

Networking

VMNICs are logical names for physical NICs on server. vNICs are virtual NICs assigned to VMs.

Maximums: vNICs per VM = 4 vNICs per vSwitch = 1016 vNICs per host = 4096

To use more than 256 vNICs per host, you must manually set MAC addresses in vmx files.

vSwitches = 127 (ESX 3.0 was 248) Port Groups = 4096 (ESX 3.0 was 512)

e100 VMNICs = 26 e1000 VMNICs = 32 Broadcom VMNICs = 20

Ethernet tagging options:

EST (External Switch Tagging) - Default. No trunking required. One-to-one relationship from the VMNICs to the physical switch ports. Each VMNIC can only see one subnet. VLAN ID of 0 or blank.

VST (Virtual Switch Tagging) - Commonly used. The VMNICs connected to a vSwitch are able to span several VLANs. Each Port Group on the vSwitch is assigned a VLAN ID of 1-4094. If you have to use the native VLAN for VMs, leave the VLAN ID blank.

VGT (Virtual Guest Tagging) - Rarely used. Install 802.1Q trunking driver software in the VMs, and the vSwitch preserve the tags given by the VMs. VLAN ID of 4095. VGT only with e1000 driver.

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

vSwitch config options (these options can also be overridden on individual Port Groups):

- General** • Number of ports. 56 by default when created in VIC, 128 when created on Service Console.
- vSwitch0 is created with 56 during the install (24 ports in ESX 3.0) - not an option on Port Groups.
- 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 drops outbound frames if source MAC address is different.

Traffic Shaping (no longer defined per VM) • Status (default Disabled) **Average Bandwidth** (default 102400 Kbps) **Peak Bandwidth** (default 102400 Kbps) **Burst size** (default 102400 KB)

NIC Teaming • Load Balancing (spreads outbound traffic from vNICs across VMNICs) - Originating virtual port ID (default) uses VMNIC based on where traffic entered - ip hash based on source & destination IP address of each packet. Use this when physical switch ports are in a team/bonded group (link aggregation) - Source MAC hash based on source MAC address - Use explicit failover order

- Network Failover Detection **Link status only** (default) detects cable pulls and switch power failures, not misconfigurations. **Beacon Probing**
- Notify Switches - No or Yes (default) updates lookup tables.
- Rolling Failover - No (default - it will return to original) or Yes. In ESX 3.5 this has been renamed to Failback - if this is set to Yes it will failback to original.
- Failover order - **Active Adapters** - **Standby Adapters** - **Unused Adapters**

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-vmknlc -l`

List VMkernel Default Gateway: `$ sudo /usr/sbin/esxcfg-route`

Must add a vSwitch (**esxcfg-vswitch -a**), then add a port group to it (**esxcfg-vswitch -A**), set the port group's VLAN ID (**esxcfg-vswitch -p -v**), and 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 and 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-vmknlc -a -i -n**) and set the VMkernel DG (**esxcfg-route**). VMotion should be enabled in VC if required.

Common networking configuration files: Name resolution order: **/etc/nsswitch.conf**

Local host file: **/etc/hosts** DNS servers: **/etc/resolv.conf** DG: **/etc/sysconfig/network**

CDP default: off in ESX 3.0, listen in ESX 3.5.

To enable full CDP: `$ sudo /usr/sbin/esxcfg-vswitch -B both vswitch_name`



MSCS Clustering

Maximums: RDM size = 2TB VMDK size = 2TB


- Configure all RDMs before configuring network settings (RDM disks need a second SCSI controller and then this bumps the NIC's PCI slot, which stuffs any IP settings you make beforehand).
- Add all RDMs to a 2nd SCSI controller i.e. SCSI(1:x). Set controller to Physical or Virtual as required
- Ensure that all VM nodes have the RDMs setup before initialising any of the LUNS within windows.
- Windows Time Service must be enable in guest (set HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Parameters\Type (REG_SZ) to "NoSync" so the VM can still use host syncing).
- When adding the second node to the cluster, select Advanced and choose "Advanced (minimum) configuration", to prevent the wizard failing when it checks the shared disks.
- Cannot use VMotion with MSCS VMs.

