

vSphere Troubleshooting

Update 1

Modified on 04 OCT 2017

VMware vSphere 6.5

VMware ESXi 6.5

vCenter Server 6.5

You can find the most up-to-date technical documentation on the VMware Web site at:

<https://docs.vmware.com/>

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

Copyright © 2010–2017 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About vSphere Troubleshooting	5
Updated Information	7
1 Troubleshooting Overview	9
Guidelines for Troubleshooting	9
Troubleshooting with Logs	11
2 Troubleshooting Virtual Machines	13
Troubleshooting Fault Tolerant Virtual Machines	13
Troubleshooting USB Passthrough Devices	18
Recover Orphaned Virtual Machines	19
Virtual Machine Does Not Power On After Cloning or Deploying from Template	19
3 Troubleshooting Hosts	21
Troubleshooting vSphere HA Host States	21
Troubleshooting vSphere Auto Deploy	25
Authentication Token Manipulation Error	31
Active Directory Rule Set Error Causes Host Profile Compliance Failure	32
Unable to Download VIBs When Using vCenter Server Reverse Proxy	32
4 Troubleshooting vCenter Server and the vSphere Web Client	35
Troubleshooting vCenter Server	35
Troubleshooting the vSphere Web Client	36
Troubleshooting vCenter Server and ESXi Host Certificates	37
5 Troubleshooting Availability	39
Troubleshooting vSphere HA Admission Control	39
Troubleshooting Heartbeat Datastores	41
Troubleshooting vSphere HA Failure Response	42
Troubleshooting vSphere Fault Tolerance in Network Partitions	44
Troubleshooting VM Component Protection	45
6 Troubleshooting Resource Management	47
Troubleshooting Storage DRS	47
Troubleshooting Storage I/O Control	52
7 Troubleshooting Storage	55
Resolving SAN Storage Display Problems	56
Resolving SAN Performance Problems	57
Virtual Machines with RDMS Need to Ignore SCSI INQUIRY Cache	62

Software iSCSI Adapter Is Enabled When Not Needed	62
Failure to Mount NFS Datastores	63
Troubleshooting Storage Adapters	63
Checking Metadata Consistency with VOMA	64
No Failover for Storage Path When TUR Command Is Unsuccessful	65
Troubleshooting Flash Devices	66
Troubleshooting Virtual Volumes	69
Troubleshooting VAIO Filters	71

8 Troubleshooting Networking 73

Troubleshooting MAC Address Allocation	74
The Conversion to the Enhanced LACP Support Fails	77
Unable to Remove a Host from a vSphere Distributed Switch	78
Hosts on a vSphere Distributed Switch 5.1 and Later Lose Connectivity to vCenter Server	79
Hosts on vSphere Distributed Switch 5.0 and Earlier Lose Connectivity to vCenter Server	80
Alarm for Loss of Network Redundancy on a Host	81
Virtual Machines Lose Connectivity After Changing the Uplink Failover Order of a Distributed Port Group	82
Unable to Add a Physical Adapter to a vSphere Distributed Switch That Has Network I/O Control Enabled	83
Troubleshooting SR-IOV Enabled Workloads	84
A Virtual Machine that Runs a VPN Client Causes Denial of Service for Virtual Machines on the Host or Across a vSphere HA Cluster	84
Low Throughput for UDP Workloads on Windows Virtual Machines	86
Virtual Machines on the Same Distributed Port Group and on Different Hosts Cannot Communicate with Each Other	88
Attempt to Power On a Migrated vApp Fails Because the Associated Protocol Profile Is Missing	89
Networking Configuration Operation Is Rolled Back and a Host Is Disconnected from vCenter Server	90

9 Troubleshooting Licensing 93

Troubleshooting Host Licensing	93
Unable to Power On a Virtual Machine	94
Unable to Configure or Use a Feature	94

Index 97

About vSphere Troubleshooting

vSphere Troubleshooting describes troubleshooting issues and procedures for VMware vCenter Server[®] implementations and related components.

Intended Audience

This information is for anyone who wants to troubleshoot virtual machines, ESXi hosts, clusters, and related storage solutions. The information in this book is for experienced Windows or Linux system administrators who are familiar with virtual machine technology and data center operations.

NOTE Not all functionality in the vSphere Web Client has been implemented for the vSphere Client in the vSphere 6.5 release. For an up-to-date list of unsupported functionality, see *Functionality Updates for the vSphere Client Guide* at <http://www.vmware.com/info?id=1413>.

Updated Information

This *vSphere Troubleshooting* is updated with each release of the product or when necessary.

This table provides the update history of the *vSphere Troubleshooting*.

Revision	Description
04 OCT 2017	<ul style="list-style-type: none">■ Updated log information in “vSphere Auto Deploy TFTP Timeout Error at Boot Time,” on page 25.■ Updated log directories in “Troubleshooting with Logs,” on page 11.
EN-002608-00	Initial release.

Troubleshooting Overview

vSphere Troubleshooting contains common troubleshooting scenarios and provides solutions for each of these problems. You can also find guidance here for resolving problems that have similar origins. For unique problems, consider developing and adopting a troubleshooting methodology.

The following approach for effective troubleshooting elaborates on how to gather troubleshooting information, such as identifying symptoms and defining the problem space. Troubleshooting with log files is also discussed.

This chapter includes the following topics:

- “[Guidelines for Troubleshooting](#),” on page 9
- “[Troubleshooting with Logs](#),” on page 11

Guidelines for Troubleshooting

To troubleshoot your implementation of vSphere, identify the symptoms of the problem, determine which of the components are affected, and test possible solutions.

Identifying Symptoms	A number of potential causes might lead to the under-performance or nonperformance of your implementation. The first step in efficient troubleshooting is to identify exactly what is going wrong.
Defining the Problem Space	After you have isolated the symptoms of the problem, you must define the problem space. Identify the software or hardware components that are affected and might be causing the problem and those components that are not involved.
Testing Possible Solutions	When you know what the symptoms of the problem are and which components are involved, test the solutions systematically until the problem is resolved.



Troubleshooting Basics (http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_vsphere_troubleshooting)

Identifying Symptoms

Before you attempt to resolve a problem in your implementation, you must identify precisely how it is failing.

The first step in the troubleshooting process is to gather information that defines the specific symptoms of what is happening. You might ask these questions when gathering this information:

- What is the task or expected behavior that is not occurring?

- Can the affected task be divided into subtasks that you can evaluate separately?
- Is the task ending in an error? Is an error message associated with it?
- Is the task completing but in an unacceptably long time?
- Is the failure consistent or sporadic?
- What has changed recently in the software or hardware that might be related to the failure?

Defining the Problem Space

After you identify the symptoms of the problem, determine which components in your setup are affected, which components might be causing the problem, and which components are not involved.

To define the problem space in an implementation of vSphere, be aware of the components present. In addition to VMware software, consider third-party software in use and which hardware is being used with the VMware virtual hardware.

Recognizing the characteristics of the software and hardware elements and how they can impact the problem, you can explore general problems that might be causing the symptoms.

- Misconfiguration of software settings
- Failure of physical hardware
- Incompatibility of components

Break down the process and consider each piece and the likelihood of its involvement separately. For example, a case that is related to a virtual disk on local storage is probably unrelated to third-party router configuration. However, a local disk controller setting might be contributing to the problem. If a component is unrelated to the specific symptoms, you can probably eliminate it as a candidate for solution testing.

Think about what changed in the configuration recently before the problems started. Look for what is common in the problem. If several problems started at the same time, you can probably trace all the problems to the same cause.

Testing Possible Solutions

After you know the problem's symptoms and which software or hardware components are most likely involved, you can systematically test solutions until you resolve the problem.

With the information that you have gained about the symptoms and affected components, you can design tests for pinpointing and resolving the problem. These tips might make this process more effective.

- Generate ideas for as many potential solutions as you can.
- Verify that each solution determines unequivocally whether the problem is fixed. Test each potential solution but move on promptly if the fix does not resolve the problem.
- Develop and pursue a hierarchy of potential solutions based on likelihood. Systematically eliminate each potential problem from the most likely to the least likely until the symptoms disappear.
- When testing potential solutions, change only one thing at a time. If your setup works after many things are changed at once, you might not be able to discern which of those things made a difference.
- If the changes that you made for a solution do not help resolve the problem, return the implementation to its previous status. If you do not return the implementation to its previous status, new errors might be introduced.
- Find a similar implementation that is working and test it in parallel with the implementation that is not working properly. Make changes on both systems at the same time until few differences or only one difference remains between them.

Troubleshooting with Logs

You can often obtain valuable troubleshooting information by looking at the logs provided by the various services and agents that your implementation is using.

Most logs are located in `C:\ProgramData\VMware\vCenterServer\logs` for Windows deployments or `/var/log/` for Linux deployments. Common logs are available in all implementations. Other logs are unique to certain deployment options (Management Node or Platform Services Controller).

Common Logs

The following logs are common to all deployments on Windows or Linux.

Table 1-1. Common Log Directories

Log Directory	Description
applmgmt	VMware Appliance Management Service
cloudvm	Logs for allotment and distribution of resources between services
cm	VMware Component Manager
firstboot	Location where first boot logs are stored
rhttpproxy	Reverse Web Proxy
sca	VMware Service Control Agent
statsmonitor	VMware Appliance Monitoring Service (Linux only)
vapi	VMware vAPI Endpoint
vmaffd	VMware Authentication Framework daemon
vmdird	VMware Directory Service daemon
vmon	VMware Service Lifecycle Manager

Management Node Logs

The following logs are available if a management node deployment is chosen.

Table 1-2. Management Node Log Directories

Log Directory	Description
autodeploy	VMware vSphere Auto Deploy Waiter
content-library	VMware Content Library Service
eam	VMware ESX Agent Manager
invsvc	VMware Inventory Service
mbcs	VMware Message Bus Config Service
netdump	VMware vSphere ESXi Dump Collector
perfcharts	VMware Performance Charts
vmcam	VMware vSphere Authentication Proxy
vmdird	VMware Directory Service daemon
vmsyslog collector	vSphere Syslog Collector (Windows only)
vmware-sps	VMware vSphere Profile-Driven Storage Service
vmware-vpx	VMware VirtualCenter Server
vpostgres	vFabric Postgres database service

Table 1-2. Management Node Log Directories (Continued)

Log Directory	Description
mbsc	VMware Message Bus Config Service
vsphere-client	VMware vSphere Web Client
vcha	VMware High Availability Service (Linux only)

Platform Services Controller Logs

You can examine the following logs if a Platform Services Controller node deployment is chosen.

Table 1-3. Platform Services Controller Node Log Directories

Log Directory	Description
cis-license	VMware Licensing Service
sso	VMware Secure Token Service
vmcad	VMware Certificate Authority daemon
vmdird	VMware Directory Service

For Platform Services Controller node deployments, additional runtime logs are located at `C:\ProgramData\VMware\CIS\runtime\VMwareSTSService\logs`.

Troubleshooting Virtual Machines

The virtual machine troubleshooting topics provide solutions to potential problems that you might encounter when using your virtual machines.

This chapter includes the following topics:

- [“Troubleshooting Fault Tolerant Virtual Machines,”](#) on page 13
- [“Troubleshooting USB Passthrough Devices,”](#) on page 18
- [“Recover Orphaned Virtual Machines,”](#) on page 19
- [“Virtual Machine Does Not Power On After Cloning or Deploying from Template,”](#) on page 19

Troubleshooting Fault Tolerant Virtual Machines

To maintain a high level of performance and stability for your fault tolerant virtual machines and also to minimize failover rates, you should be aware of certain troubleshooting issues.

The troubleshooting topics discussed focus on problems that you might encounter when using the vSphere Fault Tolerance feature on your virtual machines. The topics also describe how to resolve problems.

You can also see the VMware knowledge base article at <http://kb.vmware.com/kb/1033634> to help you troubleshoot Fault Tolerance. This article contains a list of error messages that you might encounter when you attempt to use the feature and, where applicable, advice on how to resolve each error.

Hardware Virtualization Not Enabled

You must enable Hardware Virtualization (HV) before you use vSphere Fault Tolerance.

Problem

When you attempt to power on a virtual machine with Fault Tolerance enabled, an error message might appear if you did not enable HV.

Cause

This error is often the result of HV not being available on the ESXi server on which you are attempting to power on the virtual machine. HV might not be available either because it is not supported by the ESXi server hardware or because HV is not enabled in the BIOS.

Solution

If the ESXi server hardware supports HV, but HV is not currently enabled, enable HV in the BIOS on that server. The process for enabling HV varies among BIOSes. See the documentation for your hosts' BIOSes for details on how to enable HV.

If the ESXi server hardware does not support HV, switch to hardware that uses processors that support Fault Tolerance.

Compatible Hosts Not Available for Secondary VM

If you power on a virtual machine with Fault Tolerance enabled and no compatible hosts are available for its Secondary VM, you might receive an error message.

Problem

You might encounter the following error message:

Secondary VM could not be powered on as there are no compatible hosts that can accommodate it.

Cause

This can occur for a variety of reasons including that there are no other hosts in the cluster, there are no other hosts with HV enabled, Hardware MMU Virtualization is not supported by host CPUs, data stores are inaccessible, there is no available capacity, or hosts are in maintenance mode.

Solution

If there are insufficient hosts, add more hosts to the cluster. If there are hosts in the cluster, ensure they support HV and that HV is enabled. The process for enabling HV varies among BIOSes. See the documentation for your hosts' BIOSes for details on how to enable HV. Check that hosts have sufficient capacity and that they are not in maintenance mode.

Secondary VM on Overcommitted Host Degrades Performance of Primary VM

If a Primary VM appears to be executing slowly, even though its host is lightly loaded and retains idle CPU time, check the host where the Secondary VM is running to see if it is heavily loaded.

Problem

When a Secondary VM resides on a host that is heavily loaded, the Secondary VM can affect the performance of the Primary VM.

Cause

A Secondary VM running on a host that is overcommitted (for example, with its CPU resources) might not get the same amount of resources as the Primary VM. When this occurs, the Primary VM must slow down to allow the Secondary VM to keep up, effectively reducing its execution speed to the slower speed of the Secondary VM.

Solution

If the Secondary VM is on an overcommitted host, you can move the VM to another location without resource contention problems. Or more specifically, do the following:

- For FT networking contention, use vMotion technology to move the Secondary VM to a host with fewer FT VMs contending on the FT network. Verify that the quality of the storage access to the VM is not asymmetric.
- For storage contention problems, turn FT off and on again. When you recreate the Secondary VM, change its datastore to a location with less resource contention and better performance potential.

- To resolve a CPU resources problem, set an explicit CPU reservation for the Primary VM at an MHz value sufficient to run its workload at the desired performance level. This reservation is applied to both the Primary and Secondary VMs, ensuring that both VMs can execute at a specified rate. For guidance in setting this reservation, view the performance graphs of the virtual machine (before Fault Tolerance was enabled) to see how many CPU resources it used under normal conditions.

Increased Network Latency Observed in FT Virtual Machines

If your FT network is not optimally configured, you might experience latency problems with the FT VMs.

Problem

FT VMs might see a variable increase in packet latency (on the order of milliseconds). Applications that demand very low network packet latency or jitter (for example, certain real-time applications) might see a degradation in performance.

Cause

Some increase in network latency is expected overhead for Fault Tolerance, but certain factors can add to this latency. For example, if the FT network is on a particularly high latency link, this latency is passed on to the applications. Also, if the FT network has insufficient bandwidth (fewer than 10 Gbps), greater latency might occur.

Solution

Verify that the FT network has sufficient bandwidth (10 Gbps or more) and uses a low latency link between the Primary VM and Secondary VM. These precautions do not eliminate network latency, but minimize its potential impact.

Some Hosts Are Overloaded with FT Virtual Machines

You might encounter performance problems if your cluster's hosts have an imbalanced distribution of FT VMs.

Problem

Some hosts in the cluster might become overloaded with FT VMs, while other hosts might have unused resources.

