white paper

VMware HA - Overview and Troubleshooting

Title:	VMware HA - Overview and Troubleshooting
Author(s):	Xtravirt (Robert J Wilson, Gavin Jolliffe)
Target Audience:	Technical - Intermediate
Current Revision:	1.3 (Apr 2009)
First Published:	Feb 2009
Vendor(s):	VMware
Product(s):	ESX 3.x, vCenter 2.x (previously VirtualCenter)
UID:	XD10005

Content Overview:

- Introduction to HA and how it works
- References to external links for advanced HA concepts
- Outline some of the commonly reported HA errors and resolutions

1.0 Introduction

This Xtravirt white paper outlines VMware Virtual Infrastructure High Availability (HA), how it works, and some common troubleshooting tips.

2.0 What is HA?

VMware HA is a function which provides high availability services for virtual machines in the event of an ESX server failure.

Key points:

- It is an integrated component of VMware VI3 Enterprise, or available as an add-on license to Foundation and Standard versions
- Monitors a ESX host for failure using Legato AAM heartbeat technology
- Automatically restart VM's from a failed ESX host on other ESX servers in the same cluster
- VirtualCenter 2.5 U1 extends support for HA to ESX3i hosts

Note: VMware KB1004656 provides information on the limited configurations that are supported for VMware HA and ESX3i hosts and KB1004177 notes that ESX3i hosts without swap enabled cannot be added to an HA cluster.

3.0 How does HA work?

3.1 Basic Concepts

- VMware vCenter deploys an agent on each ESX host added to an HA enabled cluster, so that the host can communicate with other hosts to co-ordinate state information, and what actions to take place in the event of an ESX host failure
- The HA agent monitors a heartbeat from each ESX server using the Service Console network
- Up to 5 hosts in a cluster are nominated as 'primary' hosts. The first primary host in cluster has

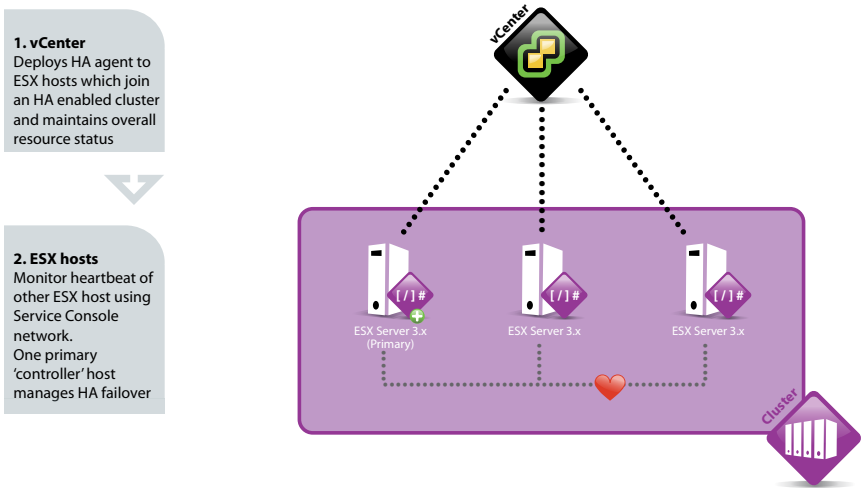

Illustration 1-1: VMware HA Overview

a special 'controller' role and manages HA failover. If this server itself fails, another primary host will be promoted to this role

- If vCenter fails, HA continues to function but cluster resource information will not be updated
- Any new host added to the cluster has to communicate with the controller in order to complete its configuration
- If a primary goes down, HA can promote another ESX server as primary once the failed primary is removed from the cluster. It will also promote a secondary when a primary is placed in Maintenance Mode, and when you select "reconfigure for HA".
- HA supports up to a maximum of 4 host failures
- All VM's and their configuration files must be located on shared storage
- HA restarts only powered on VM's which are on HA enabled hosts. Higher priority VM's are started first
- vCenter provides HA summary status information for any selected Cluster
- Individual VM behaviour regarding restart priority and power management (isolation response) can be configured within vCenter HA Cluster settings
- To use DRS (Distributed Resource Scheduling) with HA, the ESX hosts require a configured vMotion network. If this does not exist, HA will still work and DRS will efficiently place the initial failed-over VM's

3.2 Advanced Concepts

Further detailed information for HA can be found at the following links:

1. Automating High Availability (HA) Services with VMware HA, http://www.vmware.com/pdf/vmware_ha_wp.pdf, VMware
2. How does the VMware HA work, http://www.vmwarewolf.com/how-does-the-vmware-ha-work/, Rick Blythe
3. HA Failover Capacity, http://www.vmwarewolf.com/ha-failover-capacity/, Rick Blythe
4. VMware HA Admission Control, http://www.vmwarewolf.com/vmware-ha-admission-control/, Rick Blythe
5. HA Advanced Options, http://www.yellow-bricks.com/ha-advanced-options/, Duncan Epping
6. VMware High Availability Deep Dive, http://www.yellow-bricks.com/vmware-high-availability-deepdiv/, Duncan Epping
7. Special Considerations for High Availability with Blade Servers, http://www.yellow-bricks.com/2009/02/09/blades-and-ha-cluster-design/, Duncan Epping

4.0 Common Errors and Resolutions

4.1 "Unable to contact primary HA agent within a cluster"

1. In the first instance this error is sometimes resolved by either pressing **reconfigure for HA** or by disabling and then re-enabling HA allowing the process to complete each time.

2. The most common problem with HA and ESX clusters in general are DNS issues. It is important to check and recheck that all related DNS entries are correct throughout your environment and that the FQDN for all the hosts within host files, including the VC server have the same entries. Spelling and case are important for smooth communication and checking for the common mistake of 0 'zero' to O 'ohs' entries is highly recommended as well as standardising with lowercase.

Note: Previously, enabling VMware High Availability required DNS resolution of all ESX Server hosts in a High Availability cluster. This was done by configuring DNS records or by adding all of the host names and IP addresses to the /etc/hosts file on each server.

With the introduction of ESX 3.5 U2 & VC 2.5 U2 this was modified so that DNS resolution is no longer a requirement to

enable HA on ESX server hosts but it is important to note host/IP information is now provided by vCenter, refer: http://www.vmware.com/support/vi3/doc/vi3_esx35u2_vc25u2_rel_notes.html#resolvedhaissues

Also adds full support for monitoring individual virtual machine failures based on VMware tools heartbeats. This release extends support for clusters containing mixed combinations of ESX and ESX Server 3i hosts.

3. Before upgrading ESX or VC to later versions it is highly recommended that the hostname, /etc/hosts files and DNS related areas, including A & PTR records, are correct.

4. Upgrading from VC 2.0 to VC 2.5 & ESX 3.0.x to ESX 3.5 – it has been identified that renaming of a cluster can sometimes resolve this problem. To do this you should disable HA and DRS and allow the processes to complete. Rename the cluster and then re-enable HA and DRS again. You can of course rename the cluster back to the original name if required.

   Occasionally, renaming isn't sufficient and a new cluster has to be created with all hosts moved over to this new cluster. Ensure HA and DRS are disabled prior to moving any hosts to avoid unnecessary work for VC.

4.2 "An error occured during configuration of the HA Agent on the host"

1. If the **Tasks & Events** error shows something like **configuration of the host IP address is inconsistent on host ESX1.local.com: address resolved to 192.168.x.x and 172.168.x.x**, this generally indicates a DNS issue, however resolving where the problem lies can be tricky. SSH onto each of your ESX hosts and check the following to ensure that they are all consistent and up to date:

   /etc/hosts
   /etc/FT_HOSTS (if it exists)
   /etc/sysconfig/network
   /etc/vmware/esx.conf
   cat /proc/sys/kernel/hostname

   Note: Entries should follow the format of: **IP address**, **FQDN**, then **hostname**, and contain all ESX hosts, eg:

   10.0.20.15 esxserver.com esxserver

   Check that all of the DNS entries are correct as with Section 4.1 previously, including checking that the letter case is correct.

   Run the commands **hostname -i** and a **hostname –s** from the Service Console and ensure it returns good results.

   If issues persist, check if the **/etc/opt/vmware/aam/FT_HOSTS** file exists, as it can contain stale information. If it does then delete the file, disable HA on the cluster, then re-enable to complete the HA installation.

4.3 Insufficient resources to satisfy HA failover level on cluster in data center

1. It has been highlighted previously that lowercase characters when naming and configuring an ESX environment should be used. With the release of VC 2.5 U2 it became more apparent, as not only were errors arising relating to the **insufficient resources to satisfy HA failover level on cluster**, but also the error **unable to contact primary HA agent in cluster** was being recorded. The following recommends several checks, the overall goal being to keep hostnames in lowercase.

   From within VirtualCenter:

   - On each host in the cluster, go to the configuration tab, DNS and routing, hostname and ensure everything is in lowercase
   - On each of your ESX hosts, SSH into each of them and check that names are in lowercase within:

   /etc/hosts
   /etc/sysconfig/network
   cat /proc/sys/kernel/hostname

   If they are not correct then update them accordingly.

   After this has been completed you can check the settings are correct using:

   uname -n
   hostname -a
   hostname -s
   hostname -d
   hostname -f
   hostname

   In addition to this, also make sure the hostname in **/etc/vmware/esx.conf** is also lowercase (/adv/Misc/HostName = <FQDN in lowercase>)

   source: http://www.vmwarewolf.com

4.4 General HA errors on hosts when the ESX clusters have inconsistent software versions

1. When upgrading clusters it is sometimes noticed that as soon as one of the hosts is updated, HA can fail to work for that host.

