

Maximums (FT advice): Disks per VM = 16		FT VMs per host = 4		Minimum hosts per cluster = 3	
FW ports	Source	Destination	Prot (ESX port)	Description	
8100, 8200	Hosts	ESX/ESXi	UDP (SC)	FT	
8100, 8200 (out)	ESX/ESXi	Hosts	TCP/UDP(SC)	FT	

FT: uses anti-affinity rules. **Requires** - HA & host monitoring, host certificate checking (on by default), dedicated logging NIC, compatible CPU, Hardware Virtualization (HV), thick disks on shared storage, supported guest OS. **Not supported** - snapshots, storage VMotion, DRS features, hotplugging, MSCS, VCB, SMP, physical RDMs, Paravirtualized VMs, NPIV, VMDirectPath, EPT/PT.

MSCS: • Win 2000 SP4, 2003 SP2 & 2008 (Failover Clustering) • 32 & 64bit • only 2 nodes clusters **Not supported** - DRS/HA on VMs, VMotion, FT, NPIV, Round Robin NMP, iSCSI/NFS based disks

	VMDK	Virtual RDM	Physical RDM
Cluster in a box (CIB)	Yes (zeroed)	Yes	No (not supported)
Cluster across boxes (CAB)	No	Yes - not Win2008, not CCR	Yes (recommended)
Physical & VM (n+1)	No	No	Yes
Snapshots	Yes	Yes	Yes
SCSI target software	No	No	Yes

- Configure all RDMs before configuring VM's network settings, or initialising LUNs within windows.
- Add all RDMs to a 2nd SCSI controller i.e. SCSI(1:4). Set sharing to Physical or Virtual as required.
- SCSI bus sharing - CIB = Virtual & CAB or N+1 = Physical

Links: <http://kb.vmware.com/kb/101061> - Understanding FT
<http://kb.vmware.com/kb/1008072> - CPU & guest OS that support FT

Maximums (per host): 1GB VMNICs = 2 - 32 dependent on HW (e1000/e=32) 10GB VMNICs = 4
PCI VMDirectPath devices=8 vSS switches=248 VEM switches=1 vSwitch ports (vSS/vDS)=4,096
Service Console ports = 16 VMotion and IP storage (VMkernel) port group = 1
Maximums (per vCenter): vDS switches = 16 vDS port groups = 512 vDS switch ports = 6,000
Maximums (per switch): Hosts (per vDS) = 64 vSS port groups = 512 vSS switch ports = 4,088

Terminology: VMNICs - logical name for physical server NICs vNICs - virtual NICs assigned to VMs
vSS - vSphere Standard Switch vDS - virtual Distributed Switch vPortGroup - vSwitch port group
dyUplink - uplink VMNICs on a vDS NetworkVMotion - tracking of VM's network state on a vDS

Common networking commands (-h switch for options or man page for detailed description):

```
List VMNICs:                $ sudo /usr/sbin/esxcfg-nics -l
List vSwitches & Port Groups: $ sudo /usr/sbin/esxcfg-vswitch -l
List Service Console ports:  $ sudo /usr/sbin/esxcfg-vmnic -l
List VMkernel ports:         $ sudo /usr/sbin/esxcfg-vmknic -l
List VMkernel Default Gateway: $ sudo /usr/sbin/esxcfg-route
```

Common networking configuration files: Name resolution order: /etc/nsswitch.conf
Name resolution file: /etc/host.conf DNS servers: /etc/resolv.conf
Network Tagging: /etc/network/interfaces Network configuration file: /etc/sysconfig/network

Ethernet tagging: -L [External Switch Tagging] - Default No trunking required. 1:1 relationship from VMNICs to physical switch tags. Each VMNIC can only see 1 subnet. VLAN ID, 1 or blank, =
• VST (Virtual Switch Tagging) - Commonly used. VMNICs connected to a vSwitch can span several VLANs. Each Port Group has a VLAN ID of 1-4094. Set the VLAN ID to blank to use Native VLAN.
• VGT (Virtual Guest Tagging) - Rarely used. Install 802.1Q trunking driver software in the VMs, the vSwitch preserves the tags given by the VMs. VLAN ID of 4095 on vSS, VLAN policy on vDS.

SNS & VDS options (options can also be overridden on individual Port Groups):

- General • Number of ports** - by default 56 for vSS, 64 for vDS, 128 when created on Service Console. (not a Port Group option) • **Network label & VLAN ID** - only on Port Groups not vSwitches.

Security • Promiscuous mode (default Reject) - only listens to traffic destined for its MAC address.

- **MAC Address Changes** (default Accept) - accepts inbound frames when VM changes MAC address.
- **Forged Transmits** (default Accept) - won't drop outgoing frames if source MAC address is different.

Traffic Shaping • Status (default Disabled) **Average Bandwidth** (default 102400 Kbps) **Peak Bandwidth** (default 102400 Kbps) **Burst size** (default 102400 KB) - shapes out on vSS, instead on vDS.

NIC Teaming • Load Balancing (spreads outgoing traffic from vNICs across vMNICs) - & Originating port (allows incoming traffic to be sent to any vNIC on the same NIC teaming group).

VLAN • ID (default 0) - sets the VLAN of each packet (when physical switch ports are etherchannel). Source MAC hash based on source MAC address - Use explicit failover order. Incoming traffic is load balanced by physical switch.

- Network Failover Detection Link status only (default) detects cable pulls & switch power failures, not misconfigurations. **Beacon Probing** don't use with IP-hash load balancing.
- Notify Switches - No or Yes (default) updates lookup tables. Disable for MS NLB in unicast mode.
- Failback - No or Yes (default) VMNIC will return after recovering from a failure.
- Failover order **Active - Standby - Unused**. Don't use standby uplinks with IP-hash load balancing.

VLAN (vDS only) • **VLAN** - set ID. **Trunk range** - restrict IDs on trunked links. **PVLAN** - see below.

