VMware HA: Overview & Technical Best Practices

Updated 8/10/2007

What is Business Continuity?

Business Continuity = **"Always-on" uninterrupted availability of business systems and applications**

Business Continuity is about eliminating downtime (MTBF) and reducing time to recovery (MTTR)

Availability = MTBF / (MTBF + MTTR)

Type of Downtime	Business Continuity Component
Unplanned downtime **Planned downtime**	**High availability**
Disasters	**Disaster recovery**

Business Continuity in a Virtualized World

- **Virtualization can improve business continuity by:**
  - Lowering the cost and complexity of achieving business continuity
  - Protecting every workload virtualized in an OS & application independent manner
- **Business continuity related features & benefits of VMware Infrastructure:**
  - Encapsulation & HW independence: Simple disaster recovery
  - ESX NIC teaming & HBA Multipathing: Shared hw redundancy
  - VMotion: Eliminates & automates planned downtime
  - VMware HA: Detects & responds to unplanned downtime
  - DRS: Balances workloads & optimizes resource utilization
  - Consolidated Backup: Protects data & integrates with other tools

vmware

Operating System based Clustering vs. VMware HA

In practice, VMware HA and the other features of VI3 allow many customers to achieve comparable or higher levels of availability vs. legacy solutions

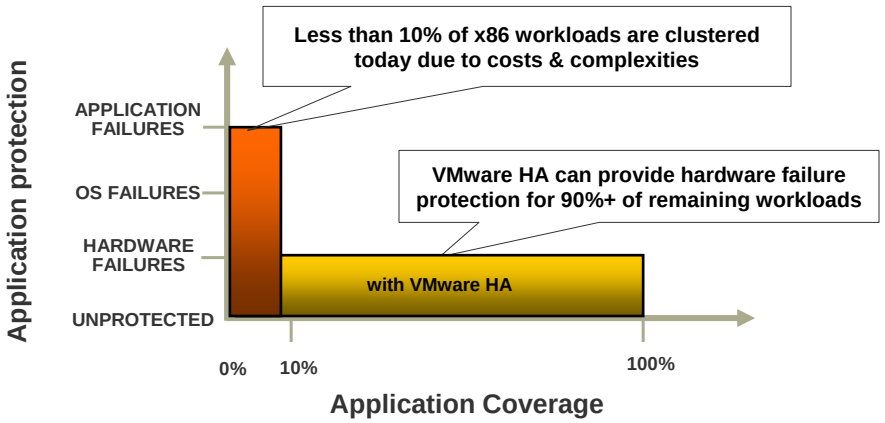

vmware

High Availability with VMware HA

What is VMware HA?

- Automatic restart of virtual machines in case of physical server failures

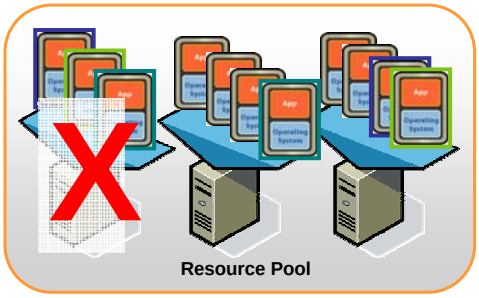

Customer Benefits:

- **High availability with minimal capital & operational costs**
  - Reduced need for passive stand-by hardware & dedicated administrators
- **Broadly applicable**
  - All x86 workloads virtualized are automatically protected without agents or scripting
- **Manageability & Flexibility**
  - Very simple to deploy and maintain
  - Application manageability unchanged
  - N-to-N failovers, hosts and VMs easily moved in and out of clusters

vmware

Detecting Individual Virtual Machine Failures

- **VMware HA does not detect individual VM failures**

- **Alternatives can be used for this purpose:**
  - VMware Tools heartbeats & VirtualCenter alarms
  - 3rd party monitoring agents running inside guests

