VMware HA Implementation Notes

V 0.1-001

ESXi / ESX 3.5 and VirtualCenter 2.5 U3

2008.10.13

Executive Summary

This document summarizes important new features introduced in the ESXi / ESX 3.5 and VirtualCenter 2.5 U3 and best practices of their implementation. For better understanding of those new features some VMware HA cluster essentials are explained as well.

Intended audience are the experienced HA cluster administrators which are plan to migrate to Update 3 as well as any HA cluster administrator interested in the explanation of some of cluster essentials.

How to use this Document

This document was originally written by one of VMware Technical Account Managers to provide a summary answer to most frequent question asked by his TAM customers about HA cluster. Significant contributors to it were engineers from VMware HA team. For details see the part "About Authors".

This document is intended for the free distribution among VMware users community and may be published or referenced in any blogs. This is not an official VMware document and should be treated as answers posted by VMware employees in the community conferences or as their blog entries.

This is a pre-release of the document. The next version will be published in VIOPS (http://viops.vmware.com).

The goal of this document is to provide detailed explanation to information you may find in the VMware documentation. It is an addition to the documentation, not its superstition. In case of any doubts the documentation is the correct source.

When some of topics discussed in this document are not covered by the documentation, they may be considered as "undocumented features". Please note, that VMware, like all others software manufactures, supports only those features, which are described in the official documentation. Authors of this document strongly discourage to use "undocumented features" for any other purpose as for self-education or for quick fixes. In the last case ("quick fix") you have to discuss your actions with VMware support specialists.

Generally, if you want to use the feature you read about in the community forums or in document like that for a long term solution, please consult your Technical Account Manager (TAM) or VMware representative. VMware always hears to wishes of its customers. TAMs filed on behalf of their customers dozens of feature requests many of them were already implemented or will be implemented soon in the VMware software. Please use your TAM properly rather than hack.

Since authors of this document have a lot of work and this document is not, what they are being paid for, we cannot promise to update or expand this document with the each version or update release of HA software. We cannot promise as well to answer any question related to this document.

For feedback or questions on the document please use the forums on VIOPS which will be created soon.

Contents

Cluster Networking

Definitions

Cluster communication network consists of one or more service console (ESX classic) or management (ESXi) portgroups which are allowed to be used for VMware HA communications and belong to the same IP network.

For example interfaces 192.168.2.12 and 192.168.2.99 will belong to the same cluster communication network and interfaces 192.168.2.12 and 192.168.3.13 to different networks. In all cases the network mask is assumed to be 255.255.255.0. It does not play any role, whether interfaces belong to ESX classic or ESXi nodes.

Since the Update 3 to the VirtualCenter server you may override using the advanced configuration option "das.bypassNetCompatCheck=true".

By default, all service console port groups will be used for cluster communications, unless the default is overridden by the advanced HA configuration parameter `das.allowNetwork[n].`

Any ESXi management portgroup will be used for the cluster communications unless it used by VMotion as well. This may be overridden by the advanced HA configuration parameter `das.allowVmotionNetworks`. If, however, there is only a single port group available for use and it is also VMotion-enabled, it will be used by HA.

All mentioned advanced parameters are explained in [1] on page 127 and in details this document.

Service console (ESX classic) and management (ESXi) interfaces which are allowed for cluster communications will be called in this document “cluster network interfaces”.

Building Cluster Networks

To understand following paragraphs we need to explain how the VMware HA forms the cluster communication networks:

1. VMware HA picks up all port groups available for HA communication services. As described above
    - On ESX classic – all service console port groups. Or, when the `das.allowNetwork` parameter is set, port group names matching the set value.
    - On ESXi all management port groups not used for VMotion. However VMotion port groups explicitly allowed by the parameter `das.allowVmotionNetworks` will be included as well, unless only one port group is defined and it also is VMotion-enabled.
2. VMware HA matches all port groups according to their network configuration (subnet) to networks.
3. VMware HA enforces the network compatibility:
   It compares the cluster communication networks which already exists in the cluster and compares them with the cluster communication networks used by node being added to the cluster. If the new node has incompatible networks, it will be rejected.
   The cluster formation begins with the first node added to the cluster.

