



# A S e r i e s F a i l u r e s

[HOME](#) [ABOUT](#) [VCAP-DCA INDEX](#) [THINAPP](#) [VCDX STUFF](#) [OTHER PAGES](#) [VCAP-DCD INDEX](#)

« VCAP-DCA Objective 4.3 : Configure a vSphere Environment to Support MSCS Clustering

VCAP-DCA Objective 2.1 : Implement and Manage Complex Virtual Networks »

## VCAP-DCA Objective 2.1 : Implement and Manage Complex Virtual Networks

Another VCAP-DCA objective here. For a complete list of study objectives for the VCAP-DCA ( VDC4-410) browse to <http://www.vfail.net/vcap-dca/>.

This is the first pass through the objective for implementing and maintaining complex VMware HA solutions. A must read for this section is [Duncan Epping's HA Deepdive](#).

If you haven't already done so, download the zip file full of additional study resources from Kendrick Coleman's blog. These can be found at his [VCAP exam landing page](#). He has compiled a lot of resources and broken them down by objective, and these make great additional reading if you are not greatly familiar with any of the following.

### Knowledge

**Identify the three admission control policies for HA**

**Identify heartbeat options and dependencies**

### Skills and Abilities

**Calculate host failure requirements**

**Configure customized isolation response settings**

The isolation response is the action HA will take when the heartbeat network is isolated. The response is either power off, leave powered on(default), or shut down.

HA will try to restart the affected Virtual Machines and by default will try up to five times. This is configurable with the parameter below.

```
das.maxvmrestartcount
```

The default value for isolation failure detection is 15 seconds. This is configurable with the below parameter.

```
das.failedetectiontime
```

At this point a restart is initiated by one of the primary hosts. The isolation response is actually initiated 1 second before the failure detection time.

Read the below excerpt from Duncan Epping's blog on why it is important to understand these different parameters and what affect configuring isolation response settings can have on your environment.

*The default value for isolation/failure detection is 15 seconds. In other words the failed or isolated host will be declared dead by the other hosts in the HA cluster on the fifteenth second and a restart will be initiated by one of the primary hosts.*

*For now let's assume the isolation response is "power off". The "power off"(isolation response) will be initiated by the isolated host 1 second before the das.failedetectiontime. A "power off" will be initiated on the fourteenth second and a restart will be initiated on the fifteenth second.*

*Does this mean that you can end up with your VMs being down and HA not restarting them?*

Search

GO

- > SRM (3)
- > ThinApp (16)
- > VCAP-DCA (38)
- > VMware (57)



### MY LATEST TWEETS

- vFAIL.NET** [vfail.net](#)  
**97 followers**
- RT @PlanetV12n: To Trunk Or Not To Trunk? (vDestination)  
<http://bit.ly/byqscT> >good read  
about 1 week ago
- New blog post: VCAP-DCA Update  
<http://bit.ly/b1DSLS>  
about 1 week ago
- Just found out lab will goto 4.1 next year. Items like those discussed in my VMTN posting not covered on the current  
<http://bit.ly/bYL2c6>  
about 1 week ago
- All VCAP-DCA Study notes now online at <http://bit.ly/9rnpHE>  
Final PDF version done and will be posted next week.  
about 1 week ago

Virtualization

Quote of the Day

Web2PDF

converted by Web2PDFConvert.com

*Yes, when the heartbeat returns between the 14th and 15th second the "power off" could already have been initiated. The restart however will not be initiated because the heartbeat indicates that the host is not isolated anymore.*

*How can you avoid this?*

*Pick "Leave VM powered on" as an isolation response. Increasing the das.failedetectiontime will also decrease the chances of running in to issues like these.*

*Basic design principle: Increase "das.failedetectiontime" to 30 seconds (30000) to decrease the likely-hood of a false positive.*

A couple of additional parameters I've found necessary

*das.isolationaddressx Used to configure multiple isolation addresses.*

*das.usedefaultisolationaddress Set to true/false and used in the cse where a default gateway is not pingable, in which case this set to false in conjunction with configuring another address for das.isolationaddress*

### Configure HA redundancy in a mixed ESX/ESXi environment

Redundant management networks are recommended for HA and there are two options to choose from, Network Redundancy Using NIC Teaming or Network Redundancy Using a Secondary Network.

### Configure HA related alarms and monitor an HA cluster

### Create a custom slot size configuration

A slot refers to a logical representation to power on a virtual machine in the cluster, taking into account memory and cpu resources that will satisfy this request.

This is calculated using the highest cpu and memory reservation of any given VM, with a default of 256mhz for cpu and OMB + memory overhead if no reservations are specified.

Set das.slotCpuInMHz or das.slotMemInMB to manually lower the slot size in cases where one VM causes a really large slot size.

### Understand interactions between DRS and HA

4.1 has brought a lot of changes to this topic and until I know with certainty that 4.1 will not be included any time soon I will shy away from speaking of those changes.