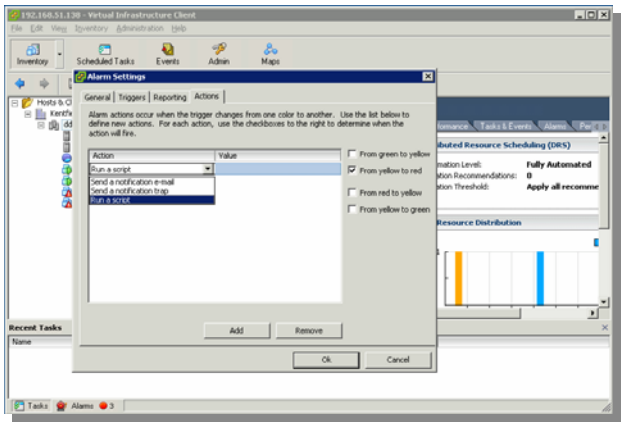

vmware

VMware HA Cluster Configuration: Step 1

VMware HA cluster configuration is composed of two steps:
1) Cluster-wide policies 2) Individual VM customizations

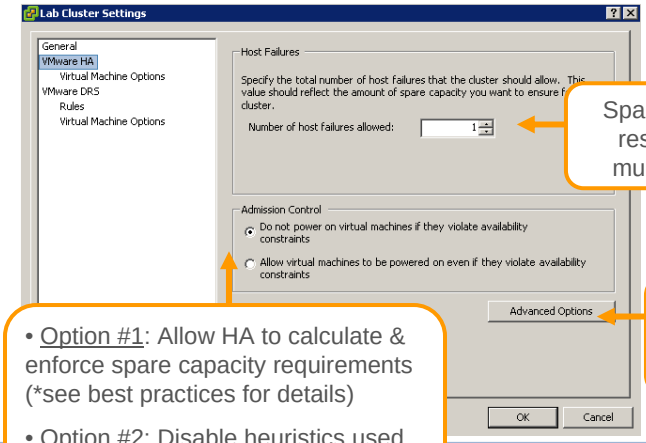

Spare capacity can be reserved to tolerate multiple host failures

See appendix for details regarding advanced settings

- Option #1: Allow HA to calculate & enforce spare capacity requirements (*see best practices for details)
- Option #2: Disable heuristics used by HA for capacity planning

vmware

Host Failure Capacity Planning

VMware HA enforces failover capacity by using memory & CPU reservations to estimate the number of VMs able to recover from multiple host failures

Examples (assuming a 3 node HA cluster):

Failover capacity of 1 limits number of running VMs to fit onto 2 hosts

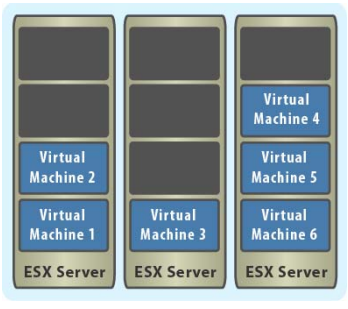

Failover capacity of 2 limits number of running VMs to fit onto 1 host

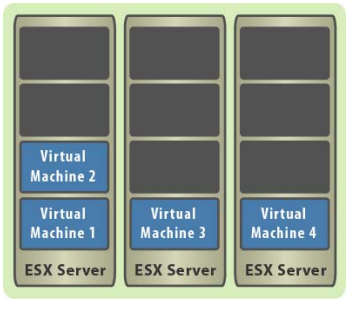

vmware

VMware HA Cluster Configuration: Step 2

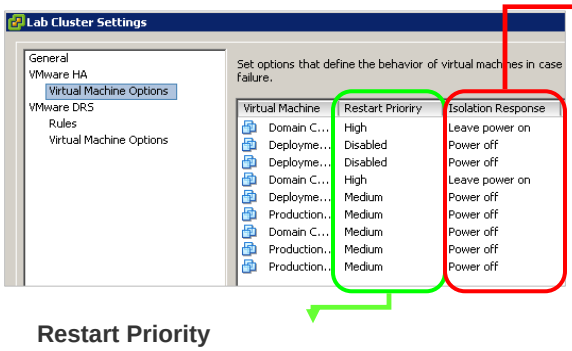

Restart Priority

- Use Low/Medium/High restart priorities to customize failover ordering
  - Default = Medium, High priority VMs are restarted first
- Non-essential VMs should be set to "Disabled" (automated restart will skip them)

