



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, DRS features, hotplugging, MSCS, VCB, SMP, physical RDMs, Paravirtualized VMs, NPV, VMDirectPath, EPT/RVI.

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

	VMMDK	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	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 (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 - virtual Standard Switch vDS - virtual Distributed Switch vPort - 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.
 • 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.
 Avoid using a **VLAN ID of 1**, as this is the native Cisco VLAN ID.

vSS & vDS options (options can and should 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 10240 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 (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 **List 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

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. 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 VM's NIC to use the 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 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



Resources

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 organisation & 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 keep VMs together or apart in a DRS cluster. Anti-affinity rule limited to 2. 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, IL0 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. Hosts reclaim memory from VMs by: • Balloon driver (vmemctl) force guest to use native algorithms (guest swap) • VM Swap files (if vmemctl 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



Storage

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 (DRS limit, 2048 for linked clones)
 VMFS = 64TB (less 64KB) NFS = 512TB 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/vmkping**
 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/vmkiscsi-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 VMMDKs	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, other 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#;C#;T#;L# - adapter:channel:target:LUN
 iSCSI: **IQN** iqn.year-month-reversed_domain_name:string or EUI eui.8string

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.](http://etc/vmware/esx.conf)

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 (8 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 format.

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.

TPGS (Target Port Group Support): storage shows path performances, so hosts can determine best path.

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)
<http://kb.vmware.com/kb/1011387> - ESX 4.x 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.



vReference.com

vSphere 4.0

by Forbes Guthrie Version 2.0 for v4.0U1 released 10 Dec 2009



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](http://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 **vmlinuz** & **initrd.img** from **/isolinux** 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 & /vms. **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
/vms	vmfs3	fill remaining 1 st disk	Logical physical partition
/ (root)	ext3	5GB (default min, may be larger)	esxconsole.vmdk file
/home	swap	600MB default (max 1600MB)	esxconsole.vmdk file
/tmp	ext3	optional - recommended 512MB	esxconsole.vmdk file
/usr	ext3	optional - recommended 1024MB	esxconsole.vmdk file
/var/log	ext3	optional - no recommendation	esxconsole.vmdk file
	ext3	optional - recommended 2000MB	esxconsole.vmdk file

vReference recommends: **/home, /opt, /tmp** - min 2GB each, **/var** (no **/var/log**) - 5GB, swap - 1600MB
 If scripting install, consider one VMFS for COS (esxconsole.vmdk) and a separate one for VMs. Disconnect Fibre Channel connections prior to installation.

Post install tasks: • Reconnect FC connections.

• Create user account & add to sudoer file (**visudo** - add to "user privilege specification").

• Test cables are in correct VMNICs: **\$ watch -n 1 'sudo /usr/sbin/esxcfg-nics -l'**

• Rearrange VMNICs in etc/vmware/esx.conf if required (reboot required).

• Adjust Service Console memory to 800MB (reboot required).

• Configure NTP (time) settings.

• Patch (VUM or **vihostupdate/esxupdate**).

• Connect vSphere Client to host (not VC) & add extra users (the sudo users) to Administrators group.

• Configure vSwitches.

• Configure storage (& set DiskMaxLUN as required).

• Connect vSphere Client to VC, add new host, move to required cluster.

• License host.

• Enable Web access if required.

Upgrade from ESX3: (cannot use DVD) • **VUM** (vCenter Update Manager) - upgrades ESX/ESXi.

• **Host Update Utility** - upgrades ESX/ESXi (& updates ESXi), small environments (< 10 hosts, no VUM). Customize in etc/PROGFILE/ESX/VMware/Infrastructure/VU/Update 4.0/settings.config

• **esxupgrade.sh** script <http://kb.vmware.com/kb/1009440> - upgrades ESX only.

Upgrade logs: esx3-installation/esx4-upgrade/ & var/log/vmware/

Unsuccessful upgrades: esx4-upgrade/ & var/log/vmware/

Post upgrade: • Upgrade VMware Tools before installing virtual hardware • Re-install 3rd party agents/apps • Convert LUN masking to claim rule format: **esxccli corestorage claimrule convert** • Successful upgrade: **cleanup-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/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/1010675> - Upgrading an ESX 3.x VM to ESX 4.0

<http://kb.vmware.com/kb/1011712> - See if Intel VT or AMD-V is BIOS enabled without rebooting



Clients

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, WebAccess	ESX	TCP	SSH access
80	WebAccess	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](http://var/log/vmware/vpx/vpxa.log) Client Install log %TEMP%/vmsmi.log

Client Service log C:\Docs and Settings\username\Local Settings\AppData\Local\vmtoolsd\client-x.log (x=0-9)

Web Access to ESX or VC: <https://hostname.domain.com/ui> • ESX - 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.jume.nl/articles/vmware/143-vcenter-client-shortcuts> - vCenter Client shortcuts



ESX Hosts

Maximums (per host): vCPUs = 512 vCPUs per physical core = 25 Logical procs (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
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](#)

VMkernel warnings: [/var/log/vmkwarning](#)

VC agent: [/var/log/vmware/vpx/vpxa.log](#) Syslog: [/var/log/messages](#)

VMkernel events: [/var/log/vmksuammary](#) 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 diagnostic 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:** host 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.

Vmware MIBs: uses embedded SNMP agent (disabled by default). Enable: **vicfg-snmp**

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://kb.vmware.com/kb/1000258](#) - Video: Configure Service Console networking from CLI

[http://kb.vmware.com/kb/1991](#) / 1992 - VMotion compatibility for Intel / AMD processors

[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 starts 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

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 administration 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.

Lockdown mode: prevents remote access by root account, but not other accounts. Disabled by default.

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.

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



vCenter

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

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\SYSTEM32\odbcad32.exe](#). 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(SMTPL),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%** 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 editions ----

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

Thin provisioning Yes Yes Yes Yes Yes Yes

vpxa, i/o Mgr, VMsafe, vStorage HA Yes Yes Yes Yes Yes Yes

Data Recovery Yes Yes Yes Yes Yes Yes

Hot Add, FT, vShield, VMotion Yes Yes Yes Yes

Storage VMotion, DRS Yes Yes Yes Yes

vDS, Host Profiles, 3rd party MMP 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, http(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/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](#) - 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 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 .vmx file	.vmem	VM's memory
	.dsk	Earlier version of .vmdk file	.vmnsd	Snapshot metadata
	.hlog	VMotion log file	.vmns	Snapshot state file
	.lck-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/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 (VM Interface) paravirtualization: standard to improve performance, only Linux 32bit guests.