Network Compatibility Enforcement

Since the Update 2 of VirtualCenter 2.5, compatibility of cluster communication networks is enforced. Update 3 adds the ability to suppress such enforcement, when the architecture solution requires it.

Compatibility requires that any node in a VMware HA cluster should be connected to all cluster networks. Nodes which are not compatible with existing cluster networks will not allowed to join the cluster.

"cluster has network X.X.X.X that does not exist on the host"
"Consider using Advanced Cluster Settings das.allowNetwork to control network usage"

For example we have the cluster with following communication networks:

1. 192.168.2.0 / 24
2. 192.30.1.0 / 24

The following nodes will not be allowed to join this cluster:

1. Node with only one cluster network interface 192.168.2.12 will be rejected because it is not connected to the cluster network 192.30.1.0 / 24
2. Node with cluster network interfaces 192.168.2.14 and 192.30.2.16 will be rejected for two reasons:

a. It is not connected to the cluster network 192.30.1.0 / 24
b. Nodes, which already formed the HA cluster, are not connected to the network 192.30.2.0 / 24

This enforcement was introduced to minimize the time necessary to detect the cluster node failure. Since the UDP protocol is used for cluster heartbeats there is a risk that such heartbeats will be lost during the routing across different networks. Thus to allow the reliable detection of the host failure event, TCP timeouts were used prior to Update 2 to VirtualCenter 2.5 to determine the host failure event for hosts from the different networks. This caused the increase of the time necessary to detect such event and to longer outages.

For such reasons the incompatibility of cluster communications networks, which was tolerated prior to Update 2, was disabled since that update.

There are no plans to relax enforcement of cluster networks compatibility. However the the new advanced HA configuration parameter introduced with Update 3 allows suppressing this enforcement.

Advanced Parameter - `das.allowVmotionNetworks`

This advanced HA configuration parameter is described in the Resource Managed Guide [1] as follows:

> Allows a NIC that is used for VMotion networks to be considered for VMware HA usage. This permits a host to have only one NIC configured for management and VMotion combined. By default, any VMotion network is ignored.

This parameter allows the sharing of the same port group by both VMotion and HA clusters communication services only for ESXi hosts. Its values may be character strings "**true**" or "**false**" (without quotes).

The standard edition of ESX distinguishes among the service console port groups (vswif***n***) used for management and HA cluster communication purposes and VMkernel port group used for VMotion and iSCSI or NFS communications. Thus the sharing of the same port group by HA cluster communications services and VMotion is not possible. The parameter `das.allowVmotionNetworks` has no meaning for such systems.

In ESX Server 3i, the VMkernel networking interface provides network connectivity for the ESXi Server host as well as handling VMotion and IP storage. This forces HA to use the VMkernel portgroups for its cluster communication services. When HA chooses portgroups, it skips, by default, those which are configured for VMotion unless the only port group available is configured for VMotion.

When the parameter das.allowVmotionNetworks is defined and set to true, that forces HA to use the port group with VMotion enabled.

Use Case 1 for `das.allowVmotionNetworks`

Classic ESX server, virtual switch vSwitch2 configured with only one NIC. This switch has both service console port group "ServiceConsole" and VMkernel port group with enabled VMotion – "VMotion".

In this case the parameter `das.allowVmotionNetworks` does not change the HA networking behavior. Independent of the value of this parameter HA cluster communication services will always use the "ServiceConsole" port group. This is caused by the architecture essentials of the classic ESX server: all VMkernel port groups are not presented to the Service Console.

Use Case 2 for `das.allowVmotionNetworks`

ESXi server. The virtual switch vSwitch2 has two VMkernel port groups "Management1" and "Management2". VMotion is enabled on "Management2"

It this case by default only the "Management1" port group will be used by the HA cluster communication services. When the parameter `das.allowVmotionNetworks` is defined to be "true" both port groups "Management1" and "Management2" will be used for the HA cluster communication services.