Cause

vSphere DRS does not load balance FT VMs (unless they are using legacy FT). This limitation might result in a cluster where hosts are unevenly distributed with FT VMs.

Solution

Manually rebalance the FT VMs across the cluster by using vSphere vMotion. Generally, the fewer FT VMs that are on a host, the better they perform, due to reduced contention for FT network bandwidth and CPU resources.

Losing Access to FT Metadata Datastore

Access to the Fault Tolerance metadata datastore is essential for the proper functioning of an FT VM. Loss of this access can cause a variety of problems.

Problem

These problems include the following:

- FT can terminate unexpectedly.

- If both the Primary VM and Secondary VM cannot access the metadata datastore, the VMs might fail unexpectedly. Typically, an unrelated failure that terminates FT must also occur when access to the FT metadata datastore is lost by both VMs. vSphere HA then tries to restart the Primary VM on a host with access to the metadata datastore.
- The VM might stop being recognized as an FT VM by vCenter Server. This failed recognition can allow unsupported operations such as taking snapshots to be performed on the VM and cause problematic behavior.

Cause

Lack of access to the Fault Tolerance metadata datastore can lead to the undesirable outcomes in the previous list.

Solution

When planning your FT deployment, place the metadata datastore on highly available storage. While FT is running, if you see that the access to the metadata datastore is lost on either the Primary VM or the Secondary VM, promptly address the storage problem before loss of access causes one of the previous problems. If a VM stops being recognized as an FT VM by vCenter Server, do not perform unsupported operations on the VM. Restore access to the metadata datastore. After access is restored for the FT VMs and the refresh period has ended, the VMs are recognizable.

Turning On vSphere FT for Powered-On VM Fails

If you try to turn on vSphere Fault Tolerance for a powered-on VM, this operation can fail.

Problem

When you select **Turn On Fault Tolerance** for a powered-on VM, the operation fails and you see an Unknown error message.

Cause

This operation can fail if the host that the VM is running on has insufficient memory resources to provide fault tolerant protection. vSphere Fault Tolerance automatically tries to allocate a full memory reservation on the host for the VM. Overhead memory is required for fault tolerant VMs and can sometimes expand to 1 to 2 GB. If the powered-on VM is running on a host that has insufficient memory resources to accommodate the full reservation plus the overhead memory, trying to turn on Fault Tolerance fails. Subsequently, the Unknown error message is returned.

Solution

Choose from these solutions:

- Free up memory resources on the host to accommodate the VM's memory reservation and the added overhead.
- Move the VM to a host with ample free memory resources and try again.

FT Virtual Machines not Placed or Evacuated by vSphere DRS

FT virtual machines in a cluster that is enabled with vSphere DRS do not function correctly if Enhanced vMotion Compatibility (EVC) is currently disabled.

Problem

Because EVC is a prerequisite for using DRS with FT VMs, DRS does not place or evacuate them if EVC has been disabled (even if it is later reenabled).

Cause

When EVC is disabled on a DRS cluster, a VM override that disables DRS on an FT VM might be added. Even if EVC is later reenabled, this override is not canceled.

Solution

If DRS does not place or evacuate FT VMs in the cluster, check the VMs for a VM override that is disabling DRS. If you find one, remove the override that is disabling DRS.

NOTE For more information on how to edit or delete VM overrides, see *vSphere Resource Management*.

Fault Tolerant Virtual Machine Failovers

A Primary or Secondary VM can fail over even though its ESXi host has not crashed. In such cases, virtual machine execution is not interrupted, but redundancy is temporarily lost. To avoid this type of failover, be aware of some of the situations when it can occur and take steps to avoid them.

Partial Hardware Failure Related to Storage

This problem can arise when access to storage is slow or down for one of the hosts. When this occurs there are many storage errors listed in the VMkernel log. To resolve this problem you must address your storage-related problems.

Partial Hardware Failure Related to Network

If the logging NIC is not functioning or connections to other hosts through that NIC are down, this can trigger a fault tolerant virtual machine to be failed over so that redundancy can be reestablished. To avoid this problem, dedicate a separate NIC each for vMotion and FT logging traffic and perform vMotion migrations only when the virtual machines are less active.

Insufficient Bandwidth on the Logging NIC Network

This can happen because of too many fault tolerant virtual machines being on a host. To resolve this problem, more broadly distribute pairs of fault tolerant virtual machines across different hosts.

Use a 10-Gbit logging network for FT and verify that the network is low latency.

vMotion Failures Due to Virtual Machine Activity Level

If the vMotion migration of a fault tolerant virtual machine fails, the virtual machine might need to be failed over. Usually, this occurs when the virtual machine is too active for the migration to be completed with only minimal disruption to the activity. To avoid this problem, perform vMotion migrations only when the virtual machines are less active.

Too Much Activity on VMFS Volume Can Lead to Virtual Machine Failovers

When a number of file system locking operations, virtual machine power ons, power offs, or vMotion migrations occur on a single VMFS volume, this can trigger fault tolerant virtual machines to be failed over. A symptom that this might be occurring is receiving many warnings about SCSI reservations in the VMkernel log. To resolve this problem, reduce the number of file system operations or ensure that the fault tolerant virtual machine is on a VMFS volume that does not have an abundance of other virtual machines that are regularly being powered on, powered off, or migrated using vMotion.

Lack of File System Space Prevents Secondary VM Startup

Check whether or not your `/(root)` or `/vmfs/datasource` file systems have available space. These file systems can become full for many reasons, and a lack of space might prevent you from being able to start a new Secondary VM.

Troubleshooting USB Passthrough Devices

Information about feature behavior can help you troubleshoot or avoid potential problems when USB devices are connected to a virtual machine.

Error Message When You Try to Migrate Virtual Machine with USB Devices Attached

Migration with vMotion cannot proceed and issues a confusing error message when you connect multiple USB devices from an ESXi host to a virtual machine and one or more devices are not enabled for vMotion.

Problem

The Migrate Virtual Machine wizard runs a compatibility check before a migration operation begins. If unsupported USB devices are detected, the compatibility check fails and an error message similar to the following appears: Currently connected device 'USB 1' uses backing 'path:1/7/1', which is not accessible.

Cause

To successfully pass vMotion compatibility checks, you must enable all USB devices that are connected to the virtual machine from a host for vMotion. If one or more devices are not enabled for vMotion, migration will fail.

Solution

- 1 Make sure that the devices are not in the process of transferring data before removing them.
- 2 Re-add and enable vMotion for each affected USB device.

Cannot Copy Data From an ESXi Host to a USB Device That Is Connected to the Host

You can connect a USB device to an ESXi host and copy data to the device from the host. For example, you might want to gather the vm-support bundle from the host after the host loses network connectivity. To perform this task, you must stop the USB arbitrator.

Problem

If the USB arbitrator is being used for USB passthrough from an ESXi host to a virtual machine the USB device appears under `lsusb` but does not mount correctly.

Cause

This problem occurs because the nonbootable USB device is reserved for the virtual machine by default. It does not appear on the host's file system, even though `lsusb` can see the device.

Solution

- 1 Stop the `usbarbitrator` service: `/etc/init.d/usbarbitrator stop`
- 2 Physically disconnect and reconnect the USB device.
By default, the device location is `/vmfs/devices/disks/mpx.vmhbaXX:C0:T0:L0`.
- 3 After you reconnect the device, restart the `usbarbitrator` service: `/etc/init.d/usbarbitrator start`
- 4 Restart `hostd` and any running virtual machines to restore access to the passthrough devices in the virtual machine.

What to do next

Reconnect the USB devices to the virtual machine.

Recover Orphaned Virtual Machines

Virtual machines appear with (orphaned) appended to their names.

Problem

Virtual machines that reside on an ESXi host that vCenter Server manages might become orphaned in rare cases. Such virtual machines exist in the vCenter Server database, but the ESXi host no longer recognizes them.

Cause

Virtual machines can become orphaned if a host failover is unsuccessful, or when the virtual machine is unregistered directly on the host. If this situation occurs, move the orphaned virtual machine to another host in the data center on which the virtual machine files are stored.

Solution

- 1 Determine the datastore where the virtual machine configuration (.vmx) file is located.
 - a Select the virtual machine in the vSphere Web Client inventory, and click the **Datastores** tab.
The datastore or datastores where the virtual machine files are stored are displayed.
 - b If more than one datastore is displayed, select each datastore and click the file browser icon to browse for the .vmx file.
 - c Verify the location of the .vmx file.
- 2 Return to the virtual machine in the vSphere Web Client, right-click it, and select **All Virtual Infrastructure Actions > Remove from Inventory**.
- 3 Click **Yes** to confirm the removal of the virtual machine.
- 4 Reregister the virtual machine with vCenter Server.
 - a Right-click the datastore where the virtual machine file is located and select **Register VM**.
 - b Browse to the .vmx file and click **OK**.
 - c Select the location for the virtual machine and click **Next**.
 - d Select the host on which to run the virtual machine and click **Next**.
 - e Click **Finish**.

Virtual Machine Does Not Power On After Cloning or Deploying from Template

Virtual machines do not power on after you complete the clone or deploy from template workflow in the vSphere Web Client.

Problem

When you clone a virtual machine or deploy a virtual machine from a template, you might not be able to power on the virtual machine after creation.

Cause

The swap file size is not reserved when the virtual machine disks are created.

Solution

- Reduce the size of the swap file that is required for the virtual machine. You can do this by increasing the virtual machine memory reservation.
 - a Right-click the virtual machine and select **Edit Settings**.
 - b Select **Virtual Hardware** and click **Memory**.
 - c Use the Reservation drop-down menu to increase the amount of memory allocated to the virtual machine.
 - d Click **OK**.
- Alternatively, you can increase the amount of space available for the swap file by moving other virtual machine disks off the datastore that is being used for the swap file.
 - a Browse to the datastore in the vSphere Web Client object navigator.
 - b Select the **VMs** tab.
 - c For each virtual machine to move, right-click the virtual machine and select **Migrate**.
 - d Select **Change storage only**.
 - e Proceed through the Migrate Virtual Machine wizard.
- You can also increase the amount of space available for the swap file by changing the swap file location to a datastore with adequate space.
 - a Browse to the host in the vSphere Web Client object navigator.
 - b Select the **Configure** tab.
 - c Under Virtual Machines, select **Swap file location**.
 - d Click **Edit**.

NOTE If the host is part of a cluster that specifies that the virtual machine swap files are stored in the same directory as the virtual machine, you cannot click **Edit**. You must use the Cluster Settings dialog box to change the swap file location policy for the cluster.

- e Select **Use a specific datastore** and select a datastore from the list.
- f Click **OK**.

Troubleshooting Hosts

The host troubleshooting topics provide solutions to potential problems that you might encounter when using your vCenter Servers and ESXi hosts.

This chapter includes the following topics:

- [“Troubleshooting vSphere HA Host States,”](#) on page 21
- [“Troubleshooting vSphere Auto Deploy,”](#) on page 25
- [“Authentication Token Manipulation Error,”](#) on page 31
- [“Active Directory Rule Set Error Causes Host Profile Compliance Failure,”](#) on page 32
- [“Unable to Download VIBs When Using vCenter Server Reverse Proxy,”](#) on page 32

Troubleshooting vSphere HA Host States

vCenter Server reports vSphere HA host states that indicate an error condition on the host. Such errors can prevent vSphere HA from fully protecting the virtual machines on the host and can impede vSphere HA's ability to restart virtual machines after a failure. Errors can occur when vSphere HA is being configured or unconfigured on a host or, more rarely, during normal operation. When this happens, you should determine how to resolve the error, so that vSphere HA is fully operational.

vSphere HA Agent Is in the Agent Unreachable State

The vSphere HA agent on a host is in the Agent Unreachable state for a minute or more. User intervention might be required to resolve this situation.

Problem

vSphere HA reports that an agent is in the Agent Unreachable state when the agent for the host cannot be contacted by the master host or by vCenter Server. Consequently, vSphere HA is not able to monitor the virtual machines on the host and might not restart them after a failure.

Cause

A vSphere HA agent can be in the Agent Unreachable state for several reasons. This condition most often indicates that a networking problem is preventing vCenter Server or the master host from contacting the agent on the host, or that all hosts in the cluster have failed. This condition can also indicate the unlikely situation that vSphere HA was disabled and then re-enabled on the cluster while vCenter Server could not communicate with the vSphere HA agent on the host, or that the ESXi host agent on the host has failed, and the watchdog process was unable to restart it. In any of these cases, a failover event is not triggered when a host goes into the Unreachable state.

Solution

Determine if vCenter Server is reporting the host as not responding. If so, there is a networking problem, an ESXi host agent failure, or a total cluster failure. After the condition is resolved, vSphere HA should work correctly. If not, reconfigure vSphere HA on the host. Similarly, if vCenter Server reports the hosts are responding but a host's state is Agent Unreachable, reconfigure vSphere HA on that host.

vSphere HA Agent is in the Uninitialized State

The vSphere HA agent on a host is in the Uninitialized state for a minute or more. User intervention might be required to resolve this situation.

Problem

vSphere HA reports that an agent is in the Uninitialized state when the agent for the host is unable to enter the run state and become the master host or to connect to the master host. Consequently, vSphere HA is not able to monitor the virtual machines on the host and might not restart them after a failure.

Cause

A vSphere HA agent can be in the Uninitialized state for one or more reasons. This condition most often indicates that the host does not have access to any datastores. Less frequently, this condition indicates that the host does not have access to its local datastore on which vSphere HA caches state information, the agent on the host is inaccessible, or the vSphere HA agent is unable to open required firewall ports. It is also possible that the ESXi host agent has stopped.

Solution

Search the list of the host's events for recent occurrences of the event `vSphere HA Agent for the host has an error`. This event indicates the reason for the host being in the uninitialized state. If the condition exists because of a datastore problem, resolve whatever is preventing the host from accessing the affected datastores. If the ESXi host agent has stopped, you must restart it. After the problem has been resolved, if the agent does not return to an operational state, reconfigure vSphere HA on the host.

NOTE If the condition exists because of a firewall problem, check if there is another service on the host that is using port 8182. If so, shut down that service, and reconfigure vSphere HA.

vSphere HA Agent is in the Initialization Error State

The vSphere HA agent on a host is in the Initialization Error state for a minute or more. User intervention is required to resolve this situation.

Problem

vSphere HA reports that an agent is in the Initialization Error state when the last attempt to configure vSphere HA for the host failed. vSphere HA does not monitor the virtual machines on such a host and might not restart them after a failure.

Cause

This condition most often indicates that vCenter Server was unable to connect to the host while the vSphere HA agent was being installed or configured on the host. This condition might also indicate that the installation and configuration completed, but the agent did not become a master host or a slave host within a timeout period. Less frequently, the condition is an indication that there is insufficient disk space on the host's local datastore to install the agent, or that there are insufficient unreserved memory resources on the host for the agent resource pool. Finally, for ESXi 5.x hosts, the configuration fails if a previous installation of another component required a host reboot, but the reboot has not yet occurred.

Solution

When a Configure HA task fails, a reason for the failure is reported.

Reason for Failure	Action
Host communication errors	Resolve any communication problems with the host and retry the configuration operation.
Timeout errors	Possible causes include that the host crashed during the configuration task, the agent failed to start after being installed, or the agent was unable to initialize itself after starting up. Verify that vCenter Server is able to communicate with the host. If so, see “vSphere HA Agent Is in the Agent Unreachable State,” on page 21 or “vSphere HA Agent is in the Uninitialized State,” on page 22 for possible solutions.
Lack of resources	Free up approximately 75MB of disk space. If the failure is due to insufficient unreserved memory, free up memory on the host by either relocating virtual machines to another host or reducing their reservations. In either case, retry the vSphere HA configuration task after resolving the problem.
Reboot pending	If an installation for a 5.0 or later host fails because a reboot is pending, reboot the host and retry the vSphere HA configuration task.

vSphere HA Agent is in the Uninitialization Error State

The vSphere HA agent on a host is in the Uninitialization Error state. User intervention is required to resolve this situation.

