

Home - VMware - VCDX Preparation - - Enterprise Admin - S4 - Continuity - O2 - Advanced HA

Objective 4.2 – Configure advanced HA deployments

Written by Matthijs van den Berg

Thursday, 26 March 2009 11:08

Knowledge

- **Describe guidelines for restart priority and isolation response.**

A host in a VMware HA cluster might lose its console network connectivity. Such a host is isolated from other hosts in the cluster. Isolation is a special case in which an ESX Server host has not actually failed, but its service console network is broken (for example, due to switch failure, Ethernet adapter failure, or some similar cause). Isolation is handled as a special case of failure in VMware HA. The isolated host intentionally shuts down the virtual machines running on that host, so they can be restarted on hosts in the cluster that are not isolated.

By default, about 12 seconds after heartbeats have ceased arriving, the isolated host itself starts a special procedure called isolation response. This typically involves shutting down its virtual machines. About 15 seconds after the start of the isolation event, other hosts in the cluster consider that isolated host as failed and attempt to restart the virtual machines. You can change both of these timeout values from their defaults using VMware HA advanced options in VirtualCenter.

You need to make sure that there are enough resources in the cluster to account for host failures. Each VM that was killed on an ESX host failure needs to be restarted on another ESX host. The HA process also uses some resources.

- **Explain how to customize a typical HA deployment**

There are a couple of customizations that you can perform on an HA deployment.

- Change the restart priority of a VM
- Change the isolation response of a VM
- Add a second service console to counter non disruptive malfunctions on the primary service console
- Use DNS name resolving instead of host files

- **Understand HA communication (heartbeat)**

HA is configured on ESX server level. The Virtual Center server is only used for configuring HA, not for the operation. For HA to work the ESX hosts communicate with each other. This communication is called the heartbeat. The VMware HA agent placed on each host maintains a heartbeat with the other hosts in the cluster using the service console network. Each server sends heartbeats to the other servers in the cluster at five-second intervals. If any servers lose heartbeat over three consecutive heartbeat intervals, VMware HA initiates the failover action of restarting all affected virtual machines on other hosts.

The HA heartbeat uses the following ports for communication from host to host:

- Incoming Port: TCP/UDP 8042-8045
- Outgoing Port: TCP/UDP 2050-2250

- **Detail impact of DRS affinity rules on an HA cluster**

DRS affinity assures that certain designated hosts are never of always running on the same ESX host. When HA kicks in on an ESX server crash this mechanism ensures that for example two VM that have an affinity rule configured to keep them together are started on the same host.

- **Describe troubleshooting techniques**

Typically, troubleshooting VMware HA involves following the steps below and taking corrective action to see if the problem is fixed:

- Verify the network is working properly. At a minimum, your VirtualCenter Management Server must be able to reach the ESX Server hosts, and the hosts must be able to ping the gateway on the service console networks. If you are using multiple service console networks for redundancy, each of the service consoles must be able to ping the gateway IP address on that network. For information on networking troubleshooting, refer to the ESX Server Configuration guide.
- Verify shared storage is accessible from all hosts in the cluster. Carefully verify connectivity at server, storage, and switch levels.
- Pay attention to cluster warnings, task details, and events. Task details, configuration issue error messages, and event histories in VirtualCenter are useful places to look for information related to an error.
- Check logs for clues. Logs on the ESX Server hosts are the best place to start troubleshooting VMware HA. Always start by looking for ESX Server service console networking errors, followed by VMware HA agent problems.
- Use the Reconfigure for HA menu in the VI Client as a last resort to reconfigure a host when the VMware HA agent on the host appears unresponsive or have encountered configuration error. This option initiates configuration and restart of the

Menu - VCDX

- Introduction
- More Info
- - Enterprise Admin
- - Design Exam

VCDX Cloud

virtual
server
vmware network
host file service
esxcfg click
configure machine
traffic vswitch
machines storage
command console
hosts cluster disk
client physical
configuration files port

Recent Blog Postings

Recently I needed to adjust the Queue Depth (Q Depth) of a HBA. When we encountered trouble using VCB 1.5 one of the first recommendations of HP / VMware was to adjust the Q-Depth of the HBA devices on the system. To adjust this first we need to look up the adaptor type. In the following example

[Read more](#)