Miscellaneous (vDS only) • **Port blocking** - selected or unselected (default) block all ports.

vDpPort options: • **Port Binding State** when initially connected **Dynamic** when connected/powered-on/**Ephemeral** no binding • **Traffic shaping Ingress** into vSwitch **Egress** out of vSwitch • Allow live port migration • **Configure network** - Host can assign ports if vCenter is down • Name format **PVLAN** (Private VLAN) extension to allow user defined names rather than standard ones. Not encapsulated. Primary PVLAN - Original VLAN divided into smaller groups. Secondary PVLAN - exists only within primary, has specific VLAN ID. **Secondary types:** Promiscuous - connect with VMs in primary. Community-connect to themselves & VMs on promiscuous Isolated-connect with VMs on promiscuous TSO (TCP Segmentation Offload): enabled by default on VMkernel ports, allows very large frames (up to 64KB), even with smaller MTU. To enable on VMs, they need enhanced vmxnet3 NIC.

Jumbo frames up to 9Kb. Must be enabled for each vSwitch. VMs need enhanced vmxnet to use it.

NetQueue enabled by default, allows certain VMNICS to spread processing across multiple CPUs.

Configure networking (for vSS): (1) add a vSwitch **esxcfg-vswitch -a** (2) add a port group to the vSwitch **esxcfg-vswitch -A** (3) set the port group's VLAN ID **esxcfg-vswitch -p -v** (4) add the VMNIC to the vSwitch **esxcfg-vswitch -l**

- VM connections: set the VMs NIC to use the port group.
- Service console: create interface & link to the port group **esxcfg-vswif -a -p -i -n**, then edit the DC in /etc/sysconfig/network, then restart network **service network restart**.
- VMkernel ports: set the port **esxcfg-vmknic -a -i -n** & set the VMkernel DG **esxcfg-route -VMotion** should be enabled in vCenter if required.

Links: <http://kb.vmware.com/kb/1010812> - Configure IPv6
<http://vmware.com/files/pdf/vsphere-network-ds-migration-configuration-wp.pdf> - vDS whitepaper
<http://kb.vmware.com/kb/1000258> - Configure networking from Service Console command line

Maximums (per DRS cluster): Hosts = 32 VMs (powered on) = 1280 (limit of 256 per host)
Maximums (per Resource Pool): Children = 1024 Tree depth = 12 (10 when in a DRS cluster)
Maximums (other): Datacenters per host = 100 RPs per host = 4096 RPs per cluster = 512
Datacenters mark organisational & vMotion boundaries. **Clusters** gather host CPU & memory resources. **Resource Pools** apply policies to clusters. A DRS cluster is also implicitly a resource pool.
Resource pools: • Shares - low, medium & high (1,2,4) • Reservations - MHz(CPU)/MB(RAM) • Limits - MHz/MB • Expandable reservation - yes (can draw from parent's pool) - no (no can't draw)
 • Shares only during resource group contention (e.g. \$ / sro /r / e/b/n / ex/c/g / r - res/grp - l)
 • Shares only while during reservations relative to siblings. Reservations guarantee a minimum are only checked when a VM is powered on. Limits are an upper bound, never exceeded; manage user expectations but can waste idle resources. Expandable reservations do not automatically hunt upwards, define if reservations are considered by admission control. Child pools actively reserve resources from parent even if VMs are powered off. Hierarchical resource pools require DRS enabled.
DRS: priority levels 1-5 (1 the highest). **DRS cluster settings:** • Manual / Partial (Initial VM placement) • Fully Automated (Initial VM placement & Dynamic balancing). **Current Host Load Standard Deviation:** load imbalance. Lower than Target value unless recommendations are unapplied.
"Grafted from" pools created when adding a host to a DRS cluster & keeping the host's resource pool hierarchy. **Maintenance mode** only clears VMs off host if DRS cluster is fully automated.
Affinity Rules keep VMs together or apart in a DRS cluster. Anti-affinity rule limited to 2. **Rule conflicts** - older wins, new rule disabled. Anti-affinity rules over affinity. Disabled rules ignored.
Current host load standard deviation: DRS load balance. Current < target unless advice unapplied.
Hosts: • Verify host health status • Verify host configuration • Verify host thresholds and dependencies • Verify host's DPM Last Time Exited Standby, DPM level - Off, Manual (makes recommendations) & Automatic. Hosts reclaim memory from VMs by: • Balloon driver (vmmemctl) force guest to use native algorithms (guest swap) • VM Swap files (if vmtoolsd not available/responsive) • Sharing memory across VMs
Links: <http://kb.vmware.com/kb/1005764> - Enhanced vMotion (EVC) FAQ
<http://kb.vmware.com/kb/1003212> - EVC CPU compatibility
<http://www.yellow-bricks.com/drs-deepdive/> - DRS Deep Dive

Maximums (per host): Volumes = 256		Paths = 1024		NAS datastores = 8 (64 with adv settings)	
FC HBAs = 8 (HBA ports = 16)		targets per HBA = 256		paths to each LUN = 32	
iSCSI HW HBAs = 4		targets per HBA = 64		Dynamic (61 Static) paths to each LUN = 8	
iSCSI SW - NICs = 8		targets = 256		paths to each LUN = 8	
Maximums (per volume): VMs = 256		Hosts = 64 (DS limit, 2048 for linked clones)			
VMFS = 64TB (less 16KB)		NFS = 16TB		File size (1/2/4/8MB blocks) = 256GB/512GB/1TB/2TB	
RDMs = 2TB (less 512B)		Extents = 32		Extent size = 2TB (less 512B)	
FW Port	Source	Destination	Prot	(ESX port)	Description
2049	NFS server	ESX/ESXi	TCP (VMMK)	NFS Client	
2049	ESX/ESXi	NFS server	TCP (VMMK)	NFS Client	
3260	ESX/ESXi	iSCSI server	UDP (SC-VMMK)	iSCSI Client	
Common storage commands (-h switch for options, or man page for detailed description):					
List all storage devices: \$ sudo /usr/sbin/esxcfg-scsidevs -c					
List LUNs, paths & multipathing plugins: \$ sudo /usr/sbin/esxcfg-mpath -l					
List all VMware SATPs: \$ sudo /usr/sbin/esxcli mp satp list					
List claim rules: \$ sudo /usr/sbin/esxcli corestorage claimrule list					
List datastores, dev names to VMFS: \$ sudo /usr/sbin/esxcfg-scsidevs -m					
List snapshot volumes: \$ sudo /usr/sbin/esxcfg-volume -l					
Test VMkernel connectivity: \$ /usr/sbin/vmknop					
Manage HW iSCSI (Qlogic) settings: \$ sudo /usr/sbin/esxcfg-hwiscsi -l					
Manage SW iSCSI settings: \$ sudo /usr/sbin/esxcfg-swiscsi -q					
List iSCSI LUNs: \$ sudo /usr/sbin/vmkniscsi-tool -l -l adapter					
Rescan iSCSI LUNs: \$ sudo /usr/sbin/esxcfg-rescan adapter					
List the NFS exports from the VMkernel: \$ sudo /usr/sbin/esxcfg-nas -l					
Storage capabilities					
VMotion, DRS, HA, FT, VCB, SRM & Thin VMMDKs	Yes	Yes			Yes
VMFS volumes, RDMs & VMware's NMP	Yes	Yes			No
Boot ESX host	Yes	Yes	Yes (HW initiator)		No
VM MSCS clustering	Yes	No			No