Problem

vSphere HA reports that an agent is in the Uninitialization Error state when vCenter Server is unable to unconfigure the agent on the host during the Unconfigure HA task. An agent left in this state can interfere with the operation of the cluster. For example, the agent on the host might elect itself as master host and lock a datastore. Locking a datastore prevents the valid cluster master host from managing the virtual machines with configuration files on that datastore.

Cause

This condition usually indicates that vCenter Server lost the connection to the host while the agent was being unconfigured.

Solution

Add the host back to vCenter Server (version 5.0 or later). The host can be added as a stand-alone host or added to any cluster.

vSphere HA Agent is in the Host Failed State

The vSphere HA agent on a host is in the Host Failed state. User intervention is required to resolve the situation.

Problem

Usually, such reports indicate that a host has actually failed, but failure reports can sometimes be incorrect. A failed host reduces the available capacity in the cluster and, in the case of an incorrect report, prevents vSphere HA from protecting the virtual machines running on the host.

Cause

This host state is reported when the vSphere HA master host to which vCenter Server is connected is unable to communicate with the host and with the heartbeat datastores that are in use for the host. Any storage failure that makes the datastores inaccessible to hosts can cause this condition if accompanied by a network failure.

Solution

Check for the noted failure conditions and resolve any that are found.

vSphere HA Agent is in the Network Partitioned State

The vSphere HA agent on a host is in the Network Partitioned state. User intervention might be required to resolve this situation.

Problem

While the virtual machines running on the host continue to be monitored by the master hosts that are responsible for them, vSphere HA's ability to restart the virtual machines after a failure is affected. First, each master host has access to a subset of the hosts, so less failover capacity is available to each host. Second, vSphere HA might be unable to restart a FT Secondary VM after a failure (see [“Primary VM Remains in the Need Secondary State,”](#) on page 44).

Cause

A host is reported as partitioned if both of the following conditions are met:

- The vSphere HA master host to which vCenter Server is connected is unable to communicate with the host by using the management (or VMware vSAN™) network, but is able to communicate with that host by using the heartbeat datastores that have been selected for it.
- The host is not isolated.

A network partition can occur for a number of reasons including incorrect VLAN tagging, the failure of a physical NIC or switch, configuring a cluster with some hosts that use only IPv4 and others that use only IPv6, or the management networks for some hosts were moved to a different virtual switch without first putting the host into maintenance mode.

Solution

Resolve the networking problem that prevents the hosts from communicating by using the management networks.

vSphere HA Agent is in the Network Isolated State

The vSphere HA agent on a host is in the Network Isolated state. User intervention is required to resolve this situation.

Problem

When a host is in the Network Isolated state, there are two things to consider -- the isolated host and the vSphere HA agent that holds the master role.

- On the isolated host, the vSphere HA agent applies the configured isolation response to the running VMs, determining if they should be shut down or powered off. It does this after checking whether a master agent is able to take responsibility for each VM (by locking the VM's home datastore.) If not, the agent defers applying the isolation response for the VM and rechecks the datastore state after a short delay.
- If the vSphere HA master agent can access one or more of the datastores, it monitors the VMs that were running on the host when it became isolated and attempts to restart any that were powered off or shut down.

Cause

A host is network isolated if both of the following conditions are met:

- Isolation addresses have been configured and the host is unable to ping them.

- The vSphere HA agent on the host is unable to access any of the agents running on the other cluster hosts.

NOTE If your vSphere HA cluster has vSAN enabled, a host is determined to be isolated if it cannot communicate with the other vSphere HA agents in the cluster and cannot reach the configured isolation addresses. Although the vSphere HA agents use the vSAN network for inter-agent communication, the default isolation address is still the gateway of the host. Hence, in the default configuration, both networks must fail for a host to be declared isolated.

Solution

Resolve the networking problem that is preventing the host from pinging its isolation addresses and communicating with other hosts.

Configuration of vSphere HA on Hosts Times Out

The configuration of a vSphere HA cluster might time out on some of the hosts added to it.

Problem

When you enable vSphere HA on an existing cluster with a large number of hosts and virtual machines, the setup of vSphere HA on some of the hosts might fail.

Cause

This failure is the result of a time out occurring before the installation of vSphere HA on the host(s) completes.

Solution

Set the vCenter Server advanced option `config.vpxd.das.electionWaitTimeSec` to `value=240`. Once this change is made, the time outs do not occur.

Troubleshooting vSphere Auto Deploy

The vSphere Auto Deploy troubleshooting topics offer solutions for situations when provisioning hosts with vSphere Auto Deploy does not work as expected.

vSphere Auto Deploy TFTP Timeout Error at Boot Time

A TFTP Timeout error message appears when a host provisioned with vSphere Auto Deploy boots. The text of the message depends on the BIOS.

Problem

A TFTP Timeout error message appears when a host provisioned with vSphere Auto Deploy boots. The text of the message depends on the BIOS.

Cause

The TFTP server is down or unreachable.

Solution

- Ensure that your TFTP service is running and reachable by the host that you are trying to boot.
- To view the diagnostic logs for details on the present error, see your TFTP service documentation.

vSphere Auto Deploy Host Boots with Wrong Configuration

A host is booting with a different ESXi image, host profile, or folder location than the one specified in the rules.

Problem

A host is booting with a different ESXi image profile or configuration than the image profile or configuration that the rules specify. For example, you change the rules to assign a different image profile, but the host still uses the old image profile.

Cause

After the host has been added to a vCenter Server system, the boot configuration is determined by the vCenter Server system. The vCenter Server system associates an image profile, host profile, or folder location with the host.

Solution

- ◆ Use the `Test-DeployRuleSetCompliance` and `Repair-DeployRuleSetCompliance` vSphere PowerCLI cmdlets to reevaluate the rules and to associate the correct image profile, host profile, or folder location with the host.

Host Is Not Redirected to vSphere Auto Deploy Server

During boot, a host that you want to provision with vSphere Auto Deploy loads iPXE. The host is not redirected to the vSphere Auto Deploy server.

Problem

During boot, a host that you want to provision with vSphere Auto Deploy loads iPXE. The host is not redirected to the vSphere Auto Deploy server.

Cause

The `tramp` file that is included in the TFTP ZIP file has the wrong IP address for the vSphere Auto Deploy server.

Solution

- ◆ Correct the IP address of the vSphere Auto Deploy server in the `tramp` file, as explained in the *vSphere Installation and Setup* documentation.

Package Warning Message When You Assign an Image Profile to a vSphere Auto Deploy Host

When you run a vSphere PowerCLI cmdlet that assigns an image profile that is not vSphere Auto Deploy ready, a warning message appears.

Problem

When you write or modify rules to assign an image profile to one or more hosts, the following error results:

Warning: Image Profile <name-here> contains one or more software packages that are not stateless-ready. You may experience problems when using this profile with Auto Deploy.

Cause

Each VIB in an image profile has a `stateless-ready` flag that indicates that the VIB is meant for use with vSphere Auto Deploy. You get the error if you attempt to write a vSphere Auto Deploy rule that uses an image profile in which one or more VIBs have that flag set to `FALSE`.

NOTE You can use hosts provisioned with vSphere Auto Deploy that include VIBs that are not `stateless-ready` without problems. However booting with an image profile that includes VIBs that are not `stateless-ready` is treated like a fresh install. Each time you boot the host, you lose any configuration data that would otherwise be available across reboots for hosts provisioned with vSphere Auto Deploy.

Solution

- 1 Use vSphere ESXi Image Builder cmdlets in a vSphere PowerCLI session to view the VIBs in the image profile.
- 2 Remove any VIBs that are not `stateless-ready`.
- 3 Rerun the vSphere Auto Deploy cmdlet.

vSphere Auto Deploy Host with a Built-In USB Flash Drive Does Not Send Coredumps to Local Disk

If your vSphere Auto Deploy host has a built-in USB flash drive, and an error results in a coredump, the coredump is lost. Set up your system to use ESXi Dump Collector to store coredumps on a networked host.

Problem

If your vSphere Auto Deploy host has a built-in USB Flash, and if it encounters an error that results in a coredump, the coredump is not sent to the local disk.

Solution

- 1 Install ESXi Dump Collector on a system of your choice.
ESXi Dump Collector is included with the vCenter Server installer.
- 2 Use ESXCLI to configure the host to use ESXi Dump Collector.

```
esxcli conn_options system coredump network set IP-addr,port
esxcli system coredump network set -e true
```
- 3 Use ESXCLI to disable local coredump partitions.

```
esxcli conn_options system coredump partition set -e false
```

vSphere Auto Deploy Host Reboots After Five Minutes

A vSphere Auto Deploy host boots and displays iPXE information, but reboots after five minutes.

Problem

A host to be provisioned with vSphere Auto Deploy boots from iPXE and displays iPXE information on the console. However, after five minutes, the host displays the following message to the console and reboots.

```
This host is attempting to network-boot using VMware
AutoDeploy. However, there is no ESXi image associated with this host.
Details: No rules containing an Image Profile match this
host. You can create a rule with the New-DeployRule PowerCLI cmdlet
and add it to the rule set with Add-DeployRule or Set-DeployRuleSet.
The rule should have a pattern that matches one or more of the attributes
listed below.
```

The host might also display the following details:

Details: This host has been added to VC, but no Image Profile is associated with it. You can use `Apply-ESXImageProfile` in the PowerCLI to associate an Image Profile with this host. Alternatively, you can reevaluate the rules for this host with the `Test-DeployRuleSetCompliance` and `Repair-DeployRuleSetCompliance` cmdlets.

The console then displays the host's machine attributes including vendor, serial number, IP address, and so on.

Cause

No image profile is currently associated with this host.

Solution

You can assign an image profile to the host by running the `Apply-EsxImageProfile` cmdlet, or by creating the following rule:

- 1 Run the `New-DeployRule` cmdlet to create a rule that includes a pattern that matches the host with an image profile.
- 2 Run the `Add-DeployRule` cmdlet to add the rule to a ruleset.
- 3 Run the `Test-DeployRuleSetCompliance` cmdlet and use the output of that cmdlet as the input to the `Repair-DeployRuleSetCompliance` cmdlet.

vSphere Auto Deploy Host Cannot Contact TFTP Server

The host that you provision with vSphere Auto Deploy cannot contact the TFTP server.

Problem

When you attempt to boot a host provisioned with vSphere Auto Deploy, the host performs a network boot and is assigned a DHCP address by the DHCP server, but the host cannot contact the TFTP server.

Cause

The TFTP server might have stopped running, or a firewall might block the TFTP port.

Solution

- If you installed the WinAgents TFTP server, open the WinAgents TFTP management console and verify that the service is running. If the service is running, check the Windows firewall's inbound rules to make sure the TFTP port is not blocked. Turn off the firewall temporarily to see whether the firewall is the problem.
- For all other TFTP servers, see the server documentation for debugging procedures.

vSphere Auto Deploy Host Cannot Retrieve ESXi Image from vSphere Auto Deploy Server

The host that you provision with vSphere Auto Deploy stops at the iPXE boot screen.

Problem

When you attempt to boot a host provisioned with vSphere Auto Deploy, the boot process stops at the iPXE boot screen and the status message indicates that the host is attempting to get the ESXi image from the vSphere Auto Deploy server.

Cause

The vSphere Auto Deploy service might be stopped or the vSphere Auto Deploy server might be inaccessible.

Solution

- 1 Log in to the system on which you installed the vSphere Auto Deploy server.
- 2 Check that the vSphere Auto Deploy server is running.
 - a Click **Start > Settings > Control Panel > Administrative Tools**.
 - b Double-click **Services** to open the Services Management panel.
 - c In the Services field, look for the VMware vSphere Auto Deploy Waiter service and restart the service if it is not running.
- 3 Open a Web browser, enter the following URL, and check whether the vSphere Auto Deploy server is accessible.

`https://Auto_Deploy_Server_IP_Address:Auto_Deploy_Server_Port/vmw/rdb`

NOTE Use this address only to check whether the server is accessible.

- 4 If the server is not accessible, a firewall problem is likely.
 - a Try setting up permissive TCP Inbound rules for the vSphere Auto Deploy server port.
The port is 6501 unless you specified a different port during installation.
 - b As a last resort, disable the firewall temporarily and enable it again after you verified whether it blocked the traffic. Do not disable the firewall on production environments.

To disable the firewall, run `netsh firewall set opmode disable`. To enable the firewall, run `netsh firewall set opmode enable`.

vSphere Auto Deploy Host Does Not Get a DHCP Assigned Address

The host you provision with vSphere Auto Deploy fails to get a DHCP Address.

Problem

When you attempt to boot a host provisioned with vSphere Auto Deploy, the host performs a network boot but is not assigned a DHCP address. The vSphere Auto Deploy server cannot provision the host with the image profile.

Cause

You might have a problem with the DHCP service or with the firewall setup.

Solution

- 1 Check that the DHCP server service is running on the Windows system on which the DHCP server is set up to provision hosts.
 - a Click **Start > Settings > Control Panel > Administrative Tools**.
 - b Double-click **Services** to open the Services Management panel.
 - c In the Services field, look for the DHCP server service and restart the service if it is not running.
- 2 If the DHCP server is running, recheck the DHCP scope and the DHCP reservations that you configured for your target hosts.

If the DHCP scope and reservations are configured correctly, the problem most likely involves the firewall.

- 3 As a temporary workaround, turn off the firewall to see whether that resolves the problem.
 - a Open the command prompt by clicking **Start > Program > Accessories > Command prompt**.
 - b Type the following command to temporarily turn off the firewall. Do not turn off the firewall in a production environment.


```
netsh firewall set opmode disable
```
 - c Attempt to provision the host with vSphere Auto Deploy.
 - d Type the following command to turn the firewall back on.


```
netsh firewall set opmode enable
```
- 4 Set up rules to allow DHCP network traffic to the target hosts.

See the firewall documentation for DHCP and for the Windows system on which the DHCP server is running for details.

vSphere Auto Deploy Host Does Not Network Boot

The host you provision with vSphere Auto Deploy comes up but does not network boot.

Problem

When you attempt to boot a host provisioned with vSphere Auto Deploy, the host does not start the network boot process.

Cause

You did not enable your host for network boot.

Solution

- 1 Reboot the host and follow the on-screen instructions to access the BIOS configuration.
- 2 In the BIOS configuration, enable Network Boot in the Boot Device configuration.

Recovering from Database Corruption on the vSphere Auto Deploy Server

In some situations, you might have a problem with the vSphere Auto Deploy database. The most efficient recovery option is to replace the existing database file with the most recent backup.

Problem

When you use vSphere Auto Deploy to provision the ESXi hosts in your environment, you might encounter a problem with the vSphere Auto Deploy database.

IMPORTANT This is a rare problem. Follow all other vSphere Auto Deploy troubleshooting strategies before you replace the current database file. Rules or associations that you created since the backup you choose are lost.

Cause

This problem happens only with hosts that are provisioned with vSphere Auto Deploy.

Solution

- 1 Stop the vSphere Auto Deploy server service.
- 2 Find the vSphere Auto Deploy log by going to the vSphere Auto Deploy page in the vSphere Web Client.

- 3 Check the logs for the following message:

DatabaseError: database disk image is malformed.

If you see the message, replace the existing database with the most recent backup.

- 4 Go to the vSphere Auto Deploy data directory.

Operating System	File Location
vCenter Server appliance	/var/lib/rbd
Microsoft Windows	%VMWARE_DATA_DIR%\autodeploy\Data

The directory contains a file named `db`, and backup files named `db-yyy-mm-dd`.

- 5 Rename the current `db` file.
VMware Support might ask for that file if you call for assistance.
- 6 Rename the most recent backup to `db`.
- 7 Restart the vSphere Auto Deploy server service.
- 8 If the message still appears in the log, repeat the steps to use the next recent backup until vSphere Auto Deploy works without database errors.