Use Case 3 for `das.allowVmotionNetworks`

ESXi server. The virtual switch vSwitch2 has the VMkernel port group "Management1" with VMotion enabled.

It this case the port group "Management1" will be considered for the HA cluster communication services even if das.allowVmotionNetworks is not defined or is explicitly set to "false". This is because HA needs some network over which to communicate, and if the only port group is VMotion-enabled, HA will use that port group.

Use Case 4 for `das.allowVmotionNetworks`

ESXi server. The virtual switch vSwitch2 has the VMkernel port group "Management1" with VMotion enabled. The virtual switch vSwitch3 has two VMkernel port groups "Management2" and "Management3". VMotion is enabled on "Management1" and "Management2"

User wants to allow cluster communications on "Management2" but not on "Management1"

In this case if the parameter `das.allowVmotionNetworks` is defined to be "true" both port groups "Management1" and "Management2" will be considered for the HA cluster communication services. User should use another advanced parameter das.allowNetwork[n] to choose the "Management2" port group.

Advanced Parameter - `das.allowNetwork<n>`

This parameter is described in the Resource Managed Guide [1] as follows:

> Enables the use of port group names to control the networks used for VMware HA. You can set the value to be "Service Console 2" or "Management Network" to use (only) the networks associated with those port group names in the networking configuration.

This parameter allows HA to filter port groups used for the cluster communications. When we will go back to the description of how the VMware HA build cluster networks made in the paragraph Building Cluster Networks – this parameter is used during the first step.

This parameter may be used either as a single parameter `das.allowNetwork` or as a parameter group `das.allowNetwork1, das.allowNetwork2, das.allowNetwork3` … etc. It is not supported to mix both forms of this parameter i.e. to use `das.allowNetwork` and `das.allowNetwork1` together.

Value of this parameter is the character string which needs to match exactly the name of the service console (classic ESX) or management (ESXi) port group on all cluster nodes. Comparison is case sensitive.

The parameter `das.allowNetwork[n]` should be defined to the name of the port group that should be used for cluster communications. When you desired to use only one group the parameter should be used in the form `das.allowNetwork.` When there is more than one group, this parameter should be used as `das.allowNetwork1`, `das.allowNetwork2` etc.

When VMware HA picks up the port groups to be used for cluster communications and the parameter `das.allowNetwork[n]` is defined, only port groups matching this value will be chosen. You must be very careful when you define this parameter. A spelling mistake in its value, for example "SevriceConsole" instead of "ServiceConsole", may prevent the nodes from joining the cluster.

The parameter is not used during the second step, when HA forms the cluster networks. That means that port groups used for different parameters, for example for `das.allowNetwork1` and `das.allowNetwork2` may build one cluster network. See the use cases below.

The parameter `das.allowNetwork[n]` will not suppress the cluster network compatibility check.

The following use cases explain the usage of this parameter:

Use Case 5 for `das.allowNetwork[n]`

Host	ESX	Port Group	VMotion	IP
ESXHOST15	ESX 3.5U2 or higher VC 2.5U2 or higher	`ServiceConsole`	n/a	192.164.11.15
		`ServiceConsole1`	n/a	10.54.20.15
ESXHOST21	ESX 3.5U2 or higher VC 2.5U2 or higher	`ServiceConsole`	n/a	192.164.11.21
		`ServiceConsole1`	n/a	10.54.20.21

Network mask is always assumed to be 255.255.255.0

Problem description:

Do we need the advanced HA parameter `das.allowNetwork[n]` here?

Solution:

No, for this use case you don't need to define the parameter `das.allowNetwork[n]` since all cluster networks are compatible.

