

Availability

Maximums (per HA cluster): Hosts = 32 VMs = 3000
Maximums (FT): Disks per VM = 16 vCPUs per VM = 1 RAM per VM = 64GB VMs per host = 4

HA: Single master, multiple slaves. If master fails, is shut down or removed from cluster an election occurs. All hosts not in standby, maint mode or disconnected participate. VMs on disconnected hosts are not protected. **Master responsibilities:** • monitors slaves • monitors VMs • restarts VMs on best host • reports cluster HA health to vCenter. **Slaves:** • monitors runtime state • report state to master. Network heartbeat every second between master & slaves. Hosts can use all VMkernel networks for heartbeat except vMotion network (unless vMotion is only available 2nd VMkernel network). If host stops responding to heartbeat, liveness checks datastore heartbeat & pings to mgt address. Host deemed Failed if both tests negative, if it passes liveness check it is Network Isolated (host cannot ping cluster isolation address) or Network Partitioned. Master monitors Network Isolated/Partitioned VMs and restarts them if they power off. Powered-on VMs on Network Isolated hosts are treated by cluster's **Host Isolation Response:** • power off • leave powered on (default) • shut down (requires VMware tools). Cluster selects 2 datastores for heartbeat (*das.heartbeatdsperhost* can increase to 5). Datastores used for datastore heartbeat must be mounted by at least 2 hosts. Disable datastore heartbeating if using Ethernet based storage.

HA Admission Control: Can VMs power on when they violate availability constraints at HA failover. Actions that change a reservation must satisfy admission control. Secondary FT VMs added to calculation. **Control policies:** • Host Failures Cluster Tolerates - 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 Failover Hosts (multiple allowed) - leaves host(s) unused until failure. 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 *das.slotCpuInMHz* or *das.slotMemInMB*.

FT: Uses anti-affinity rules to keep primary & secondary apart, checked when primary powers on. VM-VM affinity applies to primary only, VM-Host affinity applies to primary & secondary.

Requires: HA & host monitoring, host certificate checking (on by default), dedicated logging NIC (not IPv6), compatible CPUs with same extensions & power mgt features, Hardware Virtualization (HV), thick disks on shared storage (VMDKs or vRDMs), supported guest OS, min 3 hosts in cluster.

Not supported: snapshots, Storage vMotion, Linked Clones, hotplugging, MSCS, VM backups, SMP, physical RDMs, **Paravirtualized VMs (available in ESXi5?)**, NPIV, VMDirectPath, 3D video, EPT/RVI. DRS only if cluster is EVC. SiteSurvey can identify configuration issues. vLockstep interval typically < 500ms

vCenter Server Heartbeat: vCenter replication & clustering, monitors services including MSSQL, VUM & View Composer. Automated or Manual failover/failback. Uses Active/Passive nodes with heartbeat.

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

Not supported: DRS on VMs, vMotion, FT, NPIV, Round Robin NMP, FCoE/iSCSI/NFS based disks (SW iSCSI initiator in guest is possible)

	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 initializing LUNs within windows. Add RDMs to 2nd SCSI controller i.e. SCSI(1:x). Set SCSI bus sharing: • CIB = Virtual • CAB or N+1 = Physical

Firewall Ports: HA interhost - 8182 TCP/UDP, FT interhost - 8100, 8200 TCP/UDP

Links: HA Deep Dive <http://www.yellow-bricks.com/vmware-high-availability-deepdiv/>

HA and FT Error Messages <http://kb.vmware.com/kb/1033634>

CPUs & guest OSES that support FT <http://kb.vmware.com/kb/1008027>

MSCS supported configurations (vSphere 4) <http://kb.vmware.com/kb/1037959>

MSCS Support on ESX/ESXi <http://kb.vmware.com/kb/1004617>