VM's SCSI bus sharing setting: CIB - Virtual CAB or N+1 - Physical

Requirements for supported configuration (not hard rules):

- VM's OS on Direct Attached Storage of ESX host
- Cannot be part of DRS or HA cluster
- Only 32bit Windows (64bit from ESX 3.5 update 1)
- No mixed HBAs
- Shared disks must use 2GB FC (not iSCSI or NFS)
- Only miniport SCSI driver (not STORPort)
- No boot from SAN for ESX host (can with ESX 3.5 update 1)

	VMDK	Virtual RDM	Physical RDM
Cluster in a box	Yes (disk must be zeroed)	not recommended	not supported
Cluster across boxes	No	Yes	not recommended
Physical and VM	No	No	Yes
Snapshots	Yes	Yes	No
SCSI target software	No	No	Yes



ESX Hosts

System logs: ESX 2.x service log /var/log/vmware/vmware-serverd.log

ESX 3.x service log /var/log/vmware/hostd.log

Service Console log /var/log/messages

VMkernel messages /var/log/vmkernel

VMkernel warnings /var/log/vmkwarning

VMkernel events /var/log/vmksummary

VC agent log /var/log/vmware/vpx/vpxa.log

VI Client log %TEMP%\%vclient-x.log

VM log /vms/volumes/datastore_name/vm_name/vmware.log

Show description of a VMkernel error (only ESX 3.0.2 +): `$ vmkerrcode error_code_number`

Export a **detailed configuration file**:

`$ sudo /usr/sbin/esxcfg-info > /tmp/esxcfg-info.txt`

Gather **debugging report**: `$ sudo vm-support -w /tmp (-h for switches)`.

After COS changes, **refresh VI Client**: `$ sudo /sbin/service mgmt-vmware restart`

To see the status of all services: `$ sudo /sbin/service --status-all`

Restart a service: `$ sudo /sbin/service service_name restart` (start, stop, status available)

Boot process: 1) Bootloader (Normal/Debug VMkernel/Service console) - settings in **/etc/grub.conf**

2) initrd - initial RAM disk (loads VMkernel, device drivers and mounts **/root** & **/proc**).

3) VMkernel loads 4) vmmix (Service Console).

5) **/sbin/init** which runs **/etc/initab** (specifies which services run at which runlevel).

6) init script for the runlevel (**/etc/rc.d/rc3.d** for normal ESX boot) - runs scripts starting 'S' in order.

To list the service **runlevels**: `$ /sbin/chkconfig --list`

To check the **filesystem's usage**: `$ vdf -h`

Internal firewall (iptables on the Service Console)

Show all the firewall setting: `$ sudo /usr/sbin/esxcfg-firewall -q`

List the firewall named 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`

To close a port: `$ sudo /usr/sbin/esxcfg-firewall -c port, protocol, direction`

External firewall ports (from the ESX host's perspective)

Port	Incoming	Outgoing	Via	Description
80	TCP		Service Console & 3i	HTTP access: Web Access & VM console
427	UDP	UDP	3i	3i Service Location Protocol (CIM client)
443	TCP		Service Console & 3i	HTTP access (cannot change)
902	TCP	UDP	Service Console & 3i	Authentication traffic (cannot change)
903	TCP		Service Console	Remote Console traffic (cannot change)
2049	TCP	TCP	VMKernel	From NFS device
2050-5000	TCP, UDP	UDP	Service Console & 3i	HA & Autostart (3i uses just 2050-2250)
3260	TCP		Both	iSCSI
5900-5906	TCP	TCP	3i	RFB protocol for mgnt tools (e.g. VNC)
5988	TCP	TCP	3i	CIM server transactions over HTTPS
5989	TCP	TCP	3i	CIM server transactions over HTTP
8000	TCP	TCP	VMKernel	Requests from VMotion
8042-8045	TCP, UDP	UDP	Service Console & 3i	HA & EMC Autostart Mgr
8083			VirtualCenter diagnostic port	
8085, 8087 & 9080				Used internally
8086			VirtualCenter apache tomcat	
27000	TCP		Service Console & 3i	To License Server
27010	TCP		Service Console & 3i	From License Server

Rename an ESX hostname (safest way, recommended by VMware)

- 1) In VC, put host into maintenance mode (and manually migrate off VMs if required).
- 2) Release the license (if VC based licensed).
- 3) Under the DNS and Routing section, change the name of the host.
- 4) Remove the host from VC.
- 5) Login direct to the server with the VI client or via SSH and reboot the server.
- 6) Change Host (A) record on DNS servers to reflect name change.
- 7) After the reboot, add the host to VC with the new name.
- 8) Reconfigure the licensing settings and exit maintenance mode.
- 9) Migrate VMs back (ensure VMotion and HA is working as required).
- 10) If necessary, rename hardware remote management tool (iLO, RSA, DRAC).
- 11) Check that the change has taken affect:

`$ hostname` and `$ cat /etc/hosts` `$ cat /etc/sysconfig/network`

Changing the Service Console's IP address (at console or with remote management card)

`$ sudo /usr/sbin/esxcfg-vswif -i ip_address -n subnet_mask vswif0`

`$ sudo /sbin/service mgmt-vmware restart`

Edit the gateway and hostname in **/etc/sysconfig/network** and the ip address in **/etc/hosts**.