In terms of functionality today, DRS and HA while separate pieces do work together to provide availability and performance. When a host goes down and HA performs a failover and restarts a virtual machine it is looking for to provide availability. It is then DRS's job to balance these machines across the cluster.

### Create an HA solution that ensures primary node distribution across sites

The first 5 hosts that join the HA cluster are automatically selected as primary nodes.

You can manually view which nodes are primary with this command

```
cat /var/log/vmware/aam/aam_config_util_listnodes.log
```

Re-election of a primary will only occur when a primary is placed in maintenance mode, disconnected or removed from the cluster, or when a cluster is reconfigured for HA.

If all fail simultaneously, there is no HA initiated restart of VMs that occur.

In order to design an HA solution that ensures a primary is always available, the placement of your hosts is crucial. For each cluster, never put more than four hosts in a place where it could be a single point of failure, for example a chassis of blades. If you have 10 blades and two chassis, look to separate the blades amongst two chassis, and additionally make sure that no more than four blades are from one cluster in each chassis.

Read more about this concept here <http://www.yellow-bricks.com/2009/02/09/blades-and-ha-cluster-design/>

### Analyze vSphere environment to determine appropriate HA admission control policy


You have three choices for HA admission control policies

Host failures cluster tolerates

*Host with most slots is taken out of the equation, and then the next most if more than one 1 is set.*

*Your percentage of resources should be equal or larger than your largest host so that all vm's on that host can be restarted.*

*Tends to be very conservative as largest reservation dictates the slot size.*

 Configuration  
Maximums for  
VMware  
vSphere is my favorite  
VMware document. It  
answers many of the "How  
many", "How much" type  
questions about VI  
capabilities. This is one of  
the documents that will  
most often be updated as  
new releases of VMware VI  
are released so it's a good  
one to keep tabs on.

<http://www.vmware.com/pdf/vsphere4/r41>

Powered by  
[The Printed Owl](#)

## CATEGORIES

- » [SRM](#) (3)
- » [ThinApp](#) (16)
- » [VCAP-DCA](#) (38)
- » [VMware](#) (57)

Percentage of cluster resources reserved as failover spare capacity



*Instead of using slot sizes, sets a percentage of resources to be left unused for HA purposes*



*Tends to be a more realistic view of reservations as it uses actual reservations vs. slot size.*

*More flexible.*

Specify a failover host



*Specifies a single host specifically for failover purposes. You may end up with a lot of reserved capacity for failover, but you also will only get a single host for use as a failover server.*

## Analyze performance metrics to calculate host failure requirements

### Analyze Virtual Machine workload to determine optimum slot size

If you are manually specifying the memory and cpu values for the slot size, make sure the slot size is representative of typical workloads.

### Analyze HA cluster capacity to determine optimum cluster size

How many hosts in your cluster?

How many host failures can you tolerate?

What is the resource utilization of your virtual machines?

## Tools

vSphere Availability Guide

Product Documentation

vSphere Client

## Other Relevant Articles and Links Related to this Section

[http://searchsystemschannel.techtarget.com/generic/0,295582,sid99\\_qci1515486,00.html](http://searchsystemschannel.techtarget.com/generic/0,295582,sid99_qci1515486,00.html)

<http://www.dailyhypervisor.com/2009/03/31/vmware-ha-cluster-sizing-considerations/>

<http://www.yellow-bricks.com/vmware-high-availability-deepdiv/>

<http://www.b3rg.nl/vcdx/section-4-business-continuity-and-data-protection/objective-4.2-configure-advanced-ha-deployments.html>

- [Objective 9.3 – Configure vCenter Server Linked Mode](#)
- [Objective 9.2 – Plan and Execute Scripted Installations](#)
- [Objective 8.2 – Administer vCenter Orchestrator](#)
- [VCAP-DCA Brownbag Session #1](#)
- [VCAP-DCA Objective 7.1 : Secure ESX\(i\) Hosts](#)
- [VCAP-DCA Objective 9.1 : Install ESX Server with Custom Settings](#)
- [VCAP-DCA Objective 3.5 – Utilize Advanced vSphere Performance Monitoring Tools](#)
- [VCAP-DCA Objective 3.4 – Perform Capacity Planning in a vSphere Environment](#)
- [VCAP-DCA Objective 3.3 – Implement and Maintain Complex DRS Solutions](#)
- [VCAP-DCA Objective 3.2 – Optimize Virtual Machine Resources](#)

[HA](#), [VCAP-DCA](#), [VMware](#), [vsphere](#)

Share  
this  
post!

Print  
article

This entry was posted by [Sean Crookston](#) on August 27, 2010 at 7:00 am, and is filed under [VCAP-DCA](#), [VMware](#). Follow any responses to this post through [RSS 2.0](#). You can [leave a response](#) or [trackback](#) from your own site.

COMMENTS (0)

RELATED POSTS

NO COMMENTS YET.

Name (required)

E-mail (required, will not be published)

Website

---

Submit Comment

---

Mystique theme by [digitalnature](#) | Powered by [WordPress](#)

 **RSS FEEDS** | **XHTML 1.1** | **TOP** | **PR 0**  