Authentication Token Manipulation Error

Creating a password that does not meet the authentication requirements of the host causes an error.

Problem

When you create a password on the host, the following fault message appears: A general system error occurred: passwd: Authentication token manipulation error.

The following message is included: Failed to set the password. It is possible that your password does not meet the complexity criteria set by the system.

Cause

The host checks for password compliance using the default authentication plug-in, `pam_passwdqc.so`. If the password is not compliant, the error appears.

Solution

When you create a password, include a mix of characters from four character classes: lowercase letters, uppercase letters, numbers, and special characters such as an underscore or dash.

Your user password must meet the following length requirements.

- Passwords containing characters from three character classes must be at least eight characters long.
- Passwords containing characters from all four character classes must be at least seven characters long.

NOTE An uppercase character that begins a password does not count toward the number of character classes used. A number that ends a password does not count toward the number of character classes used.

For more information, see the *vSphere Security* documentation.

Active Directory Rule Set Error Causes Host Profile Compliance Failure

Applying a host profile that specifies an Active Directory domain to join causes a compliance failure.

Problem

When you apply a host profile that specifies an Active Directory domain to join, but you do not enable the **activeDirectoryAll** rule set in the firewall configuration, a compliance failure occurs. The vSphere Web Client displays the error message **Failures against the host profile: Ruleset activedirectoryAll does not match the specification**. The compliance failure also occurs when you apply a host profile to leave an Active Directory domain, but you do not disable the **activeDirectoryAll** rule set in the host profile.

Cause

Active Directory requires the **activeDirectoryAll** firewall rule set. You must enable the rule set in the firewall configuration. If you omit this setting, the system adds the necessary firewall rules when the host joins the domain, but the host will be noncompliant because of the mismatch in firewall rules. The host will also be noncompliant if you remove it from the domain without disabling the Active Directory rule set.

Solution

- 1 Browse to the host profile in the vSphere Web Client.
To find a host profile, click **Policies and Profiles > Host Profiles** on the vSphere Web Client Home page.
- 2 Right-click the host profile and select **Edit Settings**.
- 3 Click **Next**.
- 4 Select **Security and Services > Firewall Configuration > Firewall configuration > Ruleset Configuration**.
- 5 Ensure that **activeDirectoryAll** is selected.
- 6 In the right panel, select the **Flag indicating whether ruleset should be enabled** check box.
Deselect the check box if the host is leaving the domain.
- 7 Click **Next**, and then click **Finish** to complete the change to the host profile.

Unable to Download VIBs When Using vCenter Server Reverse Proxy

You are unable to download VIBs if vCenter Server is using a custom port for the reverse proxy.

Problem

If you configure vCenter Server reverse proxy to use a custom port, the VIB downloads fail.

Cause

If vCenter Server is using a custom port for the reverse proxy, the custom port is not automatically enabled in the ESXi firewall and the VIB downloads fail.

Solution

- 1 Open an SSH connection to the host and log in as root.
- 2 (Optional) List the existing firewall rules.
`esxcli network firewall ruleset list`

- 3 (Optional) Back up the `/etc/vmware/firewall/service.xml` file.


```
cp /etc/vmware/firewall/service.xml /etc/vmware/firewall/service.xml.bak
```
- 4 Edit the access permissions of the `service.xml` file to allow writes by running the `chmod` command.
 - To allow writes, run `chmod 644/etc/vmware/firewall/service.xml`.
 - To toggle the sticky bit flag, run `chmod +t /etc/vmware/firewall/service.xml`.
- 5 Open the `service.xml` file in a text editor.
- 6 Add a new rule to the `service.xml` file that enables the custom port for the vCenter Server reverse proxy .

```
<service id='id_value'>
  <id>vcenterhttpproxy</id>
  <rule id='0000'>
    <direction>outbound</direction>
    <protocol>tcp</protocol>
    <port type='dst'>custom_reverse_proxy_port</port>
  </rule>
  <enabled>true</enabled>
  <required>false</required>
</service>
```

Where `id_value` must be a unique value, for example, if the last listed service in the `service.xml` file has ID 0040, you must enter id number 0041.

- 7 Revert the access permissions of the `service.xml` file to the default read-only setting.


```
chmod 444 /etc/vmware/firewall/service.xml
```
- 8 Refresh the firewall rules for the changes to take effect.


```
esxcli network firewall refresh
```
- 9 (Optional) List the updated rule set to confirm the change.


```
esxcli network firewall ruleset list
```
- 10 (Optional) If you want the firewall configuration to persist after a reboot of the ESXi host, copy the `service.xml` onto persistent storage and modify the `local.sh` file.
 - a Copy the modified `service.xml` file onto persistent storage, for example `/store/`, or onto a VMFS volume, for example `/vmfs/volumes/volume/`.


```
cp /etc/vmware/firewall/service.xml location_of_xml_file
```

You can store a VMFS volume in a single location and copy it to multiple hosts.
 - b Add the `service.xml` file information to the `local.sh` file on the host.


```
cp location_of_xml_file /etc/vmware/firewall
esxcli network firewall refresh
```

Where `location_of_xml_file` is the location to which the file was copied.

Troubleshooting vCenter Server and the vSphere Web Client

4

The vCenter Server and vSphere Web Client troubleshooting topics provide solutions to problems you might encounter when you set up and configure vCenter Server and the vSphere Web Client, including vCenter Single Sign-On.

This chapter includes the following topics:

- [“Troubleshooting vCenter Server,”](#) on page 35
- [“Troubleshooting the vSphere Web Client,”](#) on page 36
- [“Troubleshooting vCenter Server and ESXi Host Certificates,”](#) on page 37

Troubleshooting vCenter Server

These troubleshooting topics provide solutions to problems you might encounter when you use install vCenter Server on the Windows operating system or deploy the vCenter Server Appliance on a Linux system.

vCenter Server Upgrade Fails When Unable to Stop Tomcat Service

A vCenter Server upgrade can fail when the installer is unable to stop the Tomcat service.

Problem

If the vCenter Server installer cannot stop the Tomcat service during an upgrade, the upgrade fails with an error message similar to `Unable to delete VC Tomcat service`. This problem can occur even if you stop the Tomcat service manually before the upgrade, if some files that are used by the Tomcat process are locked.

Solution

- 1 From the Windows **Start** menu, select **Settings > Control Panel > Administrative Tools > Services**.
- 2 Right-click **VMware VirtualCenter Server** and select **Manual**.
- 3 Right-click **VMware vCenter Management Webservices** and select **Manual**.
- 4 Reboot the vCenter Server machine before upgrading.

This releases any locked files that are used by the Tomcat process, and enables the vCenter Server installer to stop the Tomcat service for the upgrade.

Alternatively, you can restart the vCenter Server machine and restart the upgrade process, but select the option not to overwrite the vCenter Server data.

Microsoft SQL Database Set to Unsupported Compatibility Mode Causes vCenter Server Installation or Upgrade to Fail

vCenter Server installation with a Microsoft SQL database fails when the database is set to compatibility mode with an unsupported version.

Problem

The following error message appears: The DB User entered does not have the required permissions needed to install and configure vCenter Server with the selected DB. Please correct the following error(s): %s

Cause

The database version must be supported for vCenter Server. For SQL, even if the database is a supported version, if it is set to run in compatibility mode with an unsupported version, this error occurs. For example, if SQL 2008 is set to run in SQL 2000 compatibility mode, this error occurs.

Solution

- ◆ Make sure the vCenter Server database is a supported version and is not set to compatibility mode with an unsupported version. See the VMware Product Interoperability Matrixes at http://partnerweb.vmware.com/comp_guide2/sim/interop_matrix.php?

Troubleshooting the vSphere Web Client

The vSphere Web Client topics provide solutions to potential problems you might encounter when using the vSphere Web Client to manage vSphere components, including vCenter Single Sign-On and vCenter Server.

vCenter Server System Does Not Appear in vSphere Web Client Inventory

The vSphere Web Client does not display the vCenter Server systems that you expect to see in the inventory.

Problem

When you log in to the vSphere Web Client, the inventory appears to be empty or the vCenter Server system you expected to see does not appear.

Cause

In vSphere 5.1 and later, you log into the vSphere Web Client to view and manage multiple instances of vCenter Server. Any vCenter Server system on which you have permissions appears in the inventory, if the server is registered with the same Component Manager as the vSphere Web Client.

Solution

- Log in to the vSphere Web Client as a user with permissions on the vCenter Server system.
The vCenter Server system will not appear in the inventory if you do not have permissions on it. For example, if you log in as the vCenter Single Sign On administrator user, you might not have permissions on any vCenter Server system.
- Verify that the vCenter Server system is registered with the same Component Manager as the vSphere Web Client.

The vSphere Web Client discovers only vCenter Server systems that are registered with the same Component Manager.

Unable to Start the Virtual Machine Console

When you attempt to open a virtual machine console from the vSphere Web Client, the console does not open.

Problem

When you attempt to open a virtual machine console from the vSphere Web Client, the console does not open. The following error message appears:

```
HTTP ERROR 404
```

```
Problem accessing /. Reason:
```

```
Not Found
```

Errors similar to the following appear in the `virgo-server.log` file:

```
[2012-10-03 18:34:19.170] [ERROR] Thread-40
System.err
                                2012-10-03
18:34:19.167:WARN:oejuc.AbstractLifeCycle:FAILED org.eclipse.jetty.server.Server@315b0333:
java.net.BindException: Address already in use
[2012-10-03 18:34:19.170] [ERROR] Thread-40 System.err java.net.BindException: Address already
in use
```

Cause

Another program or process is using port 9443, the default port used by the HTML5 virtual machine console.

Solution

- ◆ Edit the `webclient.properties` file to add the line `html.console.port=port`, where `port` is the new port number.

The `webclient.properties` file is located in one of the following locations, depending on the operating system on the machine on which the vSphere Web Client is installed:

Windows 2008	C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\
vCenter Server Appliance	/etc/vmware/vsphere-client/

Troubleshooting vCenter Server and ESXi Host Certificates

Certificates are automatically generated when you install vCenter Server. These default certificates are not signed by a commercial certificate authority (CA) and might not provide strong security. You can replace default vCenter Server certificates with certificates signed by a commercial CA. When you replace vCenter Server and ESXi certificates, you might encounter errors.

vCenter Server Cannot Connect to Managed Hosts

After you replace default vCenter Server certificates and restart the system, vCenter Server might not be able to connect to managed hosts.

Problem

vCenter Server cannot connect to managed hosts after server certificates are replaced and the system is restarted.

Solution

Log into the host as the root user and reconnect the host to vCenter Server.

New vCenter Server Certificate Does Not Appear to Load

After you replace default vCenter Server certificates, the new certificates might not appear to load.

Problem

When you install new vCenter Server certificates, you might not see the new certificate.

Cause

Existing open connections to vCenter Server are not forcibly closed and might still use the old certificate.

Solution

To force all connections to use the new certificate, use one of the following methods.

- Restart the network stack or network interfaces on the server.
- Restart the vCenter Server service.

Cannot Configure vSphere HA When Using Custom SSL Certificates

After you install custom SSL certificates, attempts to enable vSphere High Availability (HA) fail.

Problem

When you attempt to enable vSphere HA on a host with custom SSL certificates installed, the following error message appears: vSphere HA cannot be configured on this host because its SSL thumbprint has not been verified.

Cause

When you add a host to vCenter Server, and vCenter Server already trusts the host's SSL certificate, `VPX_HOST.EXPECTED_SSL_THUMBPRINT` is not populated in the vCenter Server database. vSphere HA obtains the host's SSL thumbprint from this field in the database. Without the thumbprint, you cannot enable vSphere HA.

Solution

- 1 In the vSphere Web Client, disconnect the host that has custom SSL certificates installed.
- 2 Reconnect the host to vCenter Server.
- 3 Accept the host's SSL certificate.
- 4 Enable vSphere HA on the host.

Troubleshooting Availability

The availability troubleshooting topics provide solutions to potential problems that you might encounter when using your hosts and datastores in vSphere HA clusters.

You might get an error message when you try to use vSphere HA or vSphere FT. For information about these error messages, see the VMware knowledge base article at <http://kb.vmware.com/kb/1033634>.

This chapter includes the following topics:

- [“Troubleshooting vSphere HA Admission Control,”](#) on page 39
- [“Troubleshooting Heartbeat Datastores,”](#) on page 41
- [“Troubleshooting vSphere HA Failure Response,”](#) on page 42
- [“Troubleshooting vSphere Fault Tolerance in Network Partitions,”](#) on page 44
- [“Troubleshooting VM Component Protection,”](#) on page 45

Troubleshooting vSphere HA Admission Control

vCenter Server uses admission control to ensure that sufficient resources in a vSphere HA cluster are reserved for virtual machine recovery in the event of host failure.

If vSphere HA admission control does not function properly, there is no assurance that all virtual machines in the cluster can be restarted after a host failure.

Red Cluster Due to Insufficient Failover Resources

When you use the Host Failures Cluster Tolerates admission control policy, vSphere HA clusters might become invalid (red) due to insufficient failover resources.

Problem

If you select the Host Failures Cluster Tolerates admission control policy and certain problems arise, the cluster turns red.

Cause

This problem can arise when hosts in the cluster are disconnected, in maintenance mode, not responding, or have a vSphere HA error. Disconnected and maintenance mode hosts are typically caused by user action. Unresponsive or error-possessing hosts usually result from a more serious problem, for example, hosts or agents have failed or a networking problem exists.

Another possible cause of this problem is if your cluster contains any virtual machines that have much larger memory or CPU reservations than the others. The Host Failures Cluster Tolerates admission control policy is based on the calculation on a slot size consisting of two components, the CPU and memory reservations of a virtual machine. If the calculation of this slot size is skewed by outlier virtual machines, the admission control policy can become too restrictive and result in a red cluster. In this case, you can use the vSphere HA advanced options to reduce the slot size, use a different admission control policy, or modify the policy to tolerate fewer host failures.

Solution

Check that all hosts in the cluster are healthy, that is, connected, not in maintenance mode and free of vSphere HA errors. vSphere HA admission control only considers resources from healthy hosts.

Unable to Power On Virtual Machine Due to Insufficient Failover Resources

You might get a `not enough failover resources` fault when trying to power on a virtual machine in a vSphere HA cluster.

Problem

If you select the Host Failures Cluster Tolerates admission control policy and certain problems arise, you might be prevented from powering on a virtual machine due to insufficient resources.

Cause

This problem can have several causes.

- Hosts in the cluster are disconnected, in maintenance mode, not responding, or have a vSphere HA error.

Disconnected and maintenance mode hosts are typically caused by user action. Unresponsive or error-possessing hosts usually result from a more serious problem, for example, hosts or agents have failed or a networking problem exists).

- Cluster contains virtual machines that have much larger memory or CPU reservations than the others.

The Host Failures Cluster Tolerates admission control policy is based on the calculation on a slot size comprised of two components, the CPU and memory reservations of a virtual machine. If the calculation of this slot size is skewed by outlier virtual machines, the admission control policy can become too restrictive and result in the inability to power on virtual machines.

- No free slots in the cluster.

Problems occur if there are no free slots in the cluster or if powering on a virtual machine causes the slot size to increase because it has a larger reservation than existing virtual machines. In either case, you should use the vSphere HA advanced options to reduce the slot size, use a different admission control policy, or modify the policy to tolerate fewer host failures.

Solution

View the **Advanced Runtime Info** pane that appears in the vSphere HA section of the cluster's **Monitor** tab in the vSphere Web Client. This information pane shows the slot size and how many available slots there are in the cluster. If the slot size appears too high, click on the **Resource Allocation** tab of the cluster and sort the virtual machines by reservation to determine which have the largest CPU and memory reservations. If there are outlier virtual machines with much higher reservations than the others, consider using a different vSphere HA admission control policy (such as the Percentage of Cluster Resources Reserved admission control policy) or use the vSphere HA advanced options to place an absolute cap on the slot size. Both of these options, however, increase the risk of resource fragmentation.