   To alleviate this issue, it is recommended to disable HA on the cluster where the hosts will be updated and update the rest of the hosts. Once all hosts are at the same level HA can be re-enabled.

   If a VM was being moved to the newly updated host, the fact that HA was not running properly on that host can cause VMotion to fail, so disabling HA until the upgrade process is completed should fix this issue.

4.5 Incompatible HA Network

1. VMware KB1006541 identifies a HA Network issue after installation or upgrade to VirtualCenter 2.5 U2. Symptoms being that HA does not work although configured correctly. An error similar to below can appear within the **Tasks & Events** section:

   HA agent on <esxhostname> in cluster <clustername> in <datacenter> has an error Incompatible HA Networks:
   Cluster has network(s) missing on host: x.x.x.x
   Consider using the Advanced Cluster Settings das.allowNetwork to control network usage.

   See the VMware KB article, http://kb.vmware.com/kb/1006541 for more information and a resolution.

4.6 noprimaryagent: Could not find a primary host to configure host

1. Due to this error, many people to have to recreate their cluster(s) due to an upgrade of vCenter causing corruption of their HA database. It can normally be prevented by disabling HA and DRS prior to the upgrade of VC 2.0.1 Patch 2. Occasionally it may be necessary to restart the management agents on the ESX hosts to get them working again in vCenter.

4.7 HA agent has an error: Host in HA cluster must have userworld swap enabled

1. For ESX3i embedded or installable, after the release of VC 2.5 U1, hosts that do not have swap enabled cannot be added to a HA cluster. The error within **Tasks & Events** may show **HA agent has an error: Host in HA cluster must have userworld swap enabled**. VMware KB 1004177 provides details on this and how to resolve the issue.

4.8 General

Note: VirtualCenter 2.5 update 2 adds full support for monitoring individual virtual machine failures based on VMware tools heartbeats. This release also extends support for clusters containing mixed combinations of ESX and ESX Server 3i hosts, and minimizes previous configuration dependencies on DNS.

5.0 Useful References

5.1 HA Log Files

For further troubleshooting, log files for HA can be found in the ESX Service Console here:

/opt/LGTOaam512/	ESX 3.0.x
/var/log/vmware/aam/	ESX 3.5

5.2 VMware HA Knowledge Base Articles

The following is a list of VMware KB articles relating to HA (High Availability):

1. KB1004177, ESXi 3 Hosts without swap enabled cannot be added to a VMware High Availability Cluster, http://kb.vmware.com/kb/1004177
2. KB1004633, Troubleshooting the VMware High Availability error, Unable to contact the primary host in a cluster, http://kb.vmware.com/kb/1004633
3. KB1006541, After installation or upgrade to VirtualCenter 2.5.0 Update 2 an Incompatible HA Networks error is generated, http://kb.vmware.com/kb/1006541
4. KB1002117, Setting Multiple Isolation Response Addresses for VMware High Availability, http://kb.vmware.com/kb/1002117
5. KB1002478, Unable to Configure VMware High Availability because the Default Gateway cannot be reached, http://kb.vmware.com/kb/1002478
6. KB1003691, Diagnosing a VMware High Availability cluster configuration failure, http://kb.vmware.com/kb/1003691
7. KB1004939, VMware High Availability configuration fails with a "no active primaries" found error, http://kb.vmware.com/kb/1004939
8. KB1006421, Advanced Configuration options for VMware High Availability, http://kb.vmware.com/kb/1006421
9. KB1007156, An ESX host in a VMware High Availability cluster fails to enter maintenance mode and stops at 2%, http://kb.vmware.com/kb/1007156
10. KB1007234, VMware High Availability error: cannot read /var/log/vmware/aam/aam_config_util.def, http://kb.vmware.com/kb/1007234
11. KB1001596 , Troubleshooting Adding an ESX Server Host to a VMware High Availability Cluster, http://kb.vmware.com/kb/1001596
12. KB1002080, Best Practices and Advanced Features for VMware High Availability, http://kb.vmware.com/kb/1002080
13. KB1003715 - Recreating a VMware High Availability Cluster, http://kb.vmware.com/kb/1003715