`$ sudo /sbin/service network restart`

If you are using VST (VLANing) on the the Service Console, you also need to run:

`$ sudo /usr/sbin/esxcfg-vswitch -p PortGroupName -v VLANid vswitch0`

Patching: 1) Copy the patch to the server 2) Extract the file: `$ tar -xvzf patching.tgz`

- 3) Change to the newly created directory: `$ cd patch`
- 4) Install: `$ sudo /usr/sbin/esxupdate update (-n to prevent a reboot)`

Patching logs: **/var/log/vmware/esxupdate.log**

Mounting USB keys: Run `$ sudo /sbin/fdisk -l` before plugging in the key and then run once after its plugged in. The new partition listed will give you the device name.


Create a directory `$ sudo mkdir /mnt/usbkey`

Mount the key `$ sudo mount /dev/device_name /mnt/usbkey`

Before removing, unmount the key `$ sudo umount /mnt/usbkey` (umount not unmount)

Master config file: **/etc/vmware/esx.conf**

Set advanced options: `$ sudo /usr/sbin/esxcfg-advcfg option -s value (-g to get)`



Storage

Maximums: Hosts per virtual cluster = 32 Volumes per host = 256 Extents per volume = 32

RDMs = 2TB Hosts to a VMFS volume = 32

VMFS2 volume (32 extents) = 64TB File size = 64TB Files per volume = 256 + (64 x extents)

VMFS3 volume = 64TB File size = 2TB Files per volume = more than supported

LUNs per server = 256 LUN paths = 32 Devices per SCSI controller = 16

FC HBAs per host (FC) = 16 iSCSI hardware HBAs per host = 2

FC targets per HBA = 15 iSCSI targets per HBA = 64

NAS datastores = 8 (32 with adv settings)

Storage capabilities	FC	iSCSI	NAS
Boot ESX host	Yes	Yes (HW initiator)	No
VMotion, DRS & HA	Yes	Yes	Yes
VMFS volumes	Yes	Yes	No
RDMs	Yes	Yes	No
VM MSCS clustering	Yes	No	No
VCB	Yes	Yes (No in ESX 3.0)	Yes (No in ESX 3.0)

LUN addressing FC SAN: `vmhbaadapterID:targetID:LUN:partition`

iSCSI: `IQN iqn.year-mo-reversed.domain_name:string or EUI cui.string`

A VMkernel Port Group connection is required to use iSCSI or NFS storage.

A Service Console connection is required for iSCSI, even if CHAP authentication is not used.

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

Test VMkernel connectivity: `$ /usr/sbin/vmknping`

Lists datastores, dev names to VMFS: `$ sudo /usr/sbin/esxcfg-vmhba devs -m`

List LUNs and paths: `$ sudo /usr/sbin/esxcfg-mpath -l`

Software iSCSI adapter settings: `$ sudo /usr/sbin/esxcfg-swiscsi -q`

List iSCSI LUNs: `$ sudo /usr/sbin/vmkiiscsi-tool -L -l -1 adapter`

Rescans for iSCSI LUNs: `$ sudo /usr/sbin/esxcfg-rescan adapter`

Open the SC port for iSCSI: `$ sudo /sbin/esxcfg-firewall -e swISCSIclient`

List the NFS exports from the VMkernel: `$ sudo /usr/sbin/esxcfg-nas -l`

iSCSI discovery methods: Dynamic - initiators uses "SendTargets", and target responds with a list.

Static - can manually add/remove items, only with hardware initiators.

SAN multipathing policies: Fixed - default for active/active storage devices.


MRU (Most Recently Used) - default for active/passive (& iSCSI), doesn't revert back to preferred path

RR (Round Robin) - ESX 3.0 set on SC `esxcfg-mpath, ESX 3.5 set in VIC` - load balances paths.

Disk.MaxLUN setting: reduce number of LUN scanned. **Disk.MaskLUN** setting: Hide specific LUNs.