Active-active: access to the LUNs simultaneously through all ports, without performance degradation.

Active-passive: one port actively providing access, other as backup. Path thrashing can occur.

NPIV (N-Port ID Virtualization): FC HBA port assigned dedicated virtual port (WWPN) to VM (RDM)

LUN addressing: FC: `RuntimeName vmbus#CH:Tr:Ln - adapter:channel:target:LUN`

ISCSI: `IQN iqn.year-month-reversed.domain_name:string or iqn .eui.string`

ISCSI discovery methods: Static - can manually add/remove items, only with hardware initiators.

Dynamic - uses "SendTargets" command, will scan for all available targets, remove any not seen/reset

CHAP: HW ISCSI - CHAP initiator level, SW ISCSI - via optional CHAP, initiator target

VMkernel Port is required to use ISCSI or NFS storage. (S.C. port not required for ISCSI anyone)

MPP (MultiPathing Plugins): claim rules in `/etc/vmware/esx.conf` specify MPP to use for each path.

Claim rules: indicate which MPP, native or 3rd party, manages a given physical path.

NMP (Native MPP) - sATPs (Storage Array Type Plugins) - handles failovers, - PSPs (Path Selection Plugins) - handles load-balancing. NMP policies: Fixed - default for active/active, use preferred path when available. MRU (Most Recently Used) - default for active/passive (& ISCSI), first working path found at boot. RR (Round Robin) - safe for all arrays, rotates through paths (not MSCS LUNS).

Disk.MaxLUN: reduce number of LUNs created. `Disk.MaskLUN: convert to claim rule format.`

VMFS volumes: Large-Less LUNs to create, less to manage, flexible resizing & snapshots. Small-less LUNs (resizing), less waste of space, different LUNs, multiple multipathing disk shares.

TPG (Target Port Group Support): allows path performance, host can determine best path

ALUA (Asymmetric Logical Unit Access): finds/manages multiple paths for failover & load balancing

