



Availability (FT & MSCS)

Maximums (FT advice): Disks per VM =16 FT VMs per host = 4 Minimum hosts per cluster = 3

FW Port	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, hotplugging, MSCS, VCB, SMP, physical RDMs, Paravirtualized VMs, NPV, VMDirectPath, EPT/PTV. DRS only if cluster is EVC.

MSCS: • 2003 SP2 & 2008 (Failover Clustering) • 32 & 64bit • only 2 nodes clusters

Not supported - DRS on VMs, vMotion, FT, NPV, 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	Only 2003	Yes (recommended)
Physical & VM (n+1)	No	No	Yes
Snapshots	Yes	Yes	No
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:x). Set sharing to Physical or Virtual as required.
SCSI bus sharing • CIB = Virtual • CAB or N+1 = Physical

Links: <http://kb.vmware.com/kb/1010601> - Understanding FT
<http://kb.vmware.com/kb/1008027> - CPU & guest OS that support FT



Networking

Maximums (per host): 1GB VMNICs = 2 - 32 dependent on HW 10GB VMNICs = 4
PCI VMDirectPath devices=8 switches (vSS/vDS/VEEM) = 248/16/1 vSS/vDS ports=4096
Active ports = 1016 Service Console ports = 16

Maximums (per vCenter): vDS switches=32 vDS port groups=5000(1016 ephemeral) vDS ports=20000

Maximums (per switch): Hosts (per vDS) = 350 vSS port groups = 512vSS switch ports = 4,088

Terminology: VMNICs - logical name for physical server NICs vNICs - virtual NICs assigned to VMs
vSS - virtual Standard Switch vDS - virtual Distributed Switch vdPort - port group on a vDS
dvUplink - uplink VMNICs on a vDS Network vMotion - 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-vswif -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](#)

Local host file: [etc/hosts](#) DNS servers: [etc/resolv.conf](#) DG: [etc/sysconfig/network](#)

Ethernet tagging: • EST (External Switch Tagging) - Default. No trunking required. 1-1 relationship

from VMNICs to physical switch ports. Each VMNIC can only see 1 subnet. VLAN ID of 0 or blank.

• VSI (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.

Avoid using a **VLAN ID of 1**, as this is the native Cisco VLAN ID.

vSS & 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 outbound 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, in/out on vDS.

NIC Teaming - Load Balancing (spreads outbound traffic from vNICs across VMNICs). **Originating**

port ID (default) uses VMNIC based on where traffic entered. **ip hash** based on source & destination IP

address of each packet (if physical switch ports are etherchannel). **Source MAC** hash based on source

MAC address. **Route based on physical NIC** load only on vDS, dynamically redistributes load across all

VMNICs in team. 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.

dvPort options: • Port Binding Static when initially connected **Dynamic** when connected/powered-on

Ephemeral no binding • Traffic shaping **Ingress** into vSwitch **Egress** out of vSwitch • Allow live port

moving • Config reset at disconnect • Host can assign port if vCenter is down • Name format

NIOC (Network IO Control): prioritise egress vDS traffic via shares/limits (FT,iSCSI,vMotion,Mgr,NFS,VM)

PVLAN (Private VLAN): extension to VLAN standard, adds further segmentation. 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 vmxnet vNIC.

Jumbo frames up to 9kB. Must be enabled for each vSwitch. vNIC must be vmxnet2/3 or e1000.

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

Configure networking (for vSS): (1) add vSwitch **esxcfg-vswitch -a** (2) add port group to

vSwitch **esxcfg-vswitch -A** (3) set port group's VLAN ID **esxcfg-vswitch -p -v** (4) add

VMNIC to vSwitch **esxcfg-vswitch -L** • VM connections: set VM's NIC to use port group.

• Service Console: create interface & add it to the port group **esxcfg-vswif -a -p -i -n**,

set the DG in [etc/sysconfig/network](#), then restart networking **service network restart**.

• VMkernel ports: add **esxcfg-vmknic -a -i -n** & set VMkernel DG **esxcfg-route**.

vMotion enabled in vCenter if required.