Large VMFS volumes = less LUNs to create, flexible for resizing & snapshots, fewer LUNs to manage. **Small volumes** = less contention due to locking, less wasted space, different RAID settings, more flexible for multipathing and disk shares per LUN.

SAN System Design & Deployment Guide: http://www.vmware.com/pdf/vi3_san_design_deploy.pdf



Resources

Maximums: Hosts per DRS cluster = 32 Hosts per HA cluster = 32 (16 in ESX3.0)

Children per resource pool = 256 Resource pools per host = 512 (per cluster = 128)

Tree depth per resource pool = 12 Tree depth per resource pool in DRS cluster = 10

Datacenters mark organisational and VMotion boundaries. **Clusters** gather host CPU and memory resources for central management. **Resource Pools** apply policies to clusters across hosts. Every DRS cluster is also implicitly a resource pool.

Resource pools: • Shares - low, medium and high (1,2,4) • Reservations (minima) MHz (CPU) or MB (RAM) • Limits - MHz or MB • Expandable reservation - yes (can draw from parent's pool) - no (can only draw from own pool). Shares only apply during contention. Shares are relative to siblings.

Reservations are only checked when a VM is powered on.

Expandable reservations do not automatically hunt upwards. It never allows limits to be exceeded.

List the resource group settings: `$ sudo /usr/sbin/esxcfg-resgrp -l`

Child pools actively reserve resources from a parent whether or not VMs in child pool are powered on. To use hierarchical resource pools in a cluster, you must have DRS enabled.

DRS cluster settings: **manual** **partial** **fully automatic**

Initial VM placement manual automatic automatic

Dynamic balancing manual manual automatic

Affinity Rules determine whether to try to keep VMs together or apart in a DRS cluster.

Resource pools are prefixed **"Grafted from"** when adding a host to a DRS cluster and keeping the host's resource pool hierarchy.

HA logs: /var/log/vmware/aam (ESX 3.0 HA logs were in /opt/LGTOaam512/log)




Admission Control - rules if VMs can power on when they would availability constraints at HA failover. Actions that change a reservation must satisfy admission control.

A host put into **maintenance mode** is only cleared of VMs if it is in a fully automated DRS cluster.

Experimental VM HA in ESX 3.5 - http://www.vmware.com/pdf/vi3_35_25_vhma.pdf

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.






vmreference

VI3 card

by Forbes Guthrie Version 1.3 for 3.5U3

www.vmreference.com



ESX Install

ESX host requires min of two 1.5GHz CPU, 1GB RAM. Installing ESX Server on an IDE/ATA is supported, but not for storage of VMs. VMs must be stored on VMFS partitions on a SCSI drive or a SAN. SATA drives are not supported in ESX 3.0, but are supported in ESX 3.5.

The ESX3 installer only supports a maximum of 128 FC SAN LUNs (256 supported once installed).

Maximums: vCPUs per server = 192 vCPUs per core = 20 Cores/logical procs (incl HT) per host = 32

ESX host RAM = 256GB (ESX 3.0 is 64GB) RAM allocated to SC = 800MB (default 272MB)

Boot from SAN: • Shouldn't be used with MSCS in the VMs (works, but risk of I/O contention).

- Can be used with RDMs (couldn't with ESX 2.x).

CPU compatibility tool is on the ESX Install CDROM: **/images/cpuid.iso**

Pre-upgrade script ESX2 to ESX3 is on the Install CDROM: **/scripts/preupgrade.pl**

Disconnect all FC connections prior to installation.

Mount point	Format	Primary	Default	vmreference recommendation
/boot	ext3	yes	100MB (50MB in ESX 3.0)	250MB
/	swap	yes	544MB	1600MB
/	ext3	yes	5GB (2560MB in ESX 3.0)	min 5GB
/vmfs	vmfs3	no	1.2GB (1.1GB+ in ESX 3.0)	Don't create ¹
/home	ext3	no	512MB(optional)	512MB
/tmp	ext3	no	1024MB(optional)	min 2GB
/var/log	ext3	no	2000MB(optional)	min 2GB use /var
/opt	ext3	no	n/a	min 2GB
	vmkcore	no	100MB	100MB (last on disk)

¹ Use VI Client or Web Access to set up your VMFS3 partitions rather than the ESX Server installer. This ensures the starting sectors of partitions are 64K aligned, which improves storage performance.