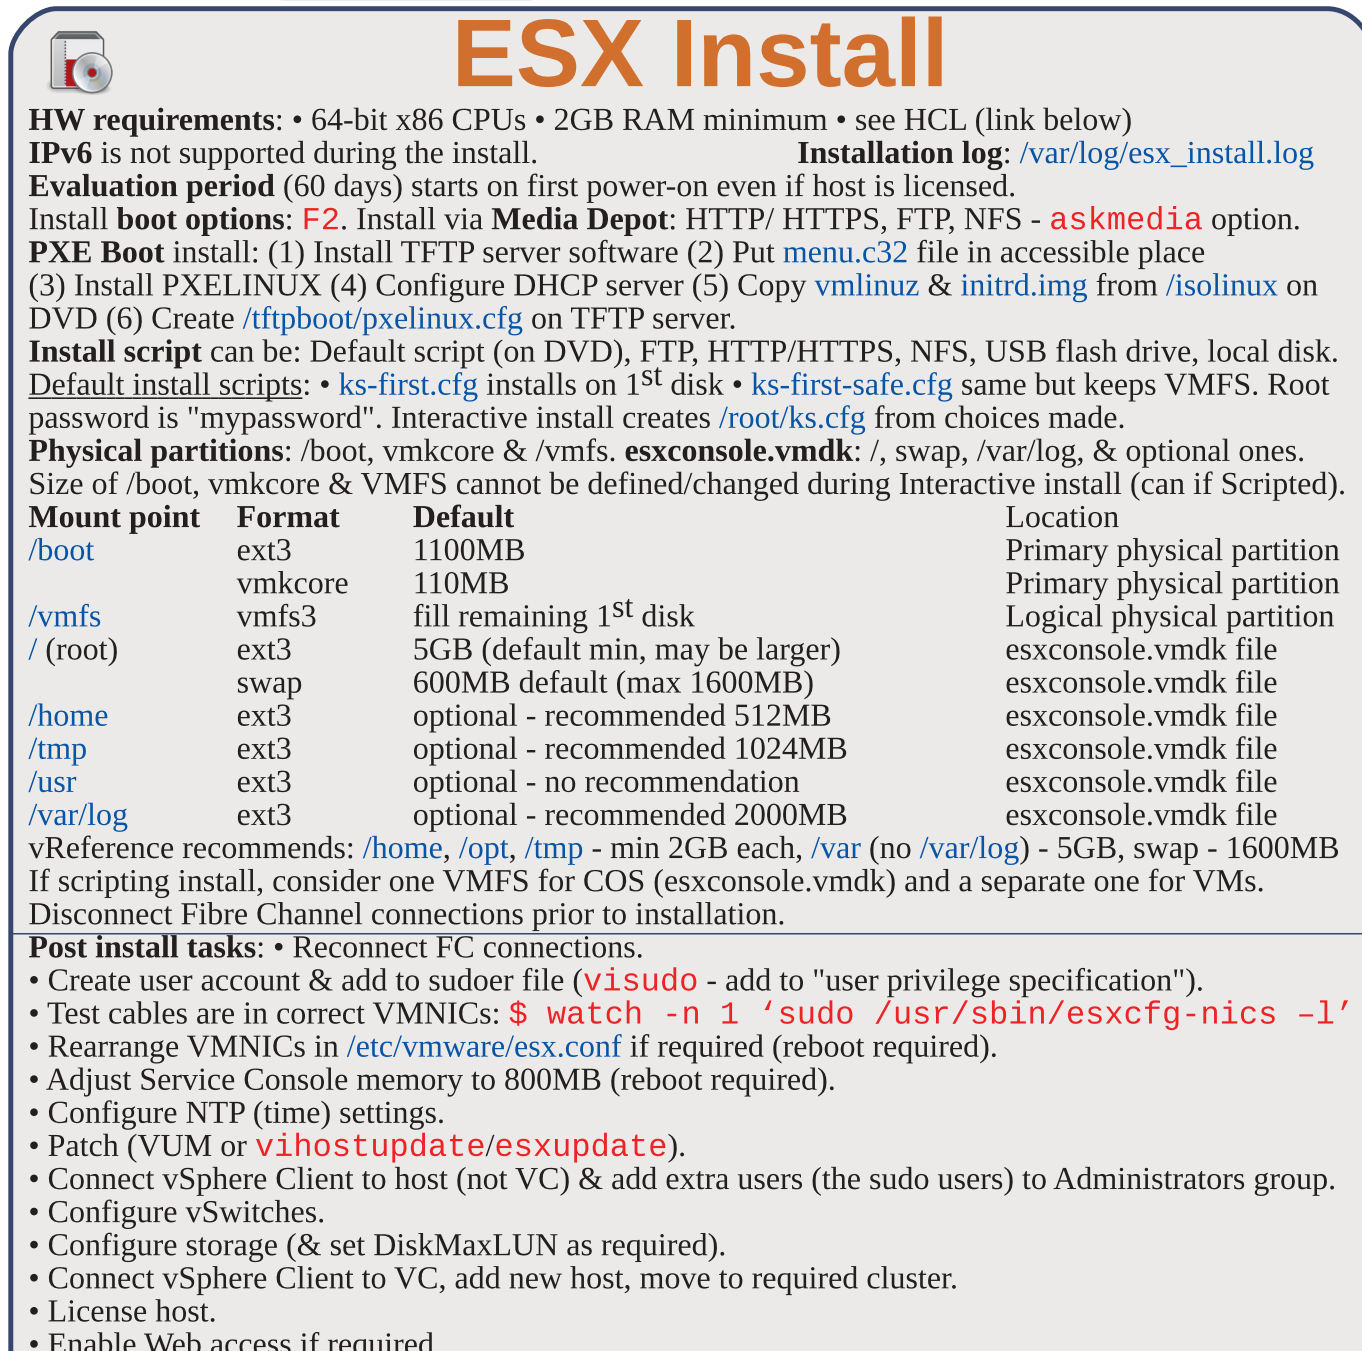
Links: <http://kb.vmware.com/kb/1009532> - Load connectivity to storage

<http://media.netapp.com/documents/r-3749.pdf> - Storage best practices whitepaper (NetApp)

<http://media.netapp.com/documents/r-3747.pdf> - File System alignment whitepaper (NetApp)

<http://kb.vmware.com/kb/1011387> - ESX 4.4 handling of LUNs detected as snapshot

This document is licensed under a Creative Commons License. Refer to <http://www.creativecommons.org> for full details. The artwork is from the Tango Project <http://tango.freedesktop.org> under their Creative Commons license.



Upgrade from ESX3: (cannot use DVD) `/VM/VMXCenter/Update Manager` - upgrades ESX3/ESXi.
 * Hosts: Update utility - upgrades ESX3/ESXi & updates ESX/ESXi, small environments (< 10 hosts, no VMs). Customized ESX3/ESXi: `ESXCLI` - updates ESX/ESXi, small environments (< 10 hosts, no VMs).
ESX upgrade: `esxupdate.sh` install <http://kb.vmware.com/kb/1009440> - upgrades ESX only.
Upgrade tools: `esx3-install esx3-upgrade/ & /var/log/vmware/`
 Unsuccessful upgrades: `esx3-upgrade/ & /var/log/vmware/`
Post upgrade: * Upgrade VMware Tools before upgrading virtual hardware * Re-install 3rd party agents/apps * Convert LUN masking to claim rule method: `esxccli corestorage claimrule convert` * Successful upgrade: `cleanups-esx3` removes ESX3 boot options & ability to roll back
Links: <http://www.vmware.com/resources/compatibility/search.php> - Hardware Compatibility Guide <http://kb.vmware.com/kb/1009393> - Upgrading to ESX 4.0 or ESXi 4.0 Upgrade Practices <http://kb.vmware.com/kb/1009039> - Upgrading to ESX 4.0 & vCenter 4.0 best practices <http://kb.vmware.com/kb/1010675> - Upgrading an ESX 3 VM to ESX 4.0 <http://kb.vmware.com/kb/1011712> - See if Intel VT or AMD-V is BIOS enabled without rebooting

SW requirements: vSphere Client: Windows with .NET 3.0 SP1 framework. Web Access: Win - 2003 SP1, XP pro SP3, XP home SP2, 2000 SP4, Linux - GTK+ - 2. Browsers - IE6, 7, or ≥, Firefox 2, 3, or ≥				
FW Port	Source	Destination	Protocol	Description
22	SSH client, Web Access	ESX	TCP	SSH access
80	Web Access	ESX, VC	TCP	Redirect to HTTPS
427	Clients, Web Access	ESX/ESXi	TCP	CIM SLP client
443	Clients, Web Access	ESX/ESXi, VC	TCP	HTTPS
902	Clients, Web Access	ESX/ESXi	TCP	Authentication
903	Clients, Web Access	ESX/ESXi	TCP	VM Console
5989	Clients, Web Access	ESX/ESXi	TCP	CIM transactions