However you may consider defining the following parameters

```
das.allowNetwork1 =     ServiceConsole
das.allowNetwork2 =     ServiceConsole1
```

This will have the following effect on the cluster bahaviour

- Nodes with misspelled Service Console port group names will not be allowed to join the cluster
- Nodes with an additional Service Console port group, for example "ServiceConsole 3" will join this cluster without complictions

Use Case 6 for `das.allowNetwork[n]`

Host	Datacenter OS	Port Group	VMotion	IP
ESXHOST15	ESX 3.5U2 or higher VC 2.5U2 or higher	`ServiceConsole`	n/a	192.164.11.15
ESXHOST21	ESX 3.5U2 or higher VC 2.5U2 or higher	`ServiceConsole`	n/a	192.164.**22**.21

Network mask is always 255.255.255.0

Problem description:

User hopes that the advanced HA parameter `das.allowNetwork[n]` will allow him to suppress the compatibility check

Solution:

No, the parameter not intended for such usage. Please consider to put both nodes on the same network. Alternatively, consider using the das.bypassNetCompatCheck advanced configuration parameter.

Use Case 7 for `das.allowNetwork[n]`

Host	Datacenter OS	Port Group	VMotion	IP
ESXHOST15	ESX 3.5U2 or higher	`ServiceConsole1`	n/a	192.164.11.15
	VC 2.5U2 or higher	`ServiceConsole2`	n/a	10.54.20.15
ESXHOST21	ESX 3.5U2 or higher	`ServiceConsole1`	n/a	192.164.11.21
	VC 2.5U2 or higher	`ServiceConsole2`	n/a	10.54.20.21

Network mask is always 255.255.255.0

das.allowNetwork is defined to "ServiceConsole"

Problem description:

Cluster cannot be formed

Solution:

Once the parameter `das.allowNetwork` is defined only matching port groups are allowed to join the cluster networks. You need either to change `das.allowNetwork` to match one of the existing port groups or, preferably, delete the `das.allowNetwork` and define das.allowNetwork1 and `das.allowNetwork 2` to use both of them.

Use Case 8 for `das.allowNetwork[n]`

Host	Datacenter OS	Port Group	VMotion	IP
ESXHOST15	ESX 3.5U2 or higher VC 2.5U2 or higher	`ServiceConsole`	n/a	192.164.11.15
		`ServiceConsole1`	n/a	10.54.20.15
ESXHOST21	ESX 3.5U2 or higher VC 2.5U2 or higher	`ServiceConsole2`	n/a	192.164.11.21
		`ServiceConsole3`	n/a	10.54.20.21
		`ServiceConsole8`	n/a	192.164.**22**.30

Network mask is always 255.255.255.0

Problem description:

When ESXHOST21 joins the cluster formed by the ESXHOST15, it will be rejected because it has the network 192.164.**22**.0/24 which is incompatible with cluster networks 192.164.11.0/24 and 10.54.20/24.

When ESXHOST15 joins the cluster formed by the ESXHOST21 it will be rejected because it does not have network 192.164.**22**.0/24 which is present in the cluster.

Solution:

Define following advanced HA parameters:

```
das.allowNetwork1 =     ServiceConsole
das.allowNetwork2 =     ServiceConsole1
das.allowNetwork3 =     ServiceConsole2
das.allowNetwork4 =     ServiceConsole3
```

This will force the VMware HA consider only port groups ServiceConsole and ServiceConsole1 on ESXHOST15 and ServiceConsole2 and ServiceConsole3 on ESXHOST21. Incompatible network ServiceConsole8 on will be excluded.

Please note that different service console port group names were used just to illustrate that the parameter `das.allowNetwork[n]` affects the choice of the port group only. There is no such rule which requires that the first network should be build only from the port groups defined by `das.allowNetworks1`, the second only from the port groups defined by `das.allowNetworks2,` etc.

The best practice recommended by VMware is to use the same service console or management port groups names on all nodes.

Use Case 9 for `das.allowNetwork[n] and das.allowVmotionNetworks`

Host	Datacenter OS	Port Group	VMotion	IP
ESXHOST15	ESX 3.5U2 or higher VC 2.5U2 or higher	`Management1`	Yes	192.164.11.15
		`Management2`		10.54.20.15
ESXHOST21	ESX 3.5U2 or higher VC 2.5U2 or higher	`Management1`	Yes	192.164.11.21
		`Management2`		10.54.20.21
		`Management3`	Yes	192.164.**22**.30

Network mask is always 255.255.255.0

Problem description:

User want to use the port group Management1 for both VMotion and cluster communication network and port group Management3 for VMotion only.