Installation log: **/root/install.log** is a complete log of the installation.

kickstart scripts: enable in .xml file (see p94 of http://www.vmware.com/pdf/vi3_301_201_installation_guide.pdf for changes required) and create from webaccess link.

Post install tasks: • Reconnect FC connections.

- Create user account and add to sudoer file (**visudo** - add to "user privilege specification").
- Patch (see ESX Hosts section) .
- Install hardware management agents as required:

VMware overview: http://www.vmware.com/support/esxp25/doc/sys_mgmt_links.html

VMware latest technical note: http://www.vmware.com/pdf/esx302_cfg_mgmt_tools.pdf

Dell: http://www.dell.com/downloads/global/solutions/installing_dell_openmanage_on_esx.pdf


HP: installation instructions are contained in the README file within the agent's tgz file.

IBM: <http://wiki01.haw.ibm.com/collaboration/wiki/display/redwiki/ESX+Server+installation>

- Install backup agents as required.
- Configure NTP settings (can now configure this via VC2.5).
- add timeserver IP addresses to **/etc/ntp.conf**, add timeserver hostnames to **/etc/ntp/step-tickers** and add timeservers to **/etc/hosts**

```
$ sudo /sbin/chkconfig --level 345 ntpd on
$ sudo /usr/sbin/esxcfg-firewall --enableService ntpClient
$ sudo /sbin/service ntpd restart check with $ date
$ sudo /sbin/hwclock --systohc
```

- 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 if changes made) .
- Connect VIC to host (not VC) and add extra user to Administrators group.
- Configure vSwitches.
- Configure storage (and set DiskMaxLUN as required).
- Connect VIC to VC and add new host, move to required cluster.
- License host (host based licenses should be copied to **/etc/vmware/vmware.lic**).
- Test Web access.



ESXi

Unsupported: • Adv networking (TSO, NetQueue, limited Jumbo Frames, CDP listening) • Infiniband. HA is supported from ESXi 3.5 update 1.

Configurable via the DCUI (Direct Console User Interface): • Root password. Lockdown mode. Management network. Keyboard • view Support info and System logs • restart Management agents

RCLI commands (Remote Command Line Interface): most **esxcfg-*** are aliased to **vicfg-***

Maintenance & patching: `# vicfg-dumpart` top monitoring: `# vicfg-resxstop`

Backup/restore 3i config: `# vicfg-cfgbackup` Set NTP servers: `# vicfg-ntp`


Set remote syslog server: `# vicfg-syslog` Configure SNMP: `# vicfg-snmp`

Install updates: `# vihostupdates` Manage local & datastore files: `# vifs`

vicfg-* commands require `--server=x --username=x (& --vhost=x if --server=VC)`

Access to **Tech Support Mode**: • login to DCUI • Alt+F1 • **unsupported** • enter root password

Return to the DCUI: `# exit` and Alt+F2



VMs

Maximums: Registered VMs per host = 200 Powered-on VMs per host = 192 (128 in ESX 3.0)

CPUs = 4 RAM = 64GB (16GB in ESX3.0) NICs = 4 Devices per SCSI controller = 15

Floppy drives = 2 IDE devices (CD) = 4 Parallel ports = 3 (2 in ESX 3.0)

Serial ports = 4 (2 in ESX 3.0) Remote consoles = 10 Snapshots = 32

VM files

.cfg	Earlier version of .vmx file
.dsk	Earlier version of .vmdk file
.hlog	VMotion log file
.lck-XXX	Locking file used on NFS based datastore
.log	A log of VM activity can be useful in troubleshooting
.nvram	BIOS settings
.raw	Raw device such as a tape device
.rdm	Raw Disk Mapping in Virtual Compatibility mode
.rdmp	Raw Disk Mapping in Physical Compatibility mode
.REDO	Earlier version of -delta.vmdk file
.std	Earlier version of .vmsx file
.vmdk	Disk descriptor (also raw virtual disk for hosted products)
-flat.vmdk	Raw virtual disks
00000#.vmdk	Snapshot metadata
00000#-delta.vmdk	Snapshot differential file
.vmem	VM's memory
.vmsd	Metadata and information about snapshots
.vmns	Snapshot state file
.vmsx	Suspended state file
.vmtid	VC template (no longer used from ESX 3.0)
.vmtm	Team data
.vmtx	VC template header
.vmx	Primary configuration file
.vmxf	Supplemental configuration file for VMs in a team
.vswp	Swap file allowing memory over commitment

Non-VMware files

	Disk	Config	Suspended State
MS Virt PC/Server	.vhd	.vmc	.vsv
Xen	.img or .qcow1	.hvm	
VirtualBox	.vdi	.xml	
Parallels	.hdd	.pvs	.sav

1 If file based, can also use physical partitions, LVM volumes or an NFS root

Domain Controllers: Normally, VMs use the VMware tools to sync guest time with the host, so disable the "Windows Time Service". However, DCs need the Windows Time Service to be running, so that they can be authoritative for the domain. Set the service to Automatic and change HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Parameters\Type (REG_SZ) to "NoSync" so the VM can still use host syncing.