Logs: Client Agent log `/var/log/vmware/vpx/vpxa.log` Client Install log `%TEMP%\vmmis.log`
 Client Service log `C:\Docs and Settings\username\Local Settings\AppData\local\vmclient-x.log (x=0-9)`
Web Access to ESX or VC: `http://hostname.domain.com/vi/` - ESXi - no WebAccess • ESX - disabled
 Web Access status check: `$ sudo /sbin/service vmware-webAccess status`
 Web Access Remote Console URLs: • Limit view to remote console - hides details like event logs
 • Limit view to single VM - disables inventory navigation. Permission to VMs still granted in ESX or vCenter. Alarms tab available connected to vCenter (not ESX). Web Access allows only viewing tasks.
Links: <http://www.june.nl/articles/vmware/143-vcenter-client-shortcuts> - vCenter client shortcuts

Maximums (per host): vCPUs = 512		vCPUs per physical core = 25		Logical cores (incl HT) = 64	
RAM = 1TB		Service Console RAM = 800MB (min=300MB)		VMs = 320 (HA can limit)	
FW Port	Source	Destination	Prot (ESX port)	Description	
22	SSH client	ESX	TCP (SC)	SSH server	
53 (out)	ESX/ESXi	DNS server(s)	UDP (SC)	DNS requests	
80	Clients	ESX/ESXi	TCP (SC)	HTTP access	
123 (out)	ESX/ESXi	NTP source	UDP (SC)	NTP (time) client	
427	Hosts, Client	ESX/ESXi	UDP (SC)	CIM SLP client/server	
427 (out)	ESX/ESXi	Hosts	UDP (SC)	CIM SLP client/server	
443	Hosts, Clients, VC	ESX/ESXi	TCP (SC)	HTTPS access	
902	Hosts, Clients, VC	ESX/ESXi	TCP (SC)	Auth, migrate, provision	
902 (out)	ESX/ESXi	Hosts, VC	UDP (SC)	Auth, migrate, provision	
9303	Clients	ESX/ESXi	TCP (SC)	VM Console	
5900-5964	?	ESX/ESXi	TCP (SC)	RFB for mgt tools (VNC)	
5900-5964 (out)	?	ESX/ESXi	TCP (SC)	RFB for mgt tools (VNC)	
5989	Clients	ESX/ESXi	TCP (SC)	CIM server over HTTPS	
5989 (out)	ESX/ESXi	Hosts	TCP (SC)	CIM server over HTTPS	
8000	Hosts	ESX/ESXi	TCP (VMK)	VMotion requests	
8000 (out)	ESX/ESXi	Hosts	TCP (VMK)	VMotion requests	
Possible exceptions: 21(FW), 22(out)(SSH), 53(DNS), 88/389/464(AD), 161/162(SNMP), 445(SMB), 5988(CIM)					
Logs: Service Console Availability & VMkernel Messages, Alerts, Availability: /var/log/vmmon.log					
ESX service log: /var/log/vmware/hostd.log			Syslog: /var/log/messages		
VMkernel warnings: /var/log/vmkwarning			VMkernel events: /var/log/vmksummary		
VC agent: /var/log/vmware/vpxa.log			Patching: /var/log/vmware/esxupdate.log		
Common ESX host commands (-h switch for options or man page for detailed description):					
List status of all services: \$ sudo /sbin/service --status-all					
List the service runlevels: \$ chkconfig --list					
Restart a service: \$ sudo /sbin/service service_name restart (start, stop, status available)					
Common services: mgmt-vmware (hostd) vmware-vpxa (vCenter agent) vmware-vmkauthd (authentication) network (vsfif changes) vmware-webAccess (Web Access)					
Show build number: \$ vmware -v					
Check the filesystem usage: \$ sudo vdf -h					
List diagnostic partitions: \$ sudo /usr/sbin/esxcfg-dumpart -l					
Show vcenter log: \$ vmkerrcode errcode_number					
Export detailed config file: \$ sudo esxconfig -info -w /tmp/esxconfig-info.txt					
Gather debugging report: \$ sudo /usr/bin/vm-support -w /tmp					
Configure authentication settings: \$ sudo /usr/sbin/esxconfig-auth					
Lists drivers loaded at startup: \$ sudo /usr/sbin/esxconfig-module -l					
Set advanced options: \$ sudo /usr/sbin/esxconfig-advconfig option -s value (-g to get)					
Update bootstrap settings: \$ sudo /usr/sbin/esxconfig-boot (treat with caution)					
Initialization routines (resets things): \$ sudo /usr/sbin/esxconfig-init (treat with caution)					
Internal firewall commands (ipables on Service Console):					