Isolation Response

- Initiated when a host experiences network isolation from the rest of the cluster
- **"Power off" is the default**
  - Assumes proper network redundancy makes isolation events very rare
- **"Leave power on" is intended for cases where:**
  - Lack of redundancy & environmental factors make outages likely
  - VM networks are separate from service console (and more reliable)

vmware

VMware HA Under the Covers

HA agents are installed & configured on ESX Hosts via VirtualCenter

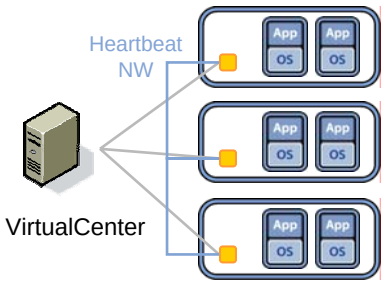

- Post configuration, HA agents maintain a heartbeat network
  - Ability to perform failovers is independent from VirtualCenter availability
- VMFS & shared storage allows hosts to access & power-on virtual machines
  - Distributed locking prevents simultaneous access to protect data integrity
- During a failover, quick restart is the primary goal
  - DRS algorithms balance workloads after HA has recovered virtual machines

vmware

Cluster Configuration & Communication Details

- **Cluster nodes are designated as Primary or Secondary**
  - Primary nodes maintain a synchronized view of entire cluster
    - Up to 5 primary nodes per cluster
  - Secondary nodes are managed by the primary nodes

- **Service console network(s) are used for heartbeats and state synchronization**
  - Minimal network activity in steady state (5 second heartbeat intervals)
  - Additional light traffic during node configuration & VM power operations
  - Incoming ports used: TCP/UDP 8042-8045
  - Outgoing ports used: TCP/UDP 2050-2250

vmware

Failure Detection & Failover Details

- **ESX host failures are detected via missing heartbeats**
  - Failure detection is initiated after 15 seconds
  - Follow-up network ping determines node failure vs. HA agent failure, and a failover is performed if no response is received

- **Failover operations are coordinated as follows:**

1. One of the primary hosts is selected to coordinate the failover, and one of the remaining hosts with spare capacity becomes failover target
2. Virtual machines affected are sorted by priority, and powered on until failover target runs out of spare capacity
3. If the host selected as coordinator fails, another primary continues the effort
4. If one of the hosts that fails was a primary node, one of the remaining secondary nodes is promoted to being a primary

vmware

Isolation Response Details

- **Network failures can cause "split-brain" conditions**
  - In such cases, hosts are unable to determine if the rest of cluster has failed or has become unreachable
- **Isolation response is used to prevent split-brain conditions and is started when:**
  - A host has stopped receiving heartbeats from other cluster nodes AND the isolation address cannot be pinged
  - The default isolation address is the ESX service console gateway, and the default isolation response time is 15 seconds (*can be changed, see appendix)
- **Powering virtual machines off releases VMFS locks, enables other hosts to recover**
  - When "Leave power on" option is set, virtual machines may require manual power-off / migration in case of an actual network isolation

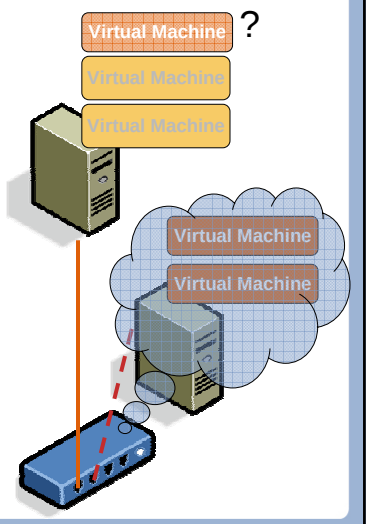

vmware

VMware HA Best Practices – Setup & Networking

1. **Proper DNS & Network settings are needed for initial configuration**
   - After configuration DNS resolutions are cached to /etc/FT_Hosts (minimizing the dependency on DNS server availability during an actual failover)
   - DNS on each host is preferred (manual editing of /etc/hosts is error prone)
2. **Redundancy to ESX Service Console is essential (several options)**
   - Choose the option that minimizes single points of failure
   - Gateways should be IP addresses that will respond via ICMP (ping)
   - Enable PortFast (or equivalent) on network switches to avoid spanning tree isolations
3. **Network maintenance activities should take into account dependencies on the ESX Service Console network(s)**
   - VMware HA can be temporarily disabled through the Cluster->Edit Settings dialog
4. **Valid VM network label names required for proper failover**
   - Virtual machines use them to re-establish network connectivity upon restart

vmware

Network Configuration

Network redundancy between the ESX service consoles is essential for reliable detection of host failures & isolation conditions