Solution:

Define following advanced HA parameters:

```
das.allowVmotionNetworks      =     true
das.allowNetwork1             =     Management1
das.allowNetwork2             =     Management2
```

Advanced Parameter - `das.bypassNetCompatCheck`

This parameter suppresses the network compatibility check. You must use it with extreme caution because the cluster which uses incompatible networks for its communications may need more time to detect the ESX/ESXi server host failure which will increase the downtime of virtual machines.

Value of this parameter may be either "**`true`**" or "**`false`**" (without quotes).

Use Case 9 for `das.bypassNetCompatCheck`

Host	Datacenter OS	Port Group	VMotion	IP
ESXHOST15	ESX 3.5U2 or higher VC 2.5U3 or higher	`ServiceConsole`	n/a	192.164.11.15
ESXHOST21	ESX 3.5U2 or higher VC 2.5U3 or higher	`ServiceConsole`	n/a	192.164.**22**.21

Network mask is always 255.255.255.0

Problem description:

User looks for the possibility to suppress the network compatibility check

Solution:

User must define the following advanced cluster parameter: `das.bypassNetCompatCheck=true`.

Side Effect:

Since there is no compatible network for establishing heartbeats, HA will rely on TCP timeouts for host failure detection. No heartbeats will be sent between hosts using networks on different subnets.

DNS Name Resolution on ESX hosts

The release notes for ESX / ESXi 3.5U2 VCU2 mention "DNS Resolution Is No Longer a Requirement to Enable VMware High Availability on ESX Server Hosts".
That means, that ESX servers do not rely upon any DNS settings and/or even the host file /etc/host to resolve the DNS names of the other cluster members.

When the new host joins an HA cluster, the VirtualCenter server provides it with the IP address of one of the HA cluster members, which initially gives the host the network information it needs to contact one of the cluster members. That one cluster member then supplies the network information for all other nodes in the cluster after initially contacted. This information is stored into the FT_HOSTS file located in

- **`/etc/opt/vmware/aam/FT_HOSTS`** on standard ESX
- **`/var/run/vmware/aam/FT_HOSTS`** on ESXi

This allows the HA cluster to bootstrap without any contact with DNS server. Once configured, nodes in the HA cluster do not rely on any name-to-IP address translation. Instead, the local FT_HOSTS cache is used.

If a cluster member becomes network-isolated, it will be able to re-join the HA cluster using the information from the FT_HOSTS file. If any cluster member's IP information had since changed, the contents of this file will then be updated by one of the primary nodes (unless all hosts have changed their IP addresses and the re-joining member can not find any of them using the old IP values).

An outage of VirtualCenter server will not prevent the HA cluster from fulfilling its functions. Failover of Virtual Machines will occur, and host resource data used by the HA policy will be shared among cluster members directly, without any reliance on the VC server.

The requirements for DNS server configuration described in the Resource Management Guide pp. 90, 125 and 172 are obsolete and can be ignored. However DNS is still necessary for other tasks important for the functionality of the ESX server. Among them are

- Translation of the NTP server name
- Translation of the license server name (this name is passed by the VirtualCenter server to ESX server when later gets connected to the VirtualCenter.

VMware HA Restart Policy

Confusion about "Host isolation response" and "VM Restart Priority"

There is a widespread confusion about those parameters. For example: which of them prevents VM from being started by VMware HA after the host failure.

The parameter "**Host isolation response**":

- Is relevant only to the host which detects its own isolation.
- Does not affect the behavior other hosts in the HA cluster
- Note that it is not used by the other HA cluster members when they decide whether to restart the Virtual Machines from the isolated host
- Defines which of the following actions will be applied to VMs running on the isolated host:
  - power off
  - shutdown
  - leave power on (do nothing)
- This parameter affects whether VM will be turned off[1] or not by the isolated host on which it is running, but does not affect how it will attempt to be restarted by the other hosts in the HA cluster.
- This parameter may be defined for both HA cluster and a single VM. The cluster-wide parameter setting will affect all VMs where the "Host isolation response" is not explicitly defined. When the host isolation response is defined for the single VM it overrides the cluster-wide setting.