List the firewall named `esxcfg`: `$ sudo /usr/sbin/esxcfg-firewall -s`
 Enable a service: `$ sudo /usr/sbin/esxcfg-firewall -e service_name (-d to disable)`
 To open a port: `$ sudo /usr/sbin/esxcfg-firewall -o open_port[,protocol, direction, name]`
Security Levels: High - in/out blocked, Medium - in blocked, out open, Low - in/out open.
 By default all traffic blocked in and out, except `22`, `47`, `443`, `902`, `5989`, `5988`, pings, DHCP & DNS.
Master config file: `/etc/vmware/esx.conf` **Certificates files:** generate new files if not present.
 Certificate public key `/etc/vmware/ssl/rui.crt` Certificate private key `/etc/vmware/ssl/rui.key`
 Set certificate location `/etc/vmware/ssl/proxy.rui` SSL settings `/etc/vmware/ssl/rui.conf`
PAM (Pluggable Authentication Modules) configuration: `/etc/pam.d/vmware-auth`
 Local authentication method is `etc/passwd`, vSphere is `vmware` for vCenter Server permissions.
Passwords: ESX user `root` is `root` to plug in the console. No default root password. Defaults for non-root users: password retries = 3, minimum password length = 9, shorter passwords if Characters Classes mixed (upper, lower, digits & other) M – C = `e, pam_pwdscd.so` provides more options.
User Password Aging: enabled by default, set to never expire (max days) & change anytime (min days = 0, warning = 7) • Change host settings: `esxcfg-auth` • Change user settings: `chage`
NUMA (Non-Uniform Memory Access): controls VM memory distribution across host memory. Only use NUMA if CPU affinity is set. **HTT** can help better utilize idle resources.
VMware MIBs: uses embedded SNMP agent (disabled by default). Enable: `vmcfg-smnpd`
syslogd: to configure = `ESX` - edit `/etc/syslog.conf` • `ESXi` - use Client or `vicfg-syslog`
Links: <http://kb.vmware.com/kb/653> - Collecting diagnostic information for ESX Servers
<http://kb.vmware.com/kb/1005144> - Decoding Machine Check Exception output after purple screen
<http://kb.vmware.com/kb/1000258> - Troubleshooting delayed boot on ESX hosts
<http://kb.vmware.com/kb/1000258> - Video: Configure Service Console networking from CLI
<http://kb.vmware.com/kb/1911992> - VMotion compatibility for Intel / AMD processors
http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.0-Interp - Performance best practices
<http://communities.vmware.com/docs/DOC-9279> - Interpreting esxtop Statistics

HW requirements: 64bit x86 CPUs, 2GB RAM, SATA, SAS or SCSI disks. No ESXi WebAccess.
ESXi Installable starts in evml mode (60 days). If no DHCP at install, link local IP used 169.254.x/x.16.
ESXi Installable Partitions: 4GB VFAT scratch for system swap (not required, but stores vm-support), 110MB diagnostic for core dumps, VMFS3 on free space.
Not supported: ESXi Installable & Embedded on same host • Booting multiple servers from 1 image
Direct Console: • Configuring host defaults • Set up administrative access • Troubleshooting
Restarting Mtg agents *etc/init.d* processes: hostd (mgmt-vmwarea), ntpd (time), sfcbd (CIM broker), slpd (discover/advertise services), wsman (share mtg info via SOAP), vobd (error reporting) & AAM (HA agent) if installed. To isolate ESXi host from DRS/HA cluster disable mtg network.
Management Network Test: pings Dg, primary DNS nameserver, secondary DNS, resolves hostname.
Lockdown mode: prevents remote access by root account, but not other accounts. Disabled by default.
vcfg-cfgbackup • Backup host configuration: -s: Restore: -l (-f if different build number)
Repair mode on ESXi Installable CD overwrites all configuration data. VMFS is preserved if VMFS is original location on boot disk (or beyond 900MB partition), or another disk.
Tech Support mode: • login to console: # Alt-F1 • # **unsupported** • enter root password
 Return to console: # **exit** • Alt-F2 <http://kb.vmware.com/kb/1003677> • Tech Support Mode KB

Maximums (per 32bit vCenter): Hosts = 200 VMs = 3000 Running VMs = 2000 Clients = 15
Maximums (per 64bit vCenter): Hosts = 300 VMs = 4500 Running VMs = 3000 Clients = 30
Maximums (Linked mode): vCenters = 10 VMs = 15000 Running VMs = 10000 Hosts = 1000
Maximums (operations per host): provisioning = 8 VMotions = 2 Storage VMotions = 2
Maximums (operations per datastore): provisioning = 8 VMotions = 4 Storage VMotions = 4
Maximum operations per vCenter = 96

HW requirements: Min - 2 CPUs, 3GB RAM • Medium (50 hosts, 250 VMs) 2 CPUs, 4GB RAM • Large (200 hosts, 2000 VMs) 4 CPUs, 4GB RAM • Extra large (300 hosts, 3000 VMs) 4 CPUs, 8GB RAM

SW requirements: • 32bit Windows up to 200 hosts, 64bit 200-300 • hostname - 15 characters or less.

Databases: • SQL 2005 Express (up to 5 hosts & 50 VMs) • SQL 2005 (use SQL Native Client v9) • SQL 2008 (SQL Native Client v10) • Oracle 10g & 11g • IBM DB2 9.5. Not SQL 2000 nor Oracle9i.

VC needs 32-bit ODBC DSN, if VC 64-bit use [C:\WINDOWS\SYSTEM64\odbc32.exe](#). User needs DB2 rights. Default of max 10 simultaneous DB connections. MSSQL - don't use master DB.

Pre-Upgrade Checker Tool: on vCenter DVD, checks for potential issues with hosts prior to upgrade.

FW Port	Source	Destination	Protocol	Description
80	Clients	VC	TCP	Redirect to HTTPS
389	VC	AD DCs	TCP	AD lookup
443	Clients	VC	TCP	VIC & WebAccess
443	VC	Hosts	TCP	vCenter agent
902	Hosts	VC	UDP	Heartbeat
902	Hosts, Clients	VC	UDP	Heartbeat
902	Hosts, Clients	VC	TCP	VM Console

Possible extras: 22/135/137-139/445/9089(guided consolidation)/25(SMTP)/53(DNS),80(redirects),88/445(AD),161/162(SNMP),389(LDAP),636(Linked VCs),1433(MSSQL),1521(Oracle),8080/8443(webervices),8181/8182(collection server),27000/21000(license),3 hosts, x3.

Logs: DB upgrade: [%TEMP%\VCDatabaseUpgrade.log](#) VC agent: [%varlog\vmware\log\vpvpxa.log](#)
 VC install: [%TEMP%\VC directory of user installing VC](#) VC logs: [%varlog\vmvpx\vpvxd.log](#)

Default roles (System roles - permanent, cannot change privileges, ESX/ESXi & VC. Sample just vC):

- No access System - Default except users in Admin Group. Cannot view or change.
- Read only System - View state & details except console tab.
- Administrator System - All privileges. Default for members of the Admin Group.
- VM power user Sample - Interact with, change VM settings, snapshots & schedule tasks.
- VM user Sample - Interact with host, insert media & power ops. Not change VM settings.
- Resource pool admin Sample - Create, modify child pools & assign VMs, but not RP itself.
- Consolidated backup user Sample - Used by Consolidated Backup product, don't modify.
- Datastore consumer Sample - Allows use of the datastore.
- Network consumer Sample - Allows network to be assigned to hosts or VMs.