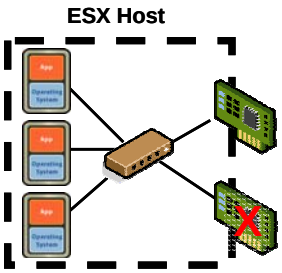

- **A single service console network with underlying redundancy is usually sufficient:**
  - Use a team of 2 NICs connected to different physical switches to avoid a single point of failure
  - Configure vNics in vSwitch for Active/Standby configuration (rolling failover = "yes", default load balancing = route based on originating port ID)
- **Consider extending timeout values & adding multiple isolation addresses (*see appendix)**
  - Timeouts of 30-60 seconds will slightly extend recovery times, but will also allow for intermittent network outages

vmware

Network Configuration (continued)

Beyond NIC teaming, a secondary service console network can be configured to provide redundant heartbeating & isolation detection

- **HA will detect and use a secondary service console network**
  - Adding a secondary service console portgroup to an existing VMotion vSwitch avoids having to dedicate an additional subnet & NIC for this purpose
  - Also need to specify an additional isolation address for the cluster to account for the added redundancy (*see appendix)
- **Continue using the primary service console network & IP address for management purposes**
  - Be careful with network maintenance that affects the primary service console network and the secondary / VMotion network

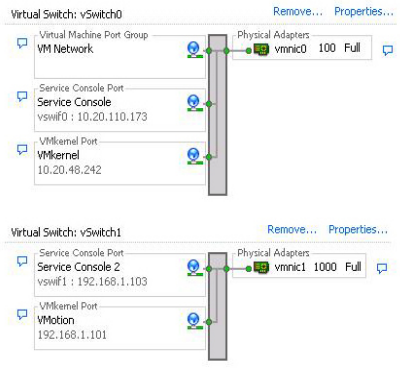

vmware

VMware HA Best Practices – Resource Management

1. **Larger groups of homogenous servers will allow higher levels of utilization across an HA/DRS enabled cluster (on average)**

- More nodes per cluster (current maximum is 16) can tolerate multiple host failures while still guaranteeing failover capacities
- Admission control heuristics are conservatively weighted; must account for extra spare capacity so that large servers with many VMs can failover to small servers

2. **To define the sizing estimates used for admission control, set reasonable reservations as the minimum resources needed**

- Admission control will allow unlimited VMs to power on without any reservations set. When reservations are set, HA will use largest reservation specified as the "slot" size.
- At a minimum, set reservations for a few virtual machines considered "average"

3. **Pay attention to cluster warnings and messages**

- Clusters marked with a YELLOW icon indicate capacity is overcommitted according to DRS; while clusters marked with RED indicate HA may be unable to restart all virtual machines subject to the availability constraints specified

vmware

Troubleshooting VMware HA

1. **Verify the following:**

- IP connectivity, DNS resolution, shared storage and networks are visible throughout the cluster, service consoles have valid & reachable gateways

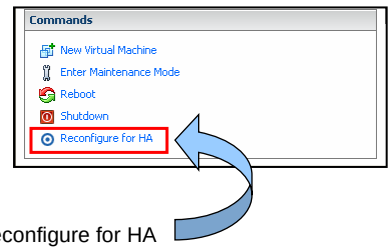

2. **Re-initialize HA cluster configuration**

- Per host: Select ESX Host->Summary Tab->Reconfigure for HA
- Per cluster: Select Cluster->Edit Settings->Uncheck HA enabled, wait for reconfiguration task to complete, and then check to re-enable

