type.
  - Updates:
    1. VMware have increased vRAM entitlements for all vSphere editions.
    2. The amount of vRAM has been capped at 96GB, so if your VMs have over this amount it will only count the first 96GB.
    3. The amount of vRAM is now calculated as a 12 month average of allocated vRAM.
- The following vRAM entitlements per edition will now be granted:

License	vRam
Enterprise Plus	96 GB
Enterprise	64 GB
Standard	32 GB
Essentials Plus	32 GB
Essentials	32 GB
Free vSphere Hypervisor	32 GB (Physical RAM per host)