Logged. Users initially granted No Access role on new objects, including databases/networks.

Involved in users removed from domain keep permissions until next validation period (default 24 hrs).

Tasks - activities that don't complete immediately. All roles allow schedule tasks by default. Can schedule tasks if user has permission when tasks created. VC Local Administrators have same rights as Administrator role by default. root & vxuser are only users not assigned No Access role on hosts by default. Propagation is per permission, not universal. Child permissions override those propagated. User permissions override Group ones. Can't set vDS permissions, set on parent & propagate.

License	ESXi Single	Essential	Essential+ ----- Essentials -----	Standard	Advanced ----- Foundation & Standard editions -----	Enterprise	Enterprise+ -----
vCenter	No	No	-----	6	12	6	12
Cores per socket	6	6	6	6	12	6	12
vSMP	4-way	4-way	4-way	4-way	4-way	4-way	8-way
Physical RAM	256GB	256GB	256GB	256GB	256GB	256GB	no limit
Thin provisioning	Yes	Yes	Yes	Yes	Yes	Yes	Yes
vxpa, iPrng, VMSSafe, vStorage HA	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Data Recovery			Yes	Yes	Yes	Yes	Yes
Hot Add, FT, vShield, VMotion			Yes	Yes	Yes/ESXi	Yes	Yes
Storage VMotion, DRS			Yes	Yes	Yes	Yes	Yes
vDS, Host Profiles, 3rd party MMP						Yes	Yes

Licensing: 25-character license keys, managed in VC. vSphere (ESX/ESXi) & vCenter Licenses.

Expiring licenses: vCenter - hosts are disconnected. ESX/ESXi - VMs run, but cannot power-on/reset.

Statistics: CPU, memory, disk, network, system, & VM ops. **Collection Intervals** (time stats - archived in DB): 5mins - 1 day, 30 mins - 1 week, 2 hrs - 1 month, 1 day - 1 year. **Real-time stats** stored in flat file on hosts & VC memory (not in DB), collected every 20 seconds. ESX - kept for 1 hr, ESXi - kept for 30 mins. **Collection level** 1-4 for each interval, 4 has most counters (default is 1). **Datstore metrics** only available in overview charts (not advanced charts). **Reports & Maps** updated every 30 mins.

Alarms: notifications of selected events, conditions & states. Composed of trigger & action. **Triggers:** condition/state triggers (monitor VMs, hosts & datastores - equal to/not equal to &/above/below) & event triggers (any object, VC or license server - arguments, operators & values). **Actions:** responses to triggered alarms. Default alarms don't have actions associated. Can disable action without disabling alarm, but effects actions on all alarms. Disable for selected object, child continues. Reduce alarms with tolerance range & trigger frequency (default 5 mins). Disconnect hosts to suspend monitoring.

Linked mode: joins VCS. Global data: IP & ports, certificates, licensing, user roles. Uses ADAM (AD App Mode) to store & sync data. Instances can run under different domain accounts. Installed by administrator who is admin on both machines. Requirements: DNS, 2-way trust if different domains, time sync, DNS name matches hostname. Roles are replicated, assignments of roles are not.

Server settings: licensing (vCenter & 3.x), statistics (intervals & DB size), runtime settings (unique ID, managed IP, name), AD (timeouts, query limit, validation period), mail, SNMP receivers, (SNMP ports, client timeouts, logging detail, DB connections, DB retention, SSL host verification, advanced settings).

Links: <http://kb.vmware.com/kb/101641> - Collecting diagnostic information for vCenter
<http://kb.vmware.com/kb/1009080> - Installing ESX 4.0 & vCenter 4.0 best practices
<http://kb.vmware.com/kb/1009039> - Upgrading to ESX 4.0 & vCenter 4.0 best practices
<http://kb.vmware.com/kb/1005593> - sypres file locations and versions
<http://kb.vmware.com/kb/1010579> - Comparison of vSphere 4.0 & VI 3.x licensing
<http://kb.vmware.com/kb/1010839> - Video: Licensing management
<http://kb.vmware.com/kb/1010550> - Setting up vCenter Server in a MSCS

Maximums (per VM):		vCPUs = 8	RAM = 255GB	Swap file = 255GB (1 per VM)
SCSI adapters = 4	Devices per SCSI adapter = 15			IDE devices (disk or CD) = 4
Floppy drives = 2	vNICs = 10		Parallel ports = 3	Serial ports = 4
Remote consoles = 40	VMDirectPath devices = 2			VMDirectPath SCSI targets = 60

Files: .cfg	Earlier version of .vmtx file	.vmsd	VM's memory
.dsk	Earlier version of .vmdk file	.vmsn	Snapshot metadata
.hlog	VMotion log file	.vmsn	Snapshot state file
.lck-XXX	Locking file on NFS datastore	.vmss	Suspended state file
.log	VM activity log	.vmtl	Earlier version of VC template
.nvram	BIOS settings	.vmtm	Team data
.raw	Raw device e.g. tape device	.vmtx	VC template header
.rdm	RDM in Virtual Compatibility mode	.vmx	Primary configuration file
.rdmp	RDM in Physical Compatibility mode	.vsw	Extra configuration file for VMs in a team
.REDO	Earlier version of -delta.vmdk file		Swap file for overcommitted memory
.std	Earlier version of .vmss file		
.vmdk	Disk descriptor (also raw virtual disk for hosted products)		
-flat.vmdk	Raw virtual disks	000000.vmdk	Snapshot metadata
-ctk.vmdk	Changed Block Tracking file	000000-delta.vmdk	Snapshot differential file