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



Resources

Maximums (per DRS cluster): Hosts = 32 VMs (powered on) = 3000 (limit of 320 per host)

Maximums (per Resource Pool): Children = 1024 Tree depth = 8

Maximums (other): Hosts per datacenter = 400 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 (can only draw

from own pool). List the resource group settings: **\$ sudo /usr/sbin/esxcfg-resgrp -l**

Shares only apply during contention. Shares are 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: VM-VM keep VMs together/apart. VM-Host keep VMs on/off specific hosts. **Should**

rule best effort. **Must rule** mandatory (for licensing). **Rule** conflicts older wins, newer rule disabled.

Anti-affinity wins over affinity. Disabled rules ignored.

Current host load standard deviation: DRS load imbalance. Current < Target unless advice unapplied

DPM: uses IPMI, ILO or WOL (in that order). DRS & DPM thresholds are independent. Verify host's

DPM **Last Time Exited Standby**. DPM level - Off, Manual (makes recommendations) & Automatic.

Links: <http://kb.vmware.com/kb/1005764> - Enhanced vMotion (EVC) FAQ

<http://www.yellow-bricks.com/hrs-deepdive/> - DRS Deep Dive

<http://kb.vmware.com/kb/1003212> - EVC CPU compatibility



Storage

Maximums (per host): Volumes = 256 Paths = 1024 NAS datastores = 64

FC - HBAs = 8 (HBA ports = 16) targets per HBA = 256 paths to each LUN = 32

iSCSI HW - HBAs = 4 targets per HBA = 64 (62 - QLogic Static) paths to each LUN = 8

iSCSI SW - NICs = 8 targets = 256 paths to each LUN = 8

Maximums (per volume): VMs = 256 Hosts = 64 (DRS 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 (VMK)	NFS Client
2049	ESX/ESXi	NFS server	TCP (VMK)	NFS Client
3260	ESX/ESXi	iSCSI server	UDP (SC+VMK)	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/esxccli mmp satp list**

List claim rules: **\$ sudo /usr/sbin/esxccli corestorage claimrule list**

Lists 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 -l**

List iSCSI LUNs: **\$ sudo /usr/sbin/vmkiiscsi-tool -L -l 1 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	FC	iSCSI	NAS
vMotion, DRS, HA, FT, VCB, SRM & Thin VMDKs	Yes	Yes	Yes
VMFS volumes, RDMs & VMware's NMP	Yes	Yes	No
Boot ESX host	Yes	Yes (HW initiator)	No
VM MSCS clustering	Yes	No	No

Zoning: at the switch. **LUN masking:** done at the SP or server.

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

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

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

LUN addressing FC: **Runtime Name** **vmhba#;Ch#;T#;L#** - adapter:channel:target:LUN

iSCSI: **IQN** **iqn.year-month-revised_domain_name:string or EUI** **ev:string**

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

Dynamic - uses "SendTargets", target responds with list. Removed targets return after HBA rescan/reset

CHAP: HW iSCSI 1-way CHAP, initiator level. SW iSCSI 1-way & mutual CHAP, initiator or target

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

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, uses 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.MaskLUN: reduce number of LUNs scanned. **Disk.MaskLUN:** convert to claim rule instead.

VMFS volumes: Large-less LUNs to create, less to manage, flexible resizing & snapshots. Small-less

contention (locking), less wasted space, different RAIDs, more flexible multipathing & disk shares.

SIOC (Storage IO Control): shares VM's disk IO across datastore's hosts. Monitors latency, adjusts

VM's host queue access. Can also enforce VM IOPS limits. Enable on datastore, set shares/limit on VM

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

Links: <http://kb.vmware.com/kb/1009553> - Lost connectivity to storage

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

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

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.



vReference.com

vSphere 4.1

by Forbes Guthrie Version 2.3.2 for v4.1
released 13 Sep 2010



ESX Install

HW requirements: • 64-bit x86 CPUs • 2GB RAM minimum • see HCL (link below)

IPv6 is not supported during the install. **Installation log:** [/var/log/esx_install.log](#)

Evaluation period (60 days) starts on first power-on even if host is licensed.

Install boot options: **F2**. Install via **Media Depot:** HTTP/ HTTPS, FTP, NFS - **askmedia** option.

PXE Boot install: (1) Install TFTP server software (2) Put **menu.c32** file in accessible place

(3) Install PXELINUX (4) Configure DHCP server (5) Copy **vmlinux** & **infrd.img** from **/isoinux** on

DVD (6) Create [/tftpboot/pxelinu.cfg](#) on TFTP server.