3. **Inspect ESX Server logs and HA agent configuration files:**

- Check for ESX Service Console networking errors first, HA agents next
- HA agent logs: /opt/LGTOaam512/log/* & /opt/LGTOaam512/vmsupport/*
- /opt/LGTOaam512/config/vmware-sites contains a list of cluster nodes

vmware

Appendix: HA Customizations

- **Setting preferred failover hosts for individual virtual machines**
  - Select Virtual Machine->Edit Settings->Advanced->Configuration Parameters
  - Add the das.defaultfailoverhost = <hostname> option/value pair to the virtual machine's parameters where <hostname> represents the short name of the preferred host
  - The host specified will be used as a first choice, but HA will automatically select an alternate in case of insufficient capacity
- **Adjusting the default timeout used for failure & isolation detection**
  - **New in VC 2.0.2:** the response time can be configured to be different than 15 seconds (15000 ms). 60 seconds (60000 ms) is an alternative commonly used.
  - Select a cluster->VMware HA->Advanced Options
  - Add the das.failuredetectiontime = <value> option/value pair to the cluster's settings where <value> represents the desired timeout value in milliseconds
  - VMware HA will not declare a host failure nor initiate an isolation detection response until the timeout value specified has been exceeded without heartbeats received

vmware

Appendix: HA Customizations

- **Changing the default isolation response address**
  - By default, HA uses the gateway specified in each hosts service console network configuration as the isolation address to query as the last step in case of isolation
  - Select a cluster->VMware HA->Advanced Options, and add the das.isolationaddress = <value> option/value pair to the cluster's settings where <value> represents the IP address to use as an override to the gateway
  - Use a reliable IP addresses reachable with fewest failure paths possible
- **Setting more than one isolation response address**
  - **New in VC 2.0.2:** more than one isolation response address can be specified, and each service console network should have a different isolation response address defined (by default, only the gateway of the first console network is used)
  - Select a cluster->VMware HA->Advanced Options , and add the das.isolationaddress2 = <value2> option/value pair to the cluster's settings where <value2> represents the secondary IP addresses to use
  - The default timeout value should also be increased to 20 seconds (20000 ms) or greater when a secondary isolation address has been specified

vmware

Appendix: Applying HA Customizations

- **To apply HA customizations:**

1. Create an HA/DRS cluster in the VirtualCenter inventory
2. Set option/value pairs
3. Disable VMware HA, wait for reconfiguration task to complete, and then re-enable HA to have settings take effect

- **Scenario #1: Single service console network with teamed NICs**
  - Some risk assumed by configuring hosts in cluster with only 1 service console network (subnet 10.20.XX.XX; 2 teamed NICs used to protect from NIC failure
  - Default timeout increased to 60 seconds (das.failuredetectiontime = 60000)

- **Scenario #2: Redundant service console networks**
  - Each host in cluster configured with 2 service console networks by leveraging an existing VMotion network (subnets 10.20.YY.YY and 192.168.ZZ.ZZ)
  - Default gateway used for the first network, das.isolationaddress2 = 192.168.1.103 specified as additional isolation address for the second network
  - Default timeout increased to 20 seconds (das.failuredetectiontime = 20000)

vmware

VMware Product & Solution Areas

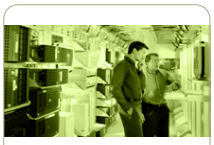

Infrastructure Optimization

Deploy virtual machines that run safely and move transparently across shared hardware.

- Consolidate servers
- Reduce data center operating costs: real estate, power, cooling
- Includes the industry's most widely deployed Virtualization Infrastructure suite

Business Continuity

Keep systems up and running through simple, reliable data protection and pervasive failover protection.

- Reduce planned and unplanned downtime
- Reduce cost and complexity of high availability
- Simplify disaster recovery

Software Lifecycle Automation

Leverage assets and speed software development by automating the setup, sharing and storage of multi-machine configurations.

- Rapidly provision machines
- Improve software quality

Virtual Clients and Desktops

Secure unmanaged PCs while retaining end-user autonomy by layering a security policy in software around desktop virtual machines.

- For enterprises and end users
- Improve security and mobility

vmware