Fewer Available Slots Shown Than Expected

The Advanced Runtime Info box might display a smaller number of available slots in the cluster than you expect.

Problem

When you select the Host Failures Cluster Tolerates admission control policy, view the **Advanced Runtime Info** pane that appears in the vSphere HA section of the cluster's **Monitor** tab in the vSphere Web Client. This pane displays information about the cluster, including the number of slots available to power on additional virtual machines in the cluster. This number might be smaller than expected under certain conditions.

Cause

Slot size is calculated using the largest reservations plus the memory overhead of any powered on virtual machines in the cluster. However, vSphere HA admission control considers only the resources on a host that are available for virtual machines. This amount is less than the total amount of physical resources on the host, because there is some overhead.

Solution

Reduce the virtual machine reservations if possible, use vSphere HA advanced options to reduce the slot size, or use a different admission control policy.

Troubleshooting Heartbeat Datastores

When the master host in a vSphere HA cluster can no longer communicate with a subordinate host over the management network, the master host uses datastore heartbeating to determine if the subordinate host might have failed or is in a network partition. If the subordinate host has stopped datastore heartbeating, that host is considered to have failed and its virtual machines are restarted elsewhere.

vCenter Server automatically selects a preferred set of datastores for heartbeating. This selection is made with the goal of maximizing the number of hosts that have access to a given datastore and minimizing the likelihood that the selected datastores are backed by the same storage array or NFS server. In most cases, this selection should not be changed. To see which datastores vSphere HA has selected for use, in the vSphere Web Client you can go to the cluster's **Monitor** tab and select vSphere HA and Heartbeat. Only datastores mounted by at least two hosts are available here.

NOTE There is no heartbeat datastore available if the only shared storage accessible to all hosts in the cluster is vSAN.

User-Preferred Datastore is Not Chosen

vCenter Server might not choose a datastore that you specify as a preference for vSphere HA storage heartbeating.

Problem

You can specify the datastores preferred for storage heartbeating, and based on this preference, vCenter Server determines the final set of datastores to use. However, vCenter Server might not choose the datastores that you specify.

Cause

This problem can occur in the following cases:

- The specified number of datastores is more than is required. vCenter Server chooses the optimal number of required datastores out of the stated user preference and ignores the rest.

- A specified datastore is not optimal for host accessibility and storage backing redundancy. More specifically, the datastore might not be chosen if it is accessible to only a small set of hosts in the cluster. A datastore also might not be chosen if it is on the same LUN or the same NFS server as datastores that vCenter Server has already chosen.
- A specified datastore is inaccessible because of storage failures, for example, storage array all paths down (APD) or permanent device loss (PDL).
- If the cluster contains a network partition, or if a host is unreachable or isolated, the host continues to use the existing heartbeat datastores even if the user preferences change.

Solution

Verify that all the hosts in the cluster are reachable and have the vSphere HA agent running. Also, ensure that the specified datastores are accessible to most, if not all, hosts in the cluster and that the datastores are on different LUNs or NFS servers.

Unmounting or Removing Datastore Fails

When you try to unmount or remove a datastore, the operation fails.

Problem

The operation to unmount or remove a datastore fails if the datastore has any opened files. For these user operations, the vSphere HA agent closes all the files that it has opened, for example, heartbeat files. If the agent is not reachable by vCenter Server or the agent cannot flush out pending I/Os to close the files, a The HA agent on host '{hostName}' failed to quiesce file activity on datastore '{dsName}' fault is triggered.

Cause

If the datastore to be unmounted or removed is used for heartbeating, vCenter Server excludes it from heartbeating and chooses a new one. However, the agent does not receive the updated heartbeat datastores if it is not reachable, that is, if the host is isolated or in a network partition. In such cases, heartbeat files are not closed and the user operation fails. The operation can also fail if the datastore is not accessible because of storage failures such as all paths down.

NOTE When you remove a VMFS datastore, the datastore is removed from all the hosts in inventory. So if there are any hosts in a vSphere HA cluster that are unreachable or that cannot access the datastore, the operation fails.

Solution

Ensure that the datastore is accessible and the affected hosts are reachable.

Troubleshooting vSphere HA Failure Response

vSphere HA provides high availability for virtual machines by pooling them and the hosts that they reside on into a cluster. Hosts in the cluster are monitored and in the event of a failure, the virtual machines on a failed host are restarted on alternate hosts.

There are various reasons why affected virtual machines might not be restarted and if this happens you should troubleshoot to determine the cause.

Incorrect Virtual Machine Protection State

A virtual machine in a vSphere HA cluster is reported as vSphere HA unprotected although it has been powered on for several minutes.

Problem

When a virtual machine is powered on for several minutes, yet its vSphere HA protection state remains as unprotected, if a failure occurs, vSphere HA might not attempt to restart the virtual machine.

Cause

vCenter Server reports a virtual machine as protected after the vSphere HA master host that is responsible for the virtual machine has saved to disk the information that the virtual machine must be restarted after a failure. This process can fail for a number of reasons.

- vSphere HA master host has not been elected or vCenter Server is unable to communicate with it.

In this situation, vCenter Server reports the vSphere HA host state for the cluster hosts as Agent Unreachable or Agent Uninitialized and reports a cluster configuration problem that a master host has not been found.

- Multiple master hosts exist and the one with which vCenter Server is communicating is not responsible for the virtual machine.

Problems occur when vCenter Server is in contact with a master host, but due to a management network partition, there are multiple master hosts, and the agent with which vCenter Server is communicating is not responsible for the virtual machine. This situation is likely if vCenter Server is reporting the vSphere HA state of some hosts as network partitioned.

- Agent is unable to access the datastore on which the configuration file of the virtual machine is stored.

vCenter Server might be in contact with the vSphere HA master host that owns the virtual machine, but the agent is unable to access the datastore on which the configuration file of the virtual machine is stored. This situation can occur if an all paths down condition affects all hosts in the cluster.

Solution

- 1 Determine whether vCenter Server is in contact with a vSphere HA master host, and if not, address this problem.
- 2 If vCenter Server is in contact with a master host, determine whether there is a network partition, and if so, address that problem.
- 3 If the problem persists, determine if other virtual machines that use the same datastore for their configuration files are also unprotected.
- 4 If these virtual machines are unprotected, verify that the vSphere HA master host can access the datastore.
- 5 If none of the previous steps resolves the problem, restore protection by reconfiguring vSphere HA on the host on which the virtual machine is running.

Virtual Machine Restart Fails

After a host or virtual machine failure, a virtual machine might not be restarted.

Problem

When a host fails or a virtual machine fails while its host continues running, the virtual machine might not restart or restarts only after a long delay.

Cause

vSphere HA might not restart a virtual machine after a failure or might delay its restart for several reasons.

- Virtual machine is not protected by vSphere HA at the time the failure occurred
- Insufficient spare capacity on hosts with which the virtual machine is compatible
- vSphere HA attempted to restart the virtual machine but encountered a fatal error each time it tried.
- Your cluster's shared storage is vSAN and one of the virtual machine's files has become inaccessible due to the occurrence of more than the specified number of host failures.
- Restart actually succeeded.

Solution

To avoid virtual machine restart failures, check that virtual machines become protected by vSphere HA after they are powered on. Also, ensure that your admission control settings match your restart expectations if a failure occurs. Maximizing the compatibility between virtual machines and hosts in the cluster can also reduce the likelihood of restart failures.

NOTE For information on the factors vSphere HA considers for virtual machine restarts, see "Determining Responses to Host Issues" in *vSphere Availability*.

Troubleshooting vSphere Fault Tolerance in Network Partitions

When a vSphere HA cluster experiences a failure of the network that vSphere uses for inter-agent communication (the management network), a subset of the cluster's hosts might be unable to communicate with other cluster hosts. In this case, the set of hosts that can communicate with each other are considered to be in a network partition.

A cluster partition impedes cluster management functions such as vMotion and can impact vSphere HA's ability to monitor and restart virtual machines after a failure. This condition must be corrected as soon as possible.

Network partitions also degrade the functionality of vSphere Fault Tolerance. For example, in a partitioned cluster, a Primary VM (or its Secondary VM) could end up in a partition managed by a master host that is not responsible for the virtual machine. When a Secondary VM must be restarted, vSphere HA does so only if the Primary VM is in a partition managed by the master host responsible for it. Ultimately, you must correct the network partition, but until that is possible, you must troubleshoot and correct any problems that arise with your fault-tolerant virtual machines to ensure that they are properly protected.

Primary VM Remains in the Need Secondary State

A fault tolerant Primary VM can remain in the need secondary state even though sufficient resources are available to start the Secondary VM.

Problem

vSphere HA might not restart the Secondary VM of a vSphere Fault Tolerance (FT) virtual machine pair even though there are sufficient resources available.

Cause

To restart a Secondary VM, vSphere HA requires that the Primary VM be running on a host that is in the same partition as the one containing the vSphere HA master host responsible for the FT pair. In addition, the vSphere HA agent on the Primary VM's host must be operating correctly. If these conditions are met, FT also requires that there be at least one other host in the same partition that is compatible with the FT pair and that has a functioning vSphere HA agent.

Solution

To fix this condition, check the vSphere HA host states reported by vCenter Server. If hosts are identified as partitioned, isolated, or unreachable, resolve the problem before proceeding. In some situations, you can resolve a restart problem by reconfiguring vSphere HA on the host that vCenter Server is reporting as the master host. However, in most situations, this step is insufficient, and you must resolve all host state problems.

After you have addressed any host state problems, check if there are any hosts in the cluster other than the Primary VM's that are compatible with the FT virtual machine pair. You can determine compatibility by trying to migrate the Primary VM to other hosts. Address any incompatibilities that are discovered.

Role Switch Behavior Problems

vCenter Server can report that the Primary VM of a vSphere Fault Tolerance virtual machine pair is powered off, but the Secondary VM is powered on.

Problem

After a failover occurs, vCenter Server might incorrectly report that the Primary VM is powered off and registered to its original host, and that the Secondary VM is powered on and registered to its original host.

Cause

This error occurs when vCenter Server is unable to communicate with the hosts on which the Primary VM and Secondary VM are actually running. vCenter Server reports these hosts as not responding and the problem persists until vCenter Server is able to communicate with the hosts.

Solution

To fix this problem, resolve the networking problem that is preventing vCenter Server from communicating with the hosts in the cluster.

Troubleshooting VM Component Protection

If you enable VM Component Protection (VMCP) for your vSphere HA cluster, it provides protection against datastore accessibility failures that can affect a virtual machine running on one of the cluster's hosts.

If the response that you have configured VMCP to make for such a failure is not run, you should troubleshoot to determine the cause.

Datastore Inaccessibility Is Not Resolved for a VM

When a datastore becomes inaccessible, VMCP might not terminate and restart the affected virtual machines.

Problem

When an All Paths Down (APD) or Permanent Device Loss (PDL) failure occurs and a datastore becomes inaccessible, VMCP might not resolve the issue for the affected virtual machines.

Cause

In an APD or PDL failure situation, VMCP might not terminate a virtual machine for the following reasons:

- VM is not protected by vSphere HA at the time of failure.
- VMCP is disabled for this virtual machine.

Furthermore, if the failure is an APD, VMCP might not terminate a VM for several reasons:

- APD failure is corrected before the VM was terminated.

- Insufficient capacity on hosts with which the virtual machine is compatible
- During a network partition or isolation, the host affected by the APD failure is not able to query the master host for available capacity. In such a case, vSphere HA defers to the user policy and terminates the VM if the VM Component Protection setting is aggressive.
- vSphere HA terminates APD-affected VMs only after the following timeouts expire:
 - APD timeout (default 140 seconds).
 - APD failover delay (default 180 seconds). For faster recovery, this can be set to 0.

NOTE Based on these default values, vSphere HA terminates the affected virtual machine after 320 seconds (APD timeout + APD failover delay)

Solution

To address this issue, check and adjust any of the following:

- Insufficient capacity to restart the virtual machine
- User-configured timeouts and delays
- User settings affecting VM termination
- VM Component Protection policy
- Host monitoring or VM restart priority must be enabled

Troubleshooting Resource Management

6

The resource management troubleshooting topics provide solutions to potential problems that you might encounter when using your hosts and datastores in vSphere DRS or vSphere Storage DRS cluster.

This chapter includes the following topics:

- [“Troubleshooting Storage DRS,”](#) on page 47
- [“Troubleshooting Storage I/O Control,”](#) on page 52

Troubleshooting Storage DRS

The Storage DRS troubleshooting topics provide solutions to potential problems that you might encounter when using Storage DRS-enabled datastores in a datastore cluster.

Storage DRS is Disabled on a Virtual Disk

Even when Storage DRS is enabled for a datastore cluster, it might be disabled on some virtual disks in the datastore cluster.

Problem

You have enabled Storage DRS for a datastore cluster, but Storage DRS is disabled on one or more virtual machine disks in the datastore cluster.

Cause

The following scenarios can cause Storage DRS to be disabled on a virtual disk.

- A virtual machine's swap file is host-local (the swap file is stored in a specified datastore that is on the host). The swap file cannot be relocated and Storage DRS is disabled for the swap file disk.
- A certain location is specified for a virtual machine's `.vmx` swap file. The swap file cannot be relocated and Storage DRS is disabled on the `.vmx` swap file disk.
- The relocate or Storage vMotion operation is currently disabled for the virtual machine in vCenter Server (for example, because other vCenter Server operations are in progress on the virtual machine). Storage DRS is disabled until the relocate or Storage vMotion operation is re-enabled in vCenter Server.
- The home disk of a virtual machine is protected by vSphere HA and relocating it will cause loss of vSphere HA protection.
- The disk is a CD-ROM/ISO file.
- If the disk is an independent disk, Storage DRS is disabled, except in the case of relocation or clone placement.

- If the virtual machine has system files on a separate datastore from the home datastore (legacy), Storage DRS is disabled on the home disk. If you use Storage vMotion to manually migrate the home disk, the system files on different datastores will be all be located on the target datastore and Storage DRS will be enabled on the home disk.
- If the virtual machine has a disk whose base/redo files are spread across separate datastores (legacy), Storage DRS for the disk is disabled. If you use Storage vMotion to manually migrate the disk, the files on different datastores will be all be located on the target datastore and Storage DRS will be enabled on the disk.
- The virtual machine has hidden disks (such as disks in previous snapshots, not in the current snapshot). This situation causes Storage DRS to be disabled on the virtual machine.
- The virtual machine is a template.
- The virtual machine is vSphere Fault Tolerance-enabled.
- The virtual machine is sharing files between its disks.
- The virtual machine is being Storage DRS-placed with manually specified datastores.

Solution

Address the problem that is causing Storage DRS to be disabled on the disk.

Datastore Cannot Enter Maintenance Mode

You place a datastore in maintenance mode when you must take it out of usage to service it. A datastore enters or leaves maintenance mode only as a result of a user request.

Problem

A datastore in a datastore cluster cannot enter maintenance mode. The Entering Maintenance Mode status remains at 1%.

Cause

One or more disks on the datastore cannot be migrated with Storage vMotion. This condition can occur in the following instances.

- Storage DRS is disabled on the disk.
- Storage DRS rules prevent Storage DRS from making migration recommendations for the disk.

Solution

- If Storage DRS is disabled, enable it or determine why it is disabled. See [“Storage DRS is Disabled on a Virtual Disk,”](#) on page 47 for reasons why Storage DRS might be disabled.
- If Storage DRS rules are preventing Storage DRS from making migration recommendations, you can remove or disable particular rules.
 - a Browse to the datastore cluster in the vSphere Web Client object navigator.
 - b Click the **Manage** tab and click **Settings**.
 - c Under Configuration, select **Rules** and click the rule.
 - d Click **Remove**.