Commands: List all registered VMs on a host: `$ sudo /usr/bin/vmware-cmd -c`
Create/modify VMXdataS, RDMs, VMFS volumes & storage devices: **vmxfs**tools (check man page)
Power Off = hard power off • Shut Down = soft with VMware tools • Reset = hard • Restart = soft
VM HW: Memory/PCI Hotplug – VMware Tools must be installed.
VMs with HW earlier than v4 had removed compatibility & capabilities. Cannot add/remove devices.
Manually MAC addresses: 00:50:56:x.y.z. Set in vmx: ethernet.network.addressType="static".
Disk types: **zerodethick** (lazy) default, pre-allocates, **eagerzerodethick** select "Support clustering features such as FT", pre-allocates & zeros, better performance, slower creation, thin allocates on-demand, monitor with "datastore usage" alarm. **NFS:** type determined by array. **Independent disks:** no snapshots. **Persistent changes** immediate & permanent. **Nonpersistent** changes lost on power-off.
RDM: Benefits User-Friendly Persistent names, Dynamic Name Resolution, Distributed File Locking, File Permissions, File System Ops, SAN Snapshots, vMotion, SAN mgmt agents & NPIV. **Limitations** not for block devices, no snapshots with physical RDMs, no partition mapping, needs whole LUN.
Snapshots: capture memory state, settings & disks. Can't snapshot physical RDMs or independent disks.
Snapshot Manager: Delete commits snapshot to parent. Delete all commits all snapshots before you are here. Go to reverts to that snapshot. Revert to snapshot back to parent's snapshot you are here.
VMotion: To vMotion a suspended VM, new host must meet CPU compatibility requirements.
Storage vMotion: can transform thick > thin or thin > thick. Limitations: VMs cannot have snapshots, only persistent VMXDS or RDMs, requires license, ESX3's hosts need vMotion licensed/configured.
VMi (VM Interface) paravirtualization: standard to improve performance, only Linux 32bit guests.
Uses a PCI slot. VMi VMs must be off to move to an unsupported host; can reduce performance.
VMDirectPath: I/O allows guest OS to access physical PCI/PCIe devices. Intel Nehalem platforms.
Restrictions: vMotion, Hot add/remove, suspend, record/replay, FT, HA, DRs (but allowed in cluster).
SCSI controllers: • BusLogic Paravirtual • LSI Logic SAS • LSI Logic Parallel • PVSCSI
• **PVSCSI** (Paravirtual SCSI): high-performance storage adapter. Recommended for DAS. Guests: Win 2k3, 2008, RH EL5. Not supported: record/replay/FT, PVSCS, (2003/8 boot disks OK since U1)
• **NPIV** (non-Port ID virtualization): enables FC HBAs to connect multiple virtual VMs, each with unique IDs.
VMs assigned 4 WWNs. Limitations: NPIV enabled FC switch, only RDMs. HBAs need access to LUN using its WWN. NPIV capable HBAs, no Storage vMotion, VM can't power on if WWNs in use.
vmnic's • Flexible - 32-bit guests, vance without VMware Tools or vmxnet with VMware Tools • **e1000** - Emulates E1000 NIC, default for 64-bit guests • Enhanced vmxnet - vmxnet with enhanced performance, requires VMware Tools • **vmxnet3** - vmxnet with enhanced performance & networking features, requires VMware Tools & HW v7, doesn't support FT.
TSO (TCP Segmentation Offload): enabled in VMkernel by default, must be enabled at VM level.
Needs enhanced vmxnet, might change the MAC. **Jumbo frame** requires enhanced vmxnet or vmxnet3.
OVF: templates can be deployed from a local file system via the Client, or from a web server.
OVF files are compressed. Client validates the OVF file before importing it.
vApp: container containing one or more VMs, can power on or off, & be cloned. Metadata in VM's DB.
IP pool - network configuration assigned to network used by vApp. VC then provides IPs to its VMs.
Links: <http://kb.vmware.com/kb/1010448> - Set all VMs to upgrade tools at next power on
<http://kb.vmware.com/kb/1002511> - Recreate missing virtual disk (VMDK) header/description file
<http://kb.vmware.com/kb/1002310> - Committing snapshots if no snapshot entries in snapshot manager
<http://kb.vmware.com/kb/1007849> - Consolidating snapshots

Maximums (per HA cluster): Hosts = 32 VMs = 1280 (max 160 per host, but > 40 limits hosts to 8)
Failover hosts = 4 (only 5 primaries), or 50% of hosts if less than 10

FW Port	Source	Destination	Port (ESX port)	Description
2050-2250	ESX/ESXi	ESX/ESXi	HA	DCP (SD)
2050-2250 (out)	ESX/ESXi	Hosts	TCP/UDP(SC)	HA
8042-8045	Hosts	ESX/ESXi	UDP (SC)	HA
8042-8045 (out)	ESX/ESXi	Hosts	TCP/UDP(SC)	HA

Logs: HA logs: [/var/log/vmware/aam/](#)

HA primary hosts (first 5): maintain & replicate cluster state and initiate failover actions.

Active primary host: decides where to restart VMs, tracks & effects failed restart attempts.

List primary hosts: `$ cat /var/log/vmware/aam/aam_config_util.lstnodeslog`

HA primary host: primary is: "maint mode - disconnected - removed. Not on failure."

Host isolated: no heartbeat for 15 seconds, then cannot ping isolation addresses. **Isolation response:**
• power off • leave powered on • shut down (default). However **Host Failure** is only after 15 seconds.

Admission Control types: • Host • Resource Pool • HA (only HA admission control can be disabled)

HA Admission Control: rules if VMs can power on when they violate availability constraints at HA failover. Actions that change a reservation must satisfy admission control. **Control policies:** Host
Failures Cluster Tolerates (1-4 hosts) - adds Advanced Runtime Info box showing slot size, total, used, available slots, total VMs on, hosts, good hosts • % of Cluster Resources (up to 50%) • Specify a reservation

Slot size: represents VM CPU & memory resources needed for any powered on VM. Distorted by large VM reservations. Avoided with advanced attributes `ha.slotCpuInMB` or `ha.slotMemInMB`

Links: [http://www.veled-bricks.com/vmware-high-availability-depdiv](#) • HA deep dive