Install script can be: Default script (on DVD), FTP, HTTP/HTTPS, NFS, USB flash drive, local disk.

Default install scripts: • **ks-first.cfg** installs on 1st disk • **ks-first-safe.cfg** same but keeps VMFS. Root

password is "mypassword". Interactive install creates [/root/ks.cfg](#) from choices made.

Physical partitions: /boot, vmkcore & /vmfs. **esxconsole.vmdk:** /, swap, /var/log, & optional ones.

Size of /boot, vmkcore & VMFS cannot be defined/changed during Interactive install (can if Scripted).

Mount point	Format	Default	Location
/boot	ext3	1100MB	Primary physical partition
	vmkcore	110MB	Primary physical partition
/vmfs	vmfs3	fill remaining 1 st disk	Logical physical partition
/ (root)	ext3	5GB (default min, may be larger)	esxconsole.vmdk file
	swap	600MB default (max 1600MB)	esxconsole.vmdk file
/home	ext3	optional - recommended 512MB	esxconsole.vmdk file
/tmp	ext3	optional - recommended 1024MB	esxconsole.vmdk file
/usr	ext3	optional - no recommendation	esxconsole.vmdk file
/var/log	ext3	optional - recommended 2000MB	esxconsole.vmdk file



ESX Hosts

Maximums (per host): vCPUs = 512 vCPUs per physical core = 25 Logical procs (incl HT) = 128
RAM = 1TB Service Console RAM = 800MB (min=272MB) VMs = 320

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
903	Clients	ESX/ESXi	TCP (SC)	VM Console
5900-5964	?	ESX/ESXi	TCP (SC)	RFB for mgt tools (VNC)
5900-5964 (out)	Hosts	?	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 extras: 21(FTP), 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/vmkernel](#)
ESX service log: [/var/log/vmware/hostd.log](#) Syslog: [/var/log/messages](#)
VMkernel warnings: [/var/log/vmkwarning](#) VMkernel events: [/var/log/vmksmmary](#)
VC agent: [/var/log/vmware/vpx/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 -hP**
List disk/partition partitions: **\$ sudo /usr/sbin/esxcfg-dumppart -l**
Show description of VMkernel error: **\$ vmkerrcode error_code number**
Export detailed config file: **\$ sudo esxcfg-info > /tmp/esxcfg-info.txt**
Gather debugging report: **\$ sudo /usr/bin/vm-support -w /tmp**
Configure authentication settings: **\$ sudo /usr/sbin/esxcfg-auth**
Lists drivers loaded at startup: **\$ sudo /usr/sbin/esxcfg-module -l**
Set advanced options: **\$ sudo /usr/sbin/esxcfg-advcfg option -s value** (-g to get)
Update bootstrap settings: **\$ sudo /usr/sbin/esxcfg-boot** (treat with caution)
Initialization routines (resets things): **\$ sudo /usr/sbin/esxcfg-init** (treat with caution)

Internal firewall commands (iptables on Service Console):

Show all firewall settings: **\$ sudo /usr/sbin/esxcfg-firewall -q**
List the firewall name services: **\$ 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 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 & out, except 22, 123, 427, 443, 902, 5989, 5988, pings, DHCP & DNS

Master config file: [/etc/vmware/esx.conf](#) **Certificate files:** hostd regenerates 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/hostd/proxy.xml](#) SSL settings [/etc/vmware/hostd/config.xml](#)
PAM (Pluggable Authentication Modules) configuration: [/etc/pam.d/vmware-authd](#)
Default authentication method is [/etc/passwd](#), vpxuser is for vCenter Server permissions.
Passwords: ESX uses [pam_cracklib.so](#) plug-in by default. No restrictions on 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 - CC = E. [pam_passwdqc.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. **HT:** can help better utilize idle resources.