The parameter "**VM restart priority**":

- Is used by the failover policy of hosts in the cluster when they detect that one of hosts appears to have failed (whether it actually failed, or can no longer communicate with any other node in the HA cluster over any HA network).
  Is valid for both isolated and failed hosts, since the other nodes in the cluster cannot tell whether a host has failed or is network isolated.
- Is relevant only for virtual machines that were running when the host on which they were running failed or became network-isolated.
- Defines the following actions to be applied to VMs
  - Restart high priority
  - Restart low priority
  - Do not restart
- This parameter affects, whether VM will be restarted and with which priority it will restarted. It does not affect VMs that were powered off.

[1] "Turned off" – includes both „Power off" and „Shutdown"

HA Cluster VM Restart Behavior

Since the nodes of HA cluster are not able to distinguish whether a node is isolated or has failed, they attempt to restart Virtual Machines designated by the parameter "VM Restart Priority" when they do not get heartbeats from this host for 15 seconds, and cannot ping any of the management networks.

This timeout value may be changed using the following advanced HA cluster parameter:

das.failuredetectiontime – This parameter specifies the time value in milliseconds. So a 15 second failure detection time would be specified as 15000.

This parameter is described on the page 127 of Resource Management Guide.

If a VM was running on a network-isolated host, the VM will not be able to be failed over to another host in the cluster until it has been shutdown or powered off on the network-isolated node. The shared storage on which the virtual machine resides will have locks on the virtual machine files, preventing them from being accessed by another node in the cluster that may be attempting to power on the virtual machine(s). The behavior of a network-isolated hosts is determined by the "Host isolation response" setting. When it set either to "Power off" or "Shutdown" the restart usually[2] will succeed. When the VM is left powered on, attempts to restart it from another node will be prevented by the VMFS lock on the virtual machine files. In this case restart attempts will fail.

The HA cluster will repeat its attempts to restart VM until

- The VM is shut down or powered off which frees its lock.
- The formerly isolated host is no longer isolated and reconnects to the VC server.
- The formerly isolated host subsequently fails.
- VM will be deregistered from the VirtualCenter server.

It is very important to understand that the VM lock is an essential part of the VMFS design. When VMs are hosted on a non-VMFS datastore, for example NFS, there is no virtual machine lock. This may cause corruption of the virtual disk if the VM is started on multiple servers. Both fibre channel and iSCSI SAN support VMFS.

When a host fails, the locks on the virtual machine files are freed and restart of VMs from the failed host will succeed.

New Value for the "Host isolation response"

VMware HA cluster parameter has the new value "Shut down" As documented in the Resource Management Guide on page 188:

[2] We do not consider all possible reasons like SAN failure etc.

> Another option is Shut down VM, which instructs the virtual machine to shut down its guest operating system before it powers off. This results in a slower failover time, compared to the Power off VM option, but also ensures greater data consistency. VMware Tools must be installed and running for this option to be available.

When the host initiates the shutdown of the Virtual Machine, it waits up to 5 minutes (by default) for shutdown to complete. When this timeout expires, the VM will be powered off, without waiting any longer for the shutdown to be completed.

The following advanced HA cluster configuration parameter: **`das.isolationShutdownTimeout -`** allows to change the time HA waits for the VM shutdown until to power off the VM. The value for this parameter should be set in seconds. The default value is 300 seconds.

The Default Value for the "Host isolation response" Changed

As documented in the Resource Management Guide page 97 and 117 the default value for the "Host isolation response" parameter of VMware HA cluster has changed to "Leave powered on" (in all previous releases it was "power off"). Existing clusters will not be affected by this change. I.e. whatever you set as the isolation response will be retained during the upgrade. If you left this value as default – it will remain as "power off". New clusters created after the upgrade to U2 will have the "leave powered on" set by default.