Add Banners

● Explain best practices for HA deployment

More info on best practice can be found here: <http://kb.vmware.com/>

The configuration of ESX Server host networking and name resolution, as well as the networking infrastructure external to ESX Server hosts (switches, routers, and firewalls), is critical to optimizing VMware HA setup. The following suggestions are best practices for configuring these components for improved HA performance:

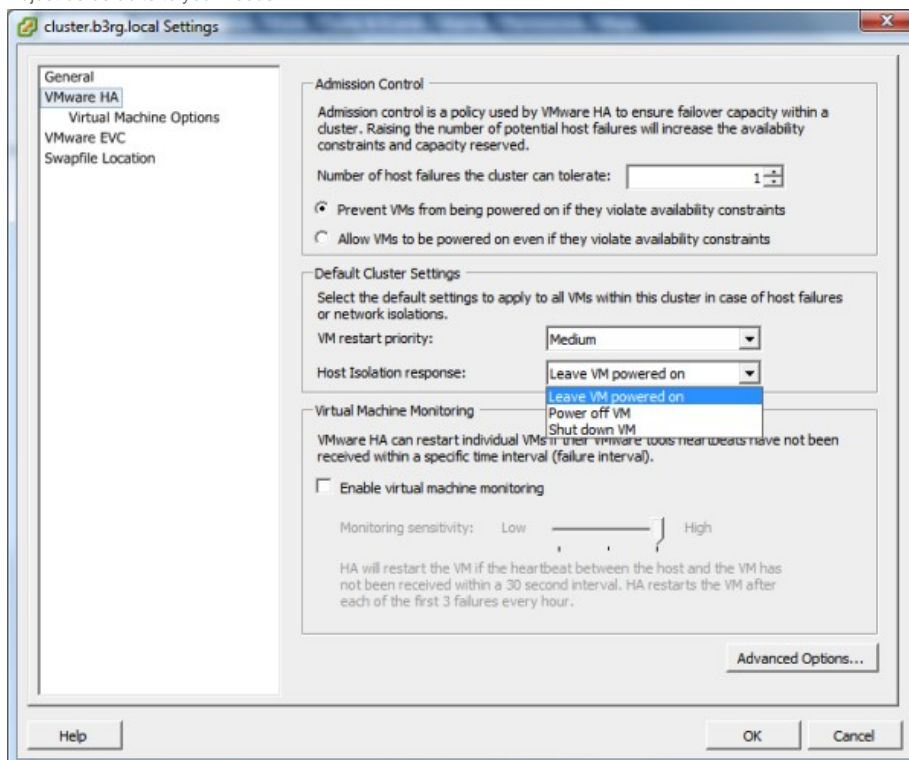
- If your switches support the PortFast (or an equivalent) setting, enable it on the physical network switches connecting servers. This helps avoid spanning tree isolation events. For more information on this option, refer to the documentation provided by your networking switch vendor.
- Ensure that the following firewall ports are open for communication by the service console for all ESX Server 3 hosts:
 - Incoming Port: TCP/UDP 8042-8045
 - Outgoing Port: TCP/UDP 2050-2250
- For better heartbeat reliability, configure end-to-end dual network paths between servers for service console networking. You should also configure shorter network paths between servers in a cluster. Routes with too many hops can cause networking packet delays for heartbeats.
- If redundant service consoles are on separate subnets, specify "isolation address" for each service console that is on its subnet. By default, gateway address for the network is used as isolation address.
- Disable VMware HA (using VirtualCenter, deselect the Enable VMware HA check box in the cluster's Settings dialog box) when performing any networking maintenance that might disable all heartbeat paths between hosts.
- Use DNS for name resolution rather than the error-prone method of manually editing the local /etc/hosts file on ESX Server hosts. If you do edit /etc/hosts, you must include both long and short names.
- Use consistent port names on VLANs for public networks on all ESX servers in the cluster. Port names are used to reconfigure access to the network by virtual machines. If the names are used on the original server and the failover server are inconsistent, virtual machines are disconnected from their networks after failover.