- Alternatively, if Storage DRS rules are preventing Storage DRS from making migration recommendations, you can set the Storage DRS advanced option `IgnoreAffinityRulesForMaintenance` to 1.
 - a Browse to the datastore cluster in the vSphere Web Client object navigator.
 - b Click the **Manage** tab and click **Settings**.
 - c Select **SDRS** and click **Edit**.
 - d In **Advanced Options > Configuration Parameters**, click **Add**.
 - e In the Option column, enter **IgnoreAffinityRulesForMaintenance**.
 - f In the Value column, enter **1** to enable the option.
 - g Click **OK**.

Storage DRS Cannot Operate on a Datastore

Storage DRS generates an alarm to indicate that it cannot operate on the datastore.

Problem

Storage DRS generates an event and an alarm and Storage DRS cannot operate.

Cause

The following scenarios can cause vCenter Server to disable Storage DRS for a datastore.

- The datastore is shared across multiple data centers.

Storage DRS is not supported on datastores that are shared across multiple data centers. This configuration can occur when a host in one data center mounts a datastore in another data center, or when a host using the datastore is moved to a different data center. When a datastore is shared across multiple data centers, Storage DRS I/O load balancing is disabled for the entire datastore cluster. However, Storage DRS space balancing remains active for all datastores in the datastore cluster that are not shared across data centers.
- The datastore is connected to an unsupported host.

Storage DRS is not supported on ESX/ESXi 4.1 and earlier hosts.
- The datastore is connected to a host that is not running Storage I/O Control.

Solution

- The datastore must be visible in only one data center. Move the hosts to the same data center or unmount the datastore from hosts that reside in other data centers.
- Ensure that all hosts associated with the datastore cluster are ESXi 5.0 or later.
- Ensure that all hosts associated with the datastore cluster have Storage I/O Control enabled.

Moving Multiple Virtual Machines into a Datastore Cluster Fails

Migrating more than one datastore into a datastore cluster fails with an error message after the first virtual machine has successfully moved into the datastore cluster.

Problem

When you attempt to migrate multiple virtual machines into a datastore cluster, some virtual machines migrate successfully, but migration of subsequent virtual machines fails. vCenter Server displays the error message, `Insufficient Disk Space on Datastore`.

Cause

Until each placement recommendation is applied, the space resources appear to be available to Storage DRS. Therefore, Storage DRS might reallocate space resources to subsequent requests for space.

Solution

Retry the failed migration operations one at a time and ensure that each recommendation is applied before requesting the next migration

Storage DRS Generates Fault During Virtual Machine Creation

When you create or clone a virtual machine on a datastore cluster, Storage DRS might generate a fault.

Problem

When you attempt to create or clone a virtual machine on a datastore cluster, you might receive the error message, `Operation Not Allowed in the Current State`.

Cause

Storage DRS checks for rule violations when you create a virtual machine on a Storage DRS-enabled datastore. If Storage DRS cannot create the new virtual machine's disks in compliance with the rules, it generates a fault. The fault is generated because Storage DRS cannot reference the virtual machine, which is in the process of being created and does not yet exist.

Solution

Revise or remove the rules and retry the create or clone virtual machine operation.

Storage DRS is Enabled on a Virtual Machine Deployed from an OVF Template

Storage DRS is enabled on a virtual machine that was deployed from an OVF template that has Storage DRS disabled. This can occur when you deploy an OVF template on a datastore cluster.

Problem

When you deploy an OVF template with Storage DRS disabled on a datastore cluster, the resulting virtual machine has Storage DRS enabled.

Cause

The vSphere Web Client applies the default automation level of the datastore cluster to virtual machines deployed from an OVF template.

Solution

- 1 To manually change the automation level of the virtual machine, browse to the datastore cluster in the vSphere Web Client object navigator.
- 2 Click the **Manage** tab and select **Settings**.
- 3 Select **VM Overrides** and click **Add**.
- 4 Select the virtual machine and click **OK**.
- 5 From the **Keep VMDKs Together** dropdown menu, select **No** and click **OK**.

Storage DRS Rule Violation Fault Is Displayed Multiple Times

When you attempt to put a datastore into maintenance mode, the same affinity or anti-affinity rule violation fault might appear to be listed more than once in the Faults dialog box.

Problem

The Faults dialog box appears to display multiple instances of identical faults, but in fact, each fault refers to a different datastore. The Faults dialog box does not list the names of the datastores, which causes the faults to appear to be redundant.

Solution

The Faults dialog box always displays a separate rule violation fault for each datastore that is considered for placement. If you want the datastore to enter maintenance mode, remove the rule that prevents the virtual machine from being migrated.

Storage DRS Rules Not Deleted from Datastore Cluster

Affinity or anti-affinity rules that apply to a virtual machine are not deleted when you remove the virtual machine from a datastore cluster.

Problem

When you remove a virtual machine from a datastore cluster, and that virtual machine is subject to an affinity or anti-affinity rule in a datastore cluster, the rule remains. This allows you to store virtual machine configurations in different datastore clusters. If the virtual machine is moved back into the datastore cluster, the rule is applied. You cannot delete the rule after you remove the virtual machine from the datastore cluster.

Cause

vCenter Server retains rules for a virtual machine that is removed from a datastore cluster if the virtual machine remains in the vCenter Server inventory.

Solution

To remove a rule from a datastore cluster configuration, you must delete the rule before you remove the virtual machine to which the rule applies from the datastore cluster.

- 1 In the vSphere Web Client, browse to the datastore cluster.
- 2 Click the **Manage** tab and select **Settings**.
- 3 Under Configuration, click **Rules**.
- 4 Select the rule to delete and click **Remove**.
- 5 Click **OK**.

Alternative Storage DRS Placement Recommendations Are Not Generated

When you create, clone, or relocate a virtual machine, Storage DRS generates only one placement recommendation.

Problem

Storage DRS generates a single placement recommendation when you create, clone, or relocate a virtual machine. No alternative recommendations are provided when multiple alternative recommendations are expected.

Cause

If the destination host explicitly specifies the virtual machine's swap file location as a datastore in the target datastore cluster, the disks to be placed in that cluster do not form a single affinity group. Storage DRS generates alternative placement recommendations only for a single item or a single affinity group.

Solution

Accept the single recommendation. To obtain multiple recommendations, choose a destination host that does not specify that the virtual machine swap file location is on a datastore that is in the target datastore cluster.

Applying Storage DRS Recommendations Fails

Storage DRS generates space or I/O load balancing recommendations, but attempts to apply the recommendations fail.

Problem

When you apply Storage DRS recommendations for space or I/O load balancing, the operation fails.

Cause

The following scenarios can prevent you from applying Storage DRS recommendations.

- A Thin Provisioning Threshold Crossed alarm might have been triggered for the target datastore, which indicates that the datastore is running out of space and no virtual machines will be migrated to it.
- The target datastore might be in maintenance mode or is entering maintenance mode.

Solution

- Address the issue that triggered the Thin Provisioning Threshold Crossed alarm.
- Verify that the target datastore is not in maintenance mode or entering maintenance mode.

Troubleshooting Storage I/O Control

The Storage I/O Control troubleshooting topics provide solutions to potential problems that you might encounter when using Storage I/O Control with datastores.

Unsupported Host Connected to Datastore

In the vSphere Web Client, an alarm is triggered when vCenter Server detects that a workload from a host might be affecting performance.

Problem

The alarm **Pre-4.1 host connected to SIOC-enabled datastore** is triggered.

Cause

The datastore is Storage I/O Control-enabled, but it cannot be fully controlled by Storage I/O Control because of the external workload.

This condition can occur if the Storage I/O Control-enabled datastore is connected to a host that does not support Storage I/O Control.

Solution

Ensure that all hosts that are connected to the datastore support Storage I/O Control.

Unmanaged Workload Detected on Datastore

In the vSphere Web Client, an alarm is triggered when vCenter Server detects that a workload from a host might be affecting performance.

Problem

The alarm **Unmanaged workload is detected on the datastore** is triggered.

Cause

The array is shared with non-vSphere workloads, or the array is performing system tasks such as replication.

Solution

There is no solution. vCenter Server does not reduce the total amount of I/O sent to the array, but continues to enforce shares.

Unable to View Performance Charts for Datastore

Performance charts for a datastore do not appear on the Performance tab.

Problem

You are unable to view performance charts for a datastore on the **Performance** tab in the vSphere Web Client.

Cause

Storage I/O Control is disabled for the datastore.

Solution

- 1 Browse to the datastore in the vSphere Web Client object navigator.
- 2 Right-click the datastore and select **Configure Storage I/O Control**.
- 3 Select the **Enable Storage I/O Control** check box.
- 4 Click **OK**.

Cannot Enable Storage I/O Control on Datastore

Storage I/O Control is disabled on a datastore and cannot be enabled.

Problem

You cannot enable Storage I/O Control on a datastore.

Cause

The following reasons might prevent you from enabling Storage I/O Control on a datastore.

- At least one host that is connected to the datastore is not running ESX/ESXi 4.1 or later.
- You do not have the appropriate license to enable Storage I/O Control.

Solution

- Verify that the hosts connected to the datastore are ESX/ESXi 4.1 or later.
- Verify that you have the appropriate license to enable Storage I/O Control.

Troubleshooting Storage

The storage troubleshooting topics provide solutions to potential problems that you might encounter when using vSphere in different storage environments that include SAN, vSAN, or Virtual Volumes.

This chapter includes the following topics:

- [“Resolving SAN Storage Display Problems,”](#) on page 56
- [“Resolving SAN Performance Problems,”](#) on page 57
- [“Virtual Machines with RDMs Need to Ignore SCSI INQUIRY Cache,”](#) on page 62
- [“Software iSCSI Adapter Is Enabled When Not Needed,”](#) on page 62
- [“Failure to Mount NFS Datastores,”](#) on page 63
- [“Troubleshooting Storage Adapters,”](#) on page 63
- [“Checking Metadata Consistency with VOMA,”](#) on page 64
- [“No Failover for Storage Path When TUR Command Is Unsuccessful,”](#) on page 65
- [“Troubleshooting Flash Devices,”](#) on page 66
- [“Troubleshooting Virtual Volumes,”](#) on page 69
- [“Troubleshooting VAIO Filters,”](#) on page 71

Resolving SAN Storage Display Problems

When you use the vSphere Web Client to display Fibre Channel SAN or iSCSI storage devices, you might not be able to see all devices available to your host. A number of troubleshooting tasks exist that you can perform to resolve storage display problems.

Resolving Fibre Channel Storage Display Problems

If Fibre Channel storage devices do not display correctly in the vSphere Web Client, perform troubleshooting tasks.

Table 7-1. Troubleshooting Fibre Channel LUN Display

Troubleshooting Task	Description
Check cable connectivity.	If you do not see a port, the problem could be cable connectivity. Check the cables first. Ensure that cables are connected to the ports and a link light indicates that the connection is good. If each end of the cable does not show a good link light, replace the cable.
Check zoning.	Zoning limits access to specific storage devices, increases security, and decreases traffic over the network. Some storage vendors allow only single-initiator zones. In that case, an HBA can be in multiple zones to only one target. Other vendors allow multiple-initiator zones. See your storage vendor's documentation for zoning requirements. Use the SAN switch software to configure and manage zoning.
Check access control configuration.	<ul style="list-style-type: none"> ■ The MASK_PATH plug-in allows you to prevent your host from accessing a specific storage array or specific LUNs on a storage array. If your host is detecting devices and paths that you do not want the host to access, path masking could have been set up incorrectly. ■ For booting from a SAN, ensure that each host sees only required LUNs. Do not allow any host to see any boot LUN other than its own. Use storage system software to make sure that the host can see only the LUNs that it is supposed to see. ■ Ensure that the Disk.MaxLUN parameter allows you to view the LUN you expect to see. For information on the parameter, see the <i>vSphere Storage</i> documentation.
Check storage processor setup.	If a disk array has more than one storage processor (SP), make sure that the SAN switch has a connection to the SP that owns the LUNs you want to access. On some disk arrays, only one SP is active and the other SP is passive until there is a failure. If you are connected to the wrong SP (the one with the passive path), you might see the LUNs but get errors when trying to access them.
Rescan your HBA.	<p>Perform a rescan each time you complete the following tasks:</p> <ul style="list-style-type: none"> ■ Create new LUNs on a SAN. ■ Change the path masking configuration on the host. ■ Reconnect a cable. ■ Make a change to a host in a cluster. <p>For information, see the <i>vSphere Storage</i> documentation.</p>

Resolving iSCSI Storage Display Problems

Perform troubleshooting tasks if iSCSI storage devices do not display correctly in the vSphere Web Client.

Table 7-2. Troubleshooting iSCSI LUN Display

Troubleshooting Task	Description
Check cable connectivity.	If you do not see a port, the problem could be cable connectivity or routing. Check the cables first. Ensure that cables are connected to the ports and a link light indicates that the connection is good. If each end of the cable does not show a good link light, replace the cable.
Check routing settings.	Controls connectivity between different subnets on your Ethernet configuration. If your ESXi system and iSCSI storage are not on the same subnet, ensure that appropriate routing exists between the subnets. Also, ensure that the subnet mask and gateway address are set correctly on the iSCSI storage and the iSCSI initiator in the ESXi host.
Check access control configuration.	If the expected LUNs do not appear after rescan, access control might not be configured correctly on the storage system side: <ul style="list-style-type: none"> ■ If CHAP is configured, ensure that it is enabled on the ESXi host and matches the storage system setup. ■ If IP-based filtering is used, ensure that the iSCSI HBA or the VMkernel port group IP address is allowed. ■ If you are using initiator name-based filtering, ensure that the name is a qualified iSCSI name and matches the storage system setup. ■ For booting from a SAN, ensure that each host sees only required LUNs. Do not allow any host to see any boot LUN other than its own. Use storage system software to make sure that the host can see only the LUNs that it is supposed to see. ■ Ensure that the Disk.MaxLUN setting allows you to view the LUN you expect to see. For information, see the <i>vSphere Storage</i> documentation.
Check storage processor setup.	If a storage system has more than one storage processor, make sure that the SAN switch has a connection to the SP that owns the LUNs you want to access. On some storage systems, only one SP is active and the other SP is passive until a failure occurs. If you are connected to the wrong SP (the one with the passive path) you might not see the expected LUNs, or you might see the LUNs but get errors when trying to access them.
For software and dependent hardware iSCSI, check network configuration.	The software iSCSI and dependent hardware adapters in ESXi require that VMkernel network port have access to the iSCSI storage. The adapters use the VMkernel for data transfer between the ESXi system and the iSCSI storage.
Rescan your iSCSI initiator.	Perform a rescan each time you complete the following tasks: <ul style="list-style-type: none"> ■ Create new LUNs on a SAN. ■ Change the LUN masking. ■ Reconnect a cable. ■ Make a change to a host in a cluster. ■ Change CHAP settings or add new discovery addresses. For information, see the <i>vSphere Storage</i> documentation.

Resolving SAN Performance Problems

A number of factors can negatively affect storage performance in the ESXi SAN environment. Among these factors are excessive SCSI reservations, path thrashing, and inadequate LUN queue depth.

To monitor storage performance in real time, use the `resxtop` and `esxtop` command-line utilities. For more information, see the *vSphere Monitoring and Performance* documentation.

Excessive SCSI Reservations Cause Slow Host Performance

When storage devices do not support the hardware acceleration, ESXi hosts use the SCSI reservations mechanism when performing operations that require a file lock or a metadata lock in VMFS. SCSI reservations lock the entire LUN. Excessive SCSI reservations by a host can cause performance degradation on other servers accessing the same VMFS.

Problem

Excessive SCSI reservations cause performance degradation and SCSI reservation conflicts.

Cause

Several operations require VMFS to use SCSI reservations.

- Creating, resignaturing, or expanding a VMFS datastore
- Powering on a virtual machine
- Creating or deleting a file
- Creating a template
- Deploying a virtual machine from a template
- Creating a new virtual machine
- Migrating a virtual machine with VMotion
- Growing a file, such as a thin provisioned virtual disk

Note For storage devices that support the hardware acceleration, the hosts use the atomic test and set (ATS) algorithm to lock the LUN. For more information on hardware acceleration, see the *vSphere Storage* documentation.