The explanation for this decision can be found on page 117:

> "This option [Leave powered on] might be preferable if your network is less redundant and network outages are more likely than single-host isolation incidents"

VMware support specialists were often confronted with outages caused by the network failures lasting longer than 15 seconds. An example of such long network failure may be rebuilding of spanning tree by CISCO switches. In this case, since all HA cluster members lost communication to all other cluster members, each HA cluster member assumes itself to be isolated and applies the "Host Isolation Response" parameter, which was "Power Off" by default. Thus non-redundant or improperly configured networks could cause HA cluster to power off all Virtual Machines.

The best way to avoid such outages is to use redundant cluster networks. VMware strongly encourages all customers to follow this best practice.

The default value of the "Host isolation response" was set to "leave power on" to allow to new HA cluster administrators to avoid outages mentioned above.

Troubleshooting

How to Delete an Advanced HA Parameter

Please note that the following issue is resolved in the Update 3. With this update you may delete misspelled parameters just by removing them from the form. Prior to update 3, you can use the following method could used to delete parameters. This paragraph is left in this document to illustrate how the HA cluster parameters are stored in the VirtualCenter database.

Sometimes it may happen that because of the spelling mistake you entered the wrong advance parameter name, for example `das.allowNetworkS1` instead of `das.allowNetwork1`. When the wrong value of a parameter can be easily corrected, the misspelled parameter cannot be removed.

This paragraph describes how to fix such a mistake. Please note that this method is unsupported and exploits undocumented features of the product which may be changed without notice. When you decide to apply it, please:

1. Consult the VMware support center or your technical account manager about alternatives methods or ask them to assist you.
2. Make the full backup of the database you use.

Please note that neither VMware nor the author of this document are responsible for any damage caused by this proposed fix.

Step1

Identify the internal identifier for the cluster you want to repair. The table VPX_ENTITY will allow you find the correlation between the cluster name and this ID.

Step 2

Retrieve (sql SELECT) the XML code for advanced cluster parameters. You need select the row from the table VPX_COMPUTE_RESOURCE using the ID obtained in the first step as a key. You need only column DAS_OPTIONS.

Step 3

You will get text which looks like this

```
<obj xmlns="urn:vim25" xsi:type="ArrayOfOptionValue">
      <OptionValue xsi:type="OptionValue">
            <key>das.allownetworks</key>
            <value xsi:type="xsd:string">this is a wrong parameter</value>
      </OptionValue>
      <OptionValue xsi:type="OptionValue">
            <key>das.allownetwork1</key>
            <value xsi:type="xsd:string">ServiceConsole1</value>
      </OptionValue>
```

```
</obj>
```

Edit it by removing the whole entry for the <OptionValue> key with the wrong parameter name – highlighted text in example above.

Step 4

Replace this text in the database using the SQL update statement.

References

1. VMware Resource Management Guide
   http://www.vmware.com/pdf/vi3_35/esx_3/r35u2/vi3_35_25_u2_resource_mgmt.pdf
2. ESX / ESXi 3.5 U2 VCU2 Release Notes
   http://www.vmware.com/support/vi3/doc/vi3_esx35u2_vc25u2_rel_notes.html
3. ESX / ESXi 3.5 U2 VCU2 Release Notes - HA Issues
   http://www.vmware.com/support/vi3/doc/vi3_esx35u2_vc25u2_rel_notes.html#knownhaissues
4. VMware VirtualCenter 2.5 Update 3 Release Notes
   http://www.vmware.com/support/vi3/doc/vi3_vc25u3_rel_notes.html#resolvedissues
5. http://kb.vmware.com/kb/1006541
6. http://kb.vmware.com/kb/1006606

About Authors

This document was initially written by Seva Semouchin – Technical Account manager based in Germany.

Author especially thanks Marc Sevigny who spend a lot of time aligning the language of this paper with English language and its contents with HA Cluster reality. Other significant contributors were Sridhar Rajagopal and Elisha Ziskind. Many thanks as well to all VMware PSO employees whose high professionalism and broad experience helped this paper to become what it is.