14. KB1003789 - iSCSI Service Console port on the same network as VMware High Availability may break configuration of VMware High Availability, http://kb.vmware.com/kb/1003789
15. KB1004656, Limited configurations are supported for VMware HA and ESX Server 3i hosts, http://kb.vmware.com/kb/1004656
16. KB1004884, DRS Might Be Disabled if DPM Is Used with HA Admission Control, http://kb.vmware.com/kb/1004884
17. KB1004495, VMware HA configuration fails with the error: VMWareClusterManager Rule not enabled, http://kb.vmware.com/kb/1004495
18. KB1004965, VMware HA fails with the error: AAM Agent did not start, http://kb.vmware.com/kb/1004965
19. KB1003713, Configuring name resolution for VMware VirtualCenter, http://kb.vmware.com/kb/1003713
20. KB1003714, Verifying and reinstalling the correct version of VMware VirtualCenter Server agent, http://kb.vmware.com/kb/1003714
21. KB1003735, Identifying issues with and setting up name resolution on ESX Server, http://kb.vmware.com/kb/1003735
22. KB1003742, Determining if a virtual machine is orphaned, http://kb.vmware.com/kb/1003742
23. KB1006336, Troubleshooting a failure to deploy or undeploy a virtual machine, http://kb.vmware.com/kb/1006336
24. KB1002641, Service Console Redundancy for VMware High Availability, http://kb.vmware.com/kb/1002641
25. KB1006013, Virtual machine violates VMware HA availability constraints even when reservation is set to 0, http://kb.vmware.com/kb/1006013
26. KB1005476, Cannot configure VMware High Availability when an ESX host has multiple service consoles, http://kb.vmware.com/kb/1005476
27. KB1003148, VMware HA Log is filling up the File System, http://kb.vmware.com/kb/1003148
28. KB1005735, VMotion, VMware HA and various service console commands fail to run after upgrading an ESX host, http://kb.vmware.com/kb/1005735
29. KB1003009, Restart or Shutdown of an ESX Server 3i Host Does Not Fail Over Virtual Machines to Other Hosts in the Cluster, http://kb.vmware.com/kb/1003009
30. KB1007372, VMware Tools stops responding or fails to install when Virtual Machine Monitoring is selected, http://kb.vmware.com/kb/1007372
31. KB1004385, UserDuct_Open Fails During Crossdup with VMK_WOULD_BLOCK or Waiters List Not Empty Warnings, http://kb.vmware.com/kb/1004385
32. KB1006129, Advisory for advanced VMkernel parameter NFS.LockDisabled, http://kb.vmware.com/kb/1006129
33. KB1000064, Fully Qualified Domain Name Length Limit is Changed in ESX Server 3.0.1 and VirtualCenter 2.0.1, http://kb.vmware.com/kb/1000064
34. KB1003301, Changes in licensing for VI3 Standard Edition when upgrading to VI 3.5, http://kb.vmware.com/kb/1003301
35. KB 1007307, Testing the Monitoring Virtual Machines functionality in VirtualCenter, http://kb.vmware.com/kb/1007307
36. KB 1008479, Edited Advanced Options In an HA-Enabled Cluster Do Not Take Effect, http://kb.vmware.com/kb/1008479
37. KB 1003654, VMware High Availability log file grows large, vmwareclustermanager.trace file, http://kb.vmware.com/kb/1003654

This concludes the white paper.

About Xtravirt

Xtravirt is a knowledge-based company that delivers its expertise in virtualization online and in person. We have developed a reputation for astute leadership and expertise through our work with an impressive array of organisations. It is this real-world experience that drives our ability to provide independent, current and free advice online.

We work with organisations whose IT staff are frustrated with how hard it is to find detailed information and skills around virtualisation. We help our clients deliver the true benefits of virtualization, resulting in cost and time savings.

For more information contact:

Dorset House, Regent Park
297 Kingston Road, Leatherhead
Surrey KT22 7PL
t +44 (0) 1372 824 296
f +44 (0) 1372 824 576
e information@xtravirt.com
w www.xtravirt.com

References

1. Automating High Availability (HA) Services with VMware HA, http://www.vmware.com/pdf/vmware_ha_wp.pdf, VMware
2. How does the VMware HA work, http://www.vmwarewolf.com/how-does-the-vmware-ha-work/, Rick Blythe
3. HA Failover Capacity, http://www.vmwarewolf.com/ha-failover-capacity/, Rick Blythe
4. VMware HA Admission Control, http://www.vmwarewolf.com/vmware-ha-admission-control/, Rick Blythe
5. HA Advanced Options, http://www.yellow-bricks.com/ha-advanced-options/, Duncan Epping
6. VMware High Availability Deep Dive, http://www.yellow-bricks.com/vmware-high-availability-deepdiv/, Duncan Epping
7. Blades and HA/Cluster Design, http://www.yellow-bricks.com/2009/02/09/blades-and-ha-cluster-design/, Duncan Epping

Useful Links

1. Refer 5.2, VMware HA Knowledge Base Articles

Tags

VMware, ESX, HA, High Availability, troubleshooting, error, log, knowledge base, kb, DNS