Reclaims VM memory by: • Balloon driver (vmmemctl) force guest to use native algorithms (guest swap) • Memory compression • vswp file (if vmmemctl unresponsive) • Sharing memory across VMs
VMware MIBs: uses embedded SNMP agent (disabled by default). Enable: **vicfg-snmpp 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/1005184](#) - Decoding Machine Check Exception output after purple screen
[http://kb.vmware.com/kb/1012514](#) - Determining detailed build number information for ESX hosts
[http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.0.pdf](#) - Performance best practices
[http://communities.vmware.com/docs/DOC-9279](#) - Interpreting esxtop Statistics



ESXi hosts

HW requirements: 64bit x86 CPUs, 2GB RAM, SATA, SAS or SCSI disks. No ESXi WebAccess.
ESXi Installable parts in eval 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-snapshot), 110MB diagnostic for core dumps, VMFS3 on free space.
Not supported: • ESXi Installable & Embedded on same host • Booting multiple servers from 1 image
DCUI (Direct Console UI): • Configuring host defaults • Set up administrative access • Troubleshooting
Restarting Mgt agents effects [/etc/init.d](#) processes: hostd (mgmt-vmware), npd (time), sfcbd (CIM broker), slpd (discover/advertise services), wsman (share mgt info via SOAP), vobd (error reporting) & AAM (HA agent) if installed. To isolate ESXi host from DRS/HA cluster disable mgt network.
Management Network Test: pings DG, primary DNS nameserver, secondary DNS, resolves hostname.
TSM (Tech Support Mode): busybox console now fully supported, remote connection via SSH.
Lockdown mode: DCUI restricted to root user, TSM disabled for all users, vSphere client and CIM monitoring only via vCenter not direct to host.
vicfg-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.



vCenter

Maximums (per vCenter): Hosts = 1000 VMs = 15000 Running VMs = 10000 Clients = 100
Maximums (Linked mode): vCenters = 10 VMs = 50000 Running VMs = 30000 Hosts = 3000
Maximums (per host): Provisioning ops = 4 vMotions = 4 (8 for 10Gbps) Storage vMotions = 2
Maximums (per datastore): Hosts = 400 Provisioning ops = 4 vMotions = 128 Storage vMotions = 8
HW requirements: Min - 2 CPU cores, 3GB RAM • Medium (50 hosts, 500 VMs) 2 cores, 4GB RAM
• Large (300 hosts, 3000 VMs) 4 cores, 8GB RAM • Extra large (1000 hosts 10000 VMs) 8 cores 16GB
SW requirements: • Only 64bit Windows now • 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 64bit DBSN. User needs DBO 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	VC	Hosts	UDP	Heartbeat
903	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(web services), 8181/8182(collector service), 27000/27010(license 3.x hosts).

Logs: DB upgrade: **%TEMP%\VCDatabaseUpgrade.log** VC agent: [/var/log/vmware/vpx/vpxa.log](#)
VC install: **%TEMP%\VC** directory of user installing VC VC logs: **%TEMP%\vpx/vpxd-#.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, 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.

Permissions: Assigning - pair user/group with role & associate with object. Role - predefined set of privileges. Users initially granted No Access role on new objects, including datastores/networks.
Logged 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 & vpxuser 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.

licenses	ESXi Single	Essential+	Standard	Advanced	Enterprise	Enterprise+
vCenter	No	Essentials	Foundation	Standard	Standard	Standard
Cores per socket	6	6	6	12	6	12
vSMP	4-way	4-way	4-way	4-way	4-way	8-way
Physical RAM	256GB	256GB	256GB	256GB	256GB	no limit
vpxa, Up Mgr, vStorage	Yes	Yes	Yes	Yes	Yes	Yes
HA, vMotion	Yes	Yes	Yes	Yes	Yes	Yes
Data Recovery	Yes	Yes	Yes	Yes	Yes	Yes
Thin Provisioning	Yes	Yes	Yes	Yes	Yes	Yes
Hot Add, FT, vShield Zones, Serial Port Concentrator	Yes	Yes	Yes	Yes	Yes	Yes
Storage vMotion, DRS, DPM, VAAI	Yes	Yes	Yes	Yes	Yes	Yes
vDS, Host Profiles, 3rd party MMP, SIOC, NIOC	Yes	Yes	Yes	Yes	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). **Datastore 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 domain user 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, (tcp/s ports, client timeouts, logging detail, DB connections, DB retention, SSL host verification, advanced settings).