Skills and Abilities

● Configure restart priority and isolation response

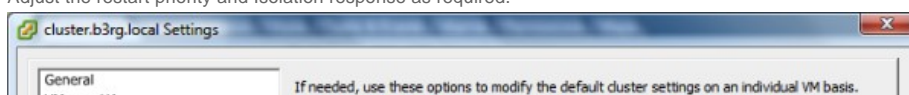
○ Cluster-wide setting

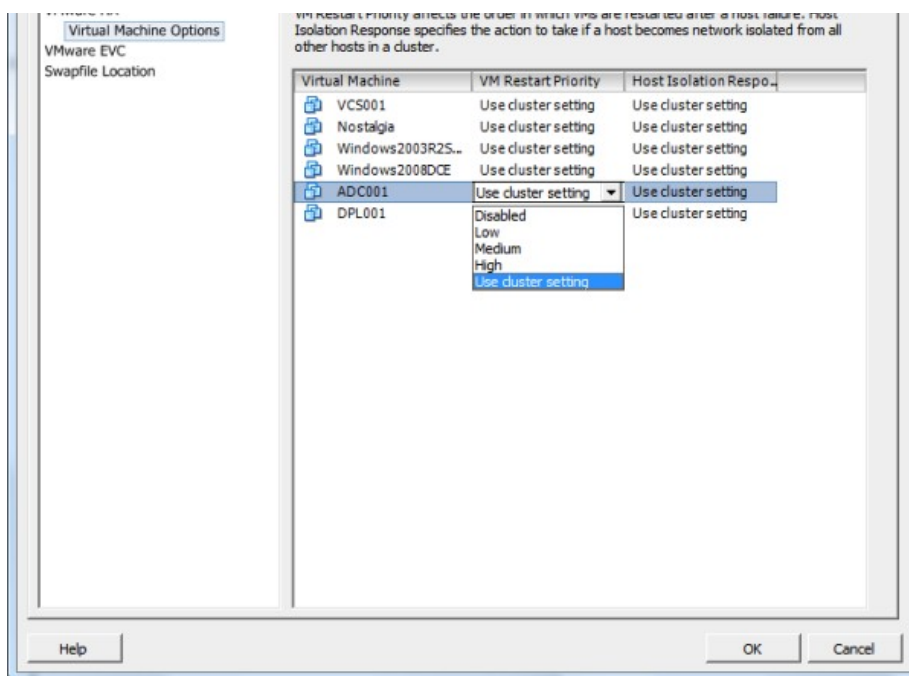
- Go to the VMware Infrastructure Client
- Right click the cluster and select "Edit Settings"
- Click "VMware HA"
- Adjust the defaults to your needs



○ Individual VM override settings

- Go to the VMware Infrastructure Client
- Right click the cluster and select "Edit Settings"
- Click in "Virtual Machine Options" under VMware HA.
- Adjust the restart priority and Isolation response as required.





- **Configure advanced HA options**

More info: <http://pubs.vmware.com/> and <http://www.yellow-bricks.com/>

- **Failure detection time**

The failure detection time is the time period a server waits before going into isolation mode after when not receiving heartbeats. This value is configured in the variable "**das.failedetectiontime**". To change this:

- In the cluster's Settings dialog box, select VMware HA.
- Click the Advanced Options button to open the Advanced Options (HA) dialog box.
- Enter each advanced attribute you want to change in a text box in the Option column and the value it should be set in the Value column.
- Click OK.

- **Redundant isolation address settings**

To add a second IP address for the HA Agent you can use the value: **das.isolationaddress2** and enter an IP address here. You can use the same way to change this as use to change the Failure Detection Time.

It is also possible to add a second IP address via the CLI:

```
esxcfg-vswif -a -i <IP address> -n <netmask> -p <port group> <vswif name>
```

- If you want it on the same physical device that the current connection is on, specify Service\ Console for the portgroup, otherwise you need to give it the name of the portgroup connected to a different nic. You can see all of the portgroup configuration by doing `esxcfg-vswitch -l`
- **Default failover host**
You can configure a hostname in the value below to configure a default failover host.: **das.defaultfailoverhost**

- **Configure physical switch settings to support HA**

? Not quite sure what is meant here. Could be the correct switch port settings, or to connect te second service console to a different physical switch. Perhaps they want to hear that you need to configure link state tracking (disabling a switchport when the uplink failes).

- **Troubleshoot HA deployments**

- **Failover capacity**
There is a good article about failover capacity on VMwarewofls blog. Read it here. <http://www.vmwarewolf.com/ha-failover-capacity/>
- **Examine log entries**
The HA agent logs in the file `/var/log/vmware/vpx/vpxa.log`.
- **Correct network issues**
There can be several issues. Be sure to have a good DNS resolution in place (FQDN lookup, short name lookup, reversed lookup). Also make sure there are no packet errors / resends on the network.

Tools

- CLI
 - `esxcfg-advcfg`
 - `hostname -s`
- VI Client

< Prev

Next >