How to grow VM disks The easiest option here is to add additional disks, this can be done on the fly. However it will force a new drive letter in a Windows guest. To expand an existing disk:

- 1) Commit or remove all snapshots. 2) Turn off the VM.
- 3) SSH to the server and cd to the directory with the VM's files (**/vms/volumes/vm_name/vm_name/**).
- 4) `$ sudo /usr/sbin/vmkfstools -X new_sizeG vm_name.vmdk`

`vm_name.vmdk` is the disk descriptor file. That is an uppercase X.


VC 2.5 allows step 4 & 5 to be completed from the VI client under the VM's settings.

To increase an existing partition, boot the VM off a Partition Magic iso disc <http://partedmagic.com> and grow the partition. (Can use **diskpart** in Windows guests in certain cases)

List all registered VMs on a host (vmx files): `$ sudo /usr/bin/vmware-cmd -l`

"Power Off" a hard power off, "Shut Down" soft with VMware tools, "Reset" = hard, "Restart" = soft.

Snapshot Manager "Delete" commits the snapshot to the parent and removes the snapshot. "Delete all" commits all the immediate snapshots before the "You are here" state. "Go to" reverts to a particular snapshot. "Revert to snapshot" takes you back to your parent's snapshot of "You are here".



Backups

Backup techniques: 1) File level (or imaging) backups using a backup agent installed within the VM.

2) Backup clients installed in Service Console to backup VMDK files (VMs must be turned off).

3) External datastore backups, using SAN snapshots or backing up NFS server's filesystem.

4) Windows VM file level backup using VCB. 5) VMDK backup using VCB.

Can also use 3rd party VMware specific applications and scripts (e.g. vRanger, xPress, vmkblp)

VCB requires a physical Win 2003 SP1 server connected to VC or single host, VCB supported software, backup hardware and can attach to FC SAN (ESX3.5 allows iSCSI).

VCB cannot backup RDMs in physical compatibility mode, or VMs without an IP address/DNS name.

VCB supports a max of 60 concurrent mounted VM partitions.

VCB currently supports: EMC NetWorker, Symantec Backup Exec, Tivoli Storage Manager, Veritas NetBackup, CA Brightstor eServe, CommVault Galaxy, EMC Avamar, HP Data Protector v5.5 & v6, Vizioncore esxRanger. See http://www.vmware.com/pdf/vi3_backup_guide.pdf for up to date details.

VCB workflow 1) backup software calls pre-backup script, queues NTFS and FAT, puts VMs into snapshot mode, takes the snapshot and makes it available to 3rd party software, image level exports the snapshot, file level mounts the snapshot. 2) Ordinary backup. 3) calls post-backup script, unmounts the snapshot, take the VM out of snapshot mode, commits any disk changes.


vcbMounter SC + Proxy Backs up entire VMs in the Service Console.

vcbRestore SC only Restores data that has been backed up using image-based backups.

mountvm Proxy only Mounts vmdk files.

On the Proxy, commands are in **C:\Program Files\VMware\VMware Consolidated Backup Framework**.

On the Service Console the commands are located in **/usr/sbin**.



Web Access

Access via ESX host or VC: <https://hostname.domain.com/ui>

Check the status with: `$ sudo /sbin/service vmware-webAccess status`

restart if it has stopped: `$ sudo /sbin/service vmware-webAccess start`

Two options are available with Remote Console URLs: 1) Limit view to the remote console - hides details such as event logs. 2) Limit view to a single VM - disables inventory navigation.

These options only affect presentation not access control. Permissions are granted in the VI client. Browser must be IE6 or Firefox 1.5 or higher to be supported. Troubleshooting the Web Browser plugin: Firefox > **about:plugins** > "VMware WebCenter Remote MKS Plug-in" should be 2.0.1.0 I.E. > Tools > Internet Options > Settings > View Objects > "QuickMksAxCtl" should be 2.0.1.0