Links: [http://kb.vmware.com/kb/1011641](#) - Collecting diagnostic information for vCenter
[http://kb.vmware.com/kb/1022101](#) - Installing ESX 4.1 & vCenter 4.1 best practices
[http://kb.vmware.com/kb/1022104](#) - Upgrading to ESX 4.1 & vCenter 4.1 best practices
[http://kb.vmware.com/kb/1005593](#) - sysprep file locations and versions
[http://kb.vmware.com/kb/1010579](#) - Comparison of vSphere 4.0 & V1 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



VMs & vApps

Maximums (per VM): vCPUs = 8 RAM = 255GB(16GB - FT VMs) Swap file = 255GB (1 per VM)
SCSI adapters = 4 Devices per SCSI adapter = 15 IDE devices (disk or CD) = 4
vNICs = 10 USB devices = 20 Floppy drives = 2 Parallel ports = 3 Serial ports = 4
Remote consoles = 40 VMDirectPath devices = 4 VMDirectPath SCSI targets = 60

Files:	Earlier version of .vmx file	.vmem	VM's memory
.dsk	Earlier version of .vmdk file	.vmnsd	Snapshot metadata
.hlog	vMotion log file	.vmns	Snapshot state file
Jck-XXX	Locking file on NFS datastore	.vmss	Suspended state file
.log	VM activity log	.vmt	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	.vmxf	Extra configuration file for VMs in a team
.REDO	Earlier version of -delta.vmdk file	.vswp	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	00000#.vmdk	Snapshot metadata
-ctk.vmdk	Changed Block Tracking file	00000#-delta.vmdk	Snapshot differential file

Logs: VM log [/vmfs/volumes/datastore_name/vm_name/vmware.log](#)

Commands: List all registered VMs on a host: **\$ sudo /usr/bin/vmware-cmd -l**
Create/modify VMDKs, RDMs, VMFS volumes & storage devices: **vmkfstools** (check man page)
Power Off = hard power off • Shut Down = soft with VMware tools • Reset = hard • Restart = soft

VM HW: Memory/CPU Hotplug - VMware Tools must be installed.
VMs with HW earlier than v4 have reduced performance & capabilities. Cannot add/remove devices.
Manually MAC addresses: 00:50:56:x.y.z. Set in vmx: ethernet.number=addressType="static".
Disk types: zeroedthick (lazy) default, pre-allocates. eagerzeroedthick 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 mgt 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 VMDKs or RDMs, requires license, ESX3.5 hosts need vMotion licensed/configured.
VM I/O (VM Interface) paravirtualization: standard to improve performance, only Linux 32bit guests. Uses a PCI slot. VM I/O 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 Parallel • LSI Logic SAS • LSI Logic Parallel • PVSCSI
PVSCSI (Paravirtualized SCSI): high-performance storage adapter. Not recommended for DAS. Guests: Win 2003, 2008, RHEL5. Not supported: Record/Replay, FT, MSCS, (2003/8 boot disks OK since U1)
NPIV (N-port ID virtualization): share FC HBA port as multiple virtual ports, each with unique IDs. VMs assigned 4 WWNs. Limitations: NPIV enabled FC switch, only RDMs, HBAs need access to LUN using its WWNs, NPIV capable HBAs, no Storage vMotion, VM can't power on if WWNs in use
vNICs: • Flexible - 32-bit guests, vance without VMware Tools • vmxnet3 - vmxnet with enhanced performance & networking features, requires VMware Tools & HW v7.
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 vmxnet2/3 or e1000.
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 & off, & be cloned. Metadata in VC'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/1010048](#) - 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



Availability (HA)

Maximums (per HA cluster): Hosts = 32 VMs = 3000
Failover hosts = 4 (only 5 primaries), or 50% of hosts if less than 8
FW Port **Source** **Destination** **Prot (ESX port)** **Description**
2050-2250 Hosts ESX/ESXi UDP (SC) HA
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_listnodes.log**
Secondary host promoted if primary is: • maint mode • disconnected • removed. Not on failure.
Host isolated: no heartbeat for 12 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 Failover Host. **Policy Factors** • resource fragmentation • flexibility • VM diversity.
Slot size: represents VM CPU & memory resources needed for any powered on VM. Distorted by large VM reservations. Avoided with advanced attributes [das.slotCpuInMHz](#) or [das.slotMemInMB](#)
Links: [http://www.yellow-bricks.com/vmware-high-availability-deepdiv](#) - HA deep dive