Solution

To eliminate potential sources of SCSI reservation conflicts, follow these guidelines:

- Serialize the operations of the shared LUNs, if possible, limit the number of operations on different hosts that require SCSI reservation at the same time.
- Increase the number of LUNs and limit the number of hosts accessing the same LUN.
- Reduce the number snapshots. Snapshots cause numerous SCSI reservations.
- Reduce the number of virtual machines per LUN. Follow recommendations in *Configuration Maximums*.
- Make sure that you have the latest HBA firmware across all hosts.
- Make sure that the host has the latest BIOS.
- Ensure a correct Host Mode setting on the SAN array.

For information about handling SCSI reservation conflicts on specific storage arrays, see the VMware knowledge base article at <http://kb.vmware.com/kb/1005009>.

Path Thrashing Causes Slow LUN Access

If your ESXi host is unable to access a LUN, or access is very slow, you might have a problem with path thrashing, also called LUN thrashing.

Problem

Your host is unable to access a LUN, or access is very slow. The host's log files might indicate frequent path state changes. For example:

```
Frequent path state changes are occurring for path vmhba2:C0:T0:L3. This may indicate a storage problem. Affected device: naa.6006000000000000edd1. Affected datastores: ds1
```

Cause

The problem might be caused by path thrashing. Path thrashing might occur when two hosts access the same LUN through different storage processors (SPs) and, as a result, the LUN is never available.

Path thrashing typically occurs on active-passive arrays. Path thrashing can also occur on a directly connected array with HBA failover on one or more nodes. Active-active arrays or arrays that provide transparent failover do not cause path thrashing.

Solution

- 1 Ensure that all hosts that share the same set of LUNs on the active-passive arrays use the same storage processor.
- 2 Correct any cabling or masking inconsistencies between different hosts and SAN targets so that all HBAs see the same targets.
- 3 Ensure that the claim rules defined on all hosts that share the LUNs are exactly the same.
- 4 Configure the path to use the Most Recently Used PSP, which is the default.

Increased Latency for I/O Requests Slows Virtual Machine Performance

If the ESXi host generates more commands to a LUN than the LUN queue depth permits, the excess commands are queued in VMkernel. This increases the latency, or the time taken to complete I/O requests.

Problem

The host takes longer to complete I/O requests and virtual machines display unsatisfactory performance.

Cause

The problem might be caused by an inadequate LUN queue depth. SCSI device drivers have a configurable parameter called the LUN queue depth that determines how many commands to a given LUN can be active at one time. If the host generates more commands to a LUN, the excess commands are queued in the VMkernel.

Solution

- 1 If the sum of active commands from all virtual machines consistently exceeds the LUN depth, increase the queue depth.

The procedure that you use to increase the queue depth depends on the type of storage adapter the host uses.

- 2 When multiple virtual machines are active on a LUN, change the Disk.SchedNumReqOutstanding (DSNRO) parameter, so that it matches the queue depth value.

Adjust Queue Depth for QLogic, Emulex, and Brocade HBAs

If you are not satisfied with the performance of your hardware bus adapters (HBAs), change the maximum queue depth on your ESXi host.

The maximum value refers to the queue depths reported for various paths to the LUN. When you lower this value, it throttles the host's throughput and alleviates SAN contention concerns if multiple hosts are overutilizing the storage and are filling its command queue.

To adjust the maximum queue depth parameter, use the vCLI commands.

In the procedure, `--server=server_name` specifies the target server. The specified target server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see *Getting Started with vSphere Command-Line Interfaces*.

Prerequisites

Install vCLI or deploy the vSphere Management Assistant (vMA) virtual machine. See *Getting Started with vSphere Command-Line Interfaces*. For troubleshooting, run `esxcli` commands in the ESXi Shell.

Procedure

- 1 Verify which HBA module is currently loaded by entering the following command:

```
esxcli --server=server_name system module list | grep module
```

Use one the following options for *module*.

Option	Description
<code>qla</code>	QLogic
<code>qln</code>	QLogic native drivers
<code>lpfc</code>	Emulex
<code>bfa</code>	Brocade

- 2 Adjust the queue depth for the appropriate module.

```
esxcli --server=server_name system module parameters set -p parameter=value -m module
```

Use the following strings for the *parameter* and *module* options.

String	Description
<code>-p ql2xmaxqdepth=value</code> <code>-m qla2xxx</code>	QLogic
<code>-p ql2xmaxqdepth=value</code> <code>-m qlnativefc</code>	QLogic native drivers
<code>-p lpfc0_lun_queue_depth=value</code> <code>-m lpfc820</code>	Emulex
<code>-p lpfc0_lun_queue_depth=value</code> <code>-m lpfc</code>	Emulex native drivers
<code>-p bfa_lun_queue_depth=value</code> <code>-m bfa</code>	Brocade

- 3 Reboot your host.
- 4 Verify your changes by running the following command:
`esxcli --server=server_name system module parameters list -m=module.`

module is an appropriate driver, such as `qlnativefc` or `bfa`.

Adjust Maximum Queue Depth for Software iSCSI

If you notice unsatisfactory performance for your software iSCSI LUNs, change their maximum queue depth by running the `esxcli` commands.

Prerequisites

- Install vCLI or deploy the vSphere Management Assistant (vMA) virtual machine. See *Getting Started with vSphere Command-Line Interfaces*. For troubleshooting, you can run `esxcli` commands in the ESXi Shell.
- In the procedure, the `--server=server_name` connection option specifies the target server. Be prepared to enter a user name and password when the target server prompts you. For a list of other possible connection options, see *Getting Started with vSphere Command-Line Interfaces*.

Procedure

- 1 Run the following command:

```
esxcli --server=server_name system module parameters set -m iscsi_vmk -p iscsi_vmk_LunQDepth=value
```

The `iscsi_vmk_LunQDepth` parameter sets the maximum number of outstanding commands, or queue depth, for each LUN accessed through the software iSCSI adapter. The default value is 128.

- 2 Reboot your system.

- 3 Verify your changes by running the

```
esxcli --server=server_name system module parameters list -m iscsi_vmk command.
```

The following output shows the queue depth for software iSCSI.

```
iscsi_vmk_LunQDepth int 64 Maximum Outstanding Commands Per LUN
```



CAUTION Setting the queue depth to a value higher than the default can decrease the total number of LUNs supported.

Change the Outstanding IO Requests Setting

If you adjusted the LUN queue depth, change the `Disk.SchedNumReqOutstanding` (DSNRO) parameter, so that its value matches the queue depth. The parameter controls the maximum number of outstanding IO requests that all virtual machines can issue to the LUN.

Change this parameter only when you have multiple virtual machines active on a LUN. The parameter does not apply when only one virtual machine is active. In that case, the bandwidth is controlled by the queue depth of the storage adapter.

The parameter is set per device.

Procedure

- 1 Enter the following command to display the current DSNRO setting for the specified device:

```
esxcli storage core device list -d device_ID
```

You get the output similar to the following:

```
No of outstanding IOs with competing worlds: 32
```

- 2 Change the DSNRO value by entering the following command:

```
esxcli storage core device set -0 | --sched-num-req-outstanding value -d device_ID
```

- 3 Verify your changes by entering the following command:

```
esxcli storage core device list -d device_ID
```

Virtual Machines with RDMs Need to Ignore SCSI INQUIRY Cache

Storage vendors might require that virtual machines with RDMs ignore SCSI INQUIRY data cached by ESXi.

Problem

Certain guest operating systems or applications run in virtual machines with RDMs display unpredictable behavior.

Cause

This behavior might be caused by cached SCSI INQUIRY data that interferes with specific guest operating systems and applications.

When the ESXi host first connects to a target storage device on a SAN, it issues the SCSI INQUIRY command to obtain basic identification data from the device. By default, ESXi caches the received SCSI INQUIRY data (Standard, page 80, and page 83) and the data remains unchanged afterwards.

Solution

- ◆ Configure the virtual machine with RDM to ignore the SCSI INQUIRY cache by adding the following parameter to the .vmx file.

```
scsix:y.ignoreDeviceInquiryCache = "true"
```

where *x* is the SCSI controller number and *y* is the SCSI target number of the RDM.

Enable this parameter only when your storage vendor recommends that you do so. This parameter is required for just a limited number of storage arrays and only for specific guest operating systems.

Software iSCSI Adapter Is Enabled When Not Needed

When your host uses a network adapter with iBFT, the software iSCSI adapter is always enabled by default.

Problem

After your ESXi host's first boot, the software iSCSI adapter is enabled and appears in the vSphere Web Client on the list of storage adapters.

Cause

The iBFT-enabled network adapter on your host causes the software iSCSI to be always present. This condition occurs even when you do not use iBFT for the iSCSI boot.

Solution

If you do not use the iBFT-enabled network adapter for the iSCSI boot and do not want the software iSCSI adapter to be enabled, remove the iBFT configuration from the network adapter. Because this process is vendor-specific, consult your vendor documentation for details.

Failure to Mount NFS Datastores

Attempts to mount NFS datastores with names in international languages result in failures.

Problem

The use of non-ASCII characters for directory and filenames on NFS storage might cause unpredictable behavior. For example, you might fail to mount an NFS datastore or not be able to power on a virtual machine.

Cause

ESXi supports the use of non-ASCII characters for directory and filenames on NFS storage, so you can create datastores and virtual machines using names in international languages. However, when the underlying NFS server does not offer internationalization support, unpredictable failures might occur.

Solution

Always make sure that the underlying NFS server offers internationalization support. If the server does not, use only ASCII characters.

Troubleshooting Storage Adapters

If your storage adapters experience performance problems, use the `esxcli storage san` commands to identify the problems.

Problem

Storage adapters experience performance and I/O problem.

Solution

Use the `esxcli storage san` commands to obtain and display events and statistics for the adapters. You can analyze the commands' output to identify the adapter problems and to find appropriate solutions.

Table 7-3. `esxcli storage san` commands

Command	Description	Options
<code>esxcli storage san fc iscsi fcoe sas list</code>	List adapter attributes. NOTE iSCSI applies to software iSCSI only.	-- adapter -A Adapter name (vmhbaX), or none, to list information for all adapters of the particular type.
<code>esxcli storage san fc iscsi fcoe sas stats get</code>	Get adapter statistics. NOTE iSCSI applies to software iSCSI only.	-- adapter -A Adapter name (vmhbaX), or none, to list information for all adapters of the particular type.
<code>esxcli storage san fc fcoe sas reset</code>	Reset a particular adapter.	-- adapter -A Adapter name (vmhbaX).
<code>esxcli storage san fc events get</code>	Retrieve events for Fibre Channel adapters.	-- adapter -A Adapter name (vmhbaX), or none, to list information for all Fibre Channel adapters on the system.

Phase 4: Checking pathname and connectivity.

Phase 5: Checking resource reference counts.

Total Errors Found: 1

Command options that the VOMA tool takes include the following.

Table 7-4. VOMA Command Options

Command Option	Description								
-m --module	The modules to run include the following: <table border="1" data-bbox="598 483 1426 798"> <thead> <tr> <th>Module</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>vmfs</td> <td>If you do not specify the name of the module, this option is used by default. You can check VMFS3, VMFS5, and VMFS6 file systems, as well as file systems that back virtual flash resources. If you specify this module, minimal checks are performed for LVM as well.</td> </tr> <tr> <td>lvm</td> <td>Check logical volumes that back VMFS datastores.</td> </tr> <tr> <td>ptck</td> <td>Check and validate VMFS partitions, such as MBR or GPT. If no partition exists, determine whether partitions should exist.</td> </tr> </tbody> </table>	Module	Description	vmfs	If you do not specify the name of the module, this option is used by default. You can check VMFS3, VMFS5, and VMFS6 file systems, as well as file systems that back virtual flash resources. If you specify this module, minimal checks are performed for LVM as well.	lvm	Check logical volumes that back VMFS datastores.	ptck	Check and validate VMFS partitions, such as MBR or GPT. If no partition exists, determine whether partitions should exist.
Module	Description								
vmfs	If you do not specify the name of the module, this option is used by default. You can check VMFS3, VMFS5, and VMFS6 file systems, as well as file systems that back virtual flash resources. If you specify this module, minimal checks are performed for LVM as well.								
lvm	Check logical volumes that back VMFS datastores.								
ptck	Check and validate VMFS partitions, such as MBR or GPT. If no partition exists, determine whether partitions should exist.								
-f --func	Functions to be performed include the following: <table border="1" data-bbox="598 850 1426 934"> <thead> <tr> <th>Function</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>query</td> <td>List functions supported by module.</td> </tr> <tr> <td>check</td> <td>Check for errors.</td> </tr> </tbody> </table>	Function	Description	query	List functions supported by module.	check	Check for errors.		
Function	Description								
query	List functions supported by module.								
check	Check for errors.								
-d --device	Device or disk to be inspected. Make sure to provide the absolute path to the device partition backing the VMFS datastore. For example, /vmfs/devices/disks/naa.00000000000000000000000000000000:1.								
-s --logfile	Specify the log file to output the results.								
-v --version	Display the version of VOMA.								
-h --help	Display the help message for the VOMA command.								

For more details, see the VMware Knowledge Base article [2036767](#).

No Failover for Storage Path When TUR Command Is Unsuccessful

A storage path does not fail over when the TUR command repeatedly returns retry requests.

Problem

Typically, when a storage path experiences problems, an ESXi host sends the Test Unit Ready (TUR) command to confirm that the path is down before initiating a path failover. However, if the TUR command is unsuccessful and repeatedly returns a retry operation request (VMK_STORAGE_RETRY_OPERATION), the host continues to retry the command without triggering the failover. Usually, the following errors cause the host to retry the TUR command:

- SCSI_HOST_BUS_BUSY 0x02
- SCSI_HOST_SOFT_ERROR 0x0b
- SCSI_HOST_RETRY 0x0c

Cause

To resolve this issue, you can use the `enable|disable_action_OnRetryErrors` parameter. When you enable this parameter, the ESXi host can mark the problematic path as dead. After marking the path as dead, the host can trigger the failover and use an alternative working path.

Solution

- 1 Set the parameter by running an appropriate command:

Action	Command
Enable the ability to mark a problematic path as dead	<code># esxcli storage nmp satp generic deviceconfig set -c enable_action_OnRetryErrors -d naa.XXX</code>
Disable the ability to mark a problematic path as dead	<code># esxcli storage nmp satp generic deviceconfig set -c disable_action_OnRetryErrors -d naa.XXX</code>

- 2 Check the status of the parameter by running the following command:

```
# esxcli storage nmp device list
```

The following example output indicates that the parameter has been enabled:

```
naa.XXX
Device Display Name: DGC Fibre Channel Disk (naa.XXX)
Storage Array Type: VMW_SATP_CX Storage Array Type Device
Config: {navireg ipfilter action_OnRetryErrors}
```

The enable|disable_action_OnRetryErrors parameter is persistent across reboots.

You can also set this parameter when configuring an SATP claim rule:

```
# esxcli storage nmp satp rule add -t device -d naa.XXX -s VMW_SATP_EXAMPLE -P VMW_PSP_FIXED -o enable_action_OnRetryErrors
```

Troubleshooting Flash Devices

vSphere uses flash drives for such storage features as vSAN, host swap cache, and Flash Read Cache.

The troubleshooting topics can help you avoid potential problems and provide solutions for issues that you might encounter when configuring flash drives.

Formatted Flash Devices Might Become Unavailable

A local flash device becomes unavailable for virtual flash resource or Virtual SAN configuration when it is formatted with VMFS or any other file system.

Problem

When you attempt to configure either vSAN or virtual flash resource, a local flash disk does not appear on the list of disks to be used.

Cause

This problem might occur when flash disk intended for use with either feature has been already formatted with VMFS. Virtual flash and vSAN cannot share the flash disk with VMFS or any other file system.

Also, because virtual flash and vSAN are mutually exclusive consumers of flash disks, both features cannot share a flash disk. If the flash disk is already claimed by one feature, for example vSAN, you cannot use it for another, such as virtual flash, unless you release the disk.

Solution

Use only unformatted flash disks for virtual flash resource and vSAN configuration.

- Avoid formatting the flash disks with VMFS during ESXi installation or Auto Deploy.
- If the flash disk is already formatted with VMFS, remove the VMFS datastore. For information, see the *vSphere Storage* documentation.

- To use the flash disk as a virtual flash resource, do not claim this disk for vSAN. If the disk is claimed by vSAN, remove the disk from vSAN. The flash disk is released from vSAN and becomes available on the list of disks to use with virtual flash. For information about removing disks from vSAN, see the *Administering VMware vSAN* documentation.
- If you intend to use the flash disk with vSAN, do not use the disk for a virtual flash resource. If the flash disk is used as the virtual flash resource, remove the virtual flash configuration. The disk becomes available for vSAN. See the *vSphere Storage* documentation.

Another reason that makes flash disk unavailable is when ESXi cannot detect the disk. See [“Local Flash Disks Are Undetectable,”](#) on page 67.

Keeping Flash Disks VMFS-Free

If you use the auto-partitioning boot option when installing or auto-deploying ESXi, the auto-partitioning option creates a VMFS datastore on your host's local storage. In certain cases, you need to keep your local storage flash disks unformatted.

Problem

By default, auto-partitioning deploys VMFS file systems on any unused local storage disks on your host, including flash disks.

However, a flash disk formatted with VMFS becomes unavailable for such features as virtual flash and vSAN. Both features require an unformatted flash disk and neither can share the disk with any other file system.

Solution

To ensure that auto-partitioning does not format the flash disk with VMFS, use the following boot options when you install ESXi or boot the ESXi host for the first time:

- `autoPartition=TRUE`
- `skipPartitioningSsds=TRUE`

If you use Auto Deploy, set these parameters on a reference host.

- 1 In the vSphere Web Client, select the host to use as a reference host and click the **Configure** tab.
- 2 Click **System** to open the system options, and click **Advanced System Settings**.
- 3 Scroll to `VMkernel.Boot.autoPartition` and set the value to true.
- 4 Scroll to `VMkernel.Boot.skipPartitioningSsds` and set the value to true.
- 5 Reboot the host.

If flash disks that you plan to use with Flash Read Cache and vSAN already have VMFS datastores, remove the datastores.

Local Flash Disks Are Undetectable

If you query for local flash disks, the ESXi host might not return a complete list of the local flash disks.

Problem

ESXi might not be able to detect flash disks, or recognize them as local. This problem can occur when you configure entities that require only local flash disks, for example, virtual flash resource or vSAN.

Cause

ESXi does not recognize certain devices as flash disks when their vendors do not support automatic flash disk detection. In other cases, some flash disks might not be detected as local, and ESXi marks them as remote. When the host does not recognize the disks as the local flash disks, it excludes them from the list of disks available for configuration.

Solution

You might need to tag the disks as flash or as local.

- If ESXi does not automatically recognize its disks as flash disks, tag them as flash disk disks.
- If ESXi does not detect flash disks as local, manually set them as local.

Mark Storage Devices as Flash

If ESXi does not recognize its devices as flash, mark them as flash devices.

ESXi does not recognize certain devices as flash when their vendors do not support automatic flash disk detection. The Drive Type column for the devices shows HDD as their type.



CAUTION Marking the HDD devices as flash might deteriorate the performance of datastores and services that use them. Mark the devices only if you are certain that they are flash devices.

Prerequisites

Verify that the device is not in use.

Procedure

- 1 Browse to the host in the vSphere Web Client object navigator.
- 2 Click the **Configure** tab.
- 3 Under **Storage**, click **Storage Devices**.
- 4 From the list of storage devices, select one or several HDD devices to mark as flash devices and click the **Mark as Flash Disks** (F) icon.
- 5 Click **Yes** to save your changes.

The type of the devices changes to flash.

What to do next

If the flash device that you mark is shared among multiple hosts, make sure that you mark the device from all hosts that share the device.

Mark Storage Devices as Local

ESXi enables you to mark devices as local. This action is useful in cases when ESXi is unable to determine whether certain devices are local.

Prerequisites

- Make sure that the device is not shared.
- Power off virtual machines that reside on the device and unmount an associated datastore.

Procedure

- 1 Browse to the host in the vSphere Web Client object navigator.
- 2 Click the **Configure** tab.
- 3 Under **Storage**, click **Storage Devices**.
- 4 From the list of storage devices, select one or several remote devices to mark as local and click the **All Actions** icon.
- 5 Click **Mark as Local**, and click **Yes** to save your changes.

Troubleshooting Virtual Volumes

Virtual volumes are encapsulations of virtual machine files, virtual disks, and their derivatives. Virtual volumes are stored natively inside a storage system that is connected through Ethernet or SAN. They are exported as objects by a compliant storage system and are managed entirely by hardware on the storage side.

For information about the Virtual Volumes functionality, see the *vSphere Storage* publication.

Virtual Volumes and esxcli Commands

You can use the `esxcli storage vvol` commands to troubleshoot your Virtual Volumes environment.

The following command options are available:

Table 7-5. `esxcli storage vvol` commands

Namespace	Command Option	Description
<code>esxcli storage core device</code>	<code>list</code>	Identify protocol endpoints. The output entry <code>Is VVOL PE: true</code> indicates that the storage device is a protocol endpoint.
<code>esxcli storage vvol daemon</code>	<code>unbindall</code>	Unbind all virtual volumes from all VASA providers known to the ESXi host.
<code>esxcli storage vvol protocolendpoint</code>	<code>list</code>	List all protocol endpoints that your host can access.
<code>esxcli storage vvol storagecontainer</code>	<code>list</code> <code>abandonedvvol scan</code>	List all available storage containers. Scan the specified storage container for abandoned VVols.
<code>esxcli storage vvol vasacontext</code>	<code>get</code>	Show the VASA context (VC UUID) associated with the host.
<code>esxcli storage vvol vasaprovider</code>	<code>list</code>	List all storage (VASA) providers associated with the host.

Virtual Datastore Is Inaccessible

After you create a virtual datastore, it remains inaccessible.

Problem

The vSphere Web Client shows the datastore as inaccessible. You cannot use the datastore for virtual machine provisioning.

Cause

This problem might occur when you fail to configure protocol endpoints for the SCSI-based storage container that is mapped to the virtual datastore. Like traditional LUNs, SCSI protocol endpoints need to be configured so that an ESXi host can detect them.

Solution

Before creating virtual datastores for SCSI-based containers, make sure to configure protocol endpoints on the storage side.

Failures When Migrating VMs or Deploying VM OVF to Virtual Volumes Datastores

Your attempts to migrate a virtual machine or to deploy a VM OVF to virtual datastores fail.

Problem

An OVF template or a VM being migrated from a nonvirtual datastore might include additional large files, such as ISO disk images, DVD images, and image files. If these additional files cause the configuration virtual volume to exceed its 4-GB limit, migration or deployment to a virtual datastore fails.

Cause

The configuration virtual volume, or config-VVol, contains various VM-related files. On traditional nonvirtual datastores, these files are stored in the VM home directory. Similar to the VM home directory, the config-VVol typically includes the VM configuration file, virtual disk and snapshot descriptor files, log files, lock files, and so on.

On virtual datastores, all other large-sized files, such as virtual disks, memory snapshots, swap, and digest, are stored as separate virtual volumes.

Config-VVols are created as 4-GB virtual volumes. Generic content of the config-VVol usually consumes only a fraction of this 4-GB allocation, so config-VVols are typically thin-provisioned to conserve backing space. Any additional large files, such as ISO disk images, DVD images, and image files, might cause the config-VVol to exceed its 4-GB limit. If such files are included in an OVF template, deployment of the VM OVF to vSphere Virtual Volumes storage fails. If these files are part of an existing VM, migration of that VM from a traditional datastore to vSphere Virtual Volumes storage also fails.

Solution

- For VM migration. Before migrating a VM from a traditional datastore to a virtual datastore, remove excess content from the VM home directory to keep the config-VVol under the 4-GB limit.
- For OVF deployment. Because you cannot deploy an OVF template that contains excess files directly to a virtual datastore, first deploy the VM to a nonvirtual datastore. Remove any excess content from the VM home directory, and migrate the resulting VM to vSphere Virtual Volumes storage.

Failed Attempts to Migrate VMs with Memory Snapshots to and from Virtual Datastores

When you attempt to migrate a VM with hardware version 10 or earlier to and from a vSphere Virtual Volumes datastore, failures occur if the VM has memory snapshots.

Problem

The following problems occur when you migrate a version 10 or earlier VM with memory snapshots:

- Migration of a version 10 or earlier VM with memory snapshots to a virtual datastore is not supported and causes a failure.
- Migration of a version 10 or earlier VM with memory snapshots from a virtual datastore to a nonvirtual datastore, such as VMFS, can succeed. If you later make additional snapshots and attempt to migrate this VM back to vSphere Virtual Volumes storage, your attempt fails.

Cause

vSphere Virtual Volumes storage does not require that you use a particular hardware version for your virtual machines. Typically, you can move a virtual machine with any hardware version to vSphere Virtual Volumes storage. However, if you have a VM with memory snapshots, and plan to migrate this VM between a virtual datastore and a nonvirtual datastore, use the VM of hardware version 11.

Non-VVols virtual machines of hardware version 11 or later use separate files to store their memory snapshots. This usage is consistent with VMs on vSphere Virtual Volumes storage, where memory snapshots are created as separate VVols instead of being stored as part of a `.vmsn` file in the VM home directory. In contrast, non-VVols VMs with hardware version 10 continue to store their memory snapshots as part of the `.vmsn` file in the VM home directory. As a result, you might experience problems or failures when attempting to migrate these VMs between virtual and nonvirtual datastores.

Solution

To avoid problems when migrating VMs with memory snapshots across virtual and nonvirtual datastores, use hardware version 11. Follow these guidelines when migrating version 10 or earlier VMs with memory snapshots:

- Migrating a version 10 or earlier VM with memory snapshots to a virtual datastore is not supported. The only workaround is to remove all snapshots. Upgrading the hardware version does not solve this problem.
- Migrating a version 10 or earlier VM with memory snapshots from a virtual datastore to a nonvirtual datastore, such as VMFS, can succeed. However, the migration might put the VM in an inconsistent state. The snapshots that were taken on the virtual datastore use the `vmem` object. Any memory snapshots taken after migrating to VMFS are stored in the `.vmsn` file. If you later attempt to migrate this VM back to vSphere Virtual Volumes storage, your attempt fails. As with the previous case, remove all snapshots to work around this problem.

Troubleshooting VAIO Filters

vSphere APIs for I/O Filtering (VAIO) provide a framework that allows third parties to create software components called I/O filters. The filters can be installed on ESXi hosts and can offer additional data services to virtual machines by processing I/O requests that move between the guest operating system of a virtual machine and virtual disks.

For information about I/O filters, see the *vSphere Storage* publication.

Handling I/O Filter Installation Failures

Typically, all ESXi hosts in a cluster have the same set of I/O filters installed. Occasionally, failures might happen during installation.

If an I/O filter installation fails on a host, the system generates events that report the failure. In addition, an alarm on the host shows the reason for the failure. Examples of failures include the following:

- The VIB URL is not accessible from the host.
- The VIB has an invalid format.
- The VIB requires the host to be in maintenance mode for an upgrade or uninstallation.
- The VIB requires the host to reboot after the installation or uninstallation.
- Attempts to put the host in maintenance mode fail because the virtual machine cannot be evacuated from the host.
- The VIB requires manual installation or uninstallation.

vCenter Server can resolve some failures. You might have to intervene for other failures. For example, you might need to edit the VIB URL, manually evacuate or power off virtual machines, or manually install or uninstall VIBs.

Install I/O Filters on a Single ESXi Host

For troubleshooting purposes, you can download an ESXi component of the I/O filter, packaged as a VIB file, and install it on the ESXi host. Use the `esxcli` command to install the VIB file.

When you specify a target server by using `--server=server_name`, the server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see *Getting Started with vSphere Command-Line Interfaces*, or run `esxcli --help` at the vCLI command prompt.

Prerequisites

Install vCLI or deploy the vSphere Management Assistant (vMA) virtual machine. See *Getting Started with vSphere Command-Line Interfaces*. For troubleshooting, run `esxcli` commands in the ESXi Shell.

Procedure

- 1 Install the VIBs by running the following command:

```
esxcli --server=server_name software vib install --depot path_to_VMware_vib_ZIP_file
```

Options for the `install` command allow you to perform a dry run, specify a specific VIB, bypass acceptance-level verification, and so on. Do not bypass verification on production systems. See the *vSphere Command-Line Interface Reference* documentation.

- 2 Verify that the VIBs are installed on your ESXi host.

```
esxcli --server=server_name software vib list
```

Troubleshooting Networking

The troubleshooting topics about networking in vSphere provide solutions to potential problems that you might encounter with the connectivity of ESXi hosts, vCenter Server and virtual machines.

This chapter includes the following topics:

- [“Troubleshooting MAC Address Allocation,”](#) on page 74
- [“The Conversion to the Enhanced LACP Support Fails,”](#) on page 77
- [“Unable to Remove a Host from a vSphere Distributed Switch,”](#) on page 78
- [“Hosts on a vSphere Distributed Switch 5.1 and Later Lose Connectivity to vCenter Server,”](#) on page 79
- [“Hosts on vSphere Distributed Switch 5.0 and Earlier Lose Connectivity to vCenter Server,”](#) on page 80
- [“Alarm for Loss of Network Redundancy on a Host,”](#) on page 81
- [“Virtual Machines Lose Connectivity After Changing the Uplink Failover Order of a Distributed Port Group,”](#) on page 82
- [“Unable to Add a Physical Adapter to a vSphere Distributed Switch That Has Network I/O Control Enabled,”](#) on page 83
- [“Troubleshooting SR-IOV Enabled Workloads,”](#) on page 84
- [“A Virtual Machine that Runs a VPN Client Causes Denial of Service for Virtual Machines on the Host or Across a vSphere HA Cluster,”](#) on page 84
- [“Low Throughput for UDP Workloads on Windows Virtual Machines,”](#) on page 86
- [“Virtual Machines on the Same Distributed Port Group and on Different Hosts Cannot Communicate with Each Other,”](#) on page 88
- [“Attempt to Power On a Migrated vApp Fails Because the Associated Protocol Profile Is Missing,”](#) on page 89
- [“Networking Configuration Operation Is Rolled Back and a Host Is Disconnected from vCenter Server,”](#) on page 90

Troubleshooting MAC Address Allocation

In vSphere, certain restrictions on the range of MAC addresses that can be assigned to virtual machines might cause loss of connectivity or inability to power on workloads.

Duplicate MAC Addresses of Virtual Machines on the Same Network

You encounter loss of packets and connectivity because virtual machines have duplicate MAC addresses generated by vCenter Server.

Problem

The MAC addresses of virtual machines on the same broadcast domain or IP subnet are in conflict, or vCenter Server generates a duplicate MAC address for a newly created virtual machine.

A virtual machine powers on and functions properly, but shares a MAC address with another virtual machine. This situation might cause packet loss and other problems.

Cause

Virtual machines might have duplicate MAC addresses due to several reasons.

- Two vCenter Server instances with identical IDs generate overlapping MAC addresses for virtual machine network adapters.

Each vCenter Server instance has an ID between 0 and 63 that is randomly generated at installation time, but can be reconfigured after installation. vCenter Server uses the instance ID to generate MAC addresses for the network adapters of the machine.

- A virtual machine has been transferred in power-off state from one vCenter Server instance to another in the same network, for example, by using shared storage, and a new virtual machine network adapter on the first vCenter Server receives the freed MAC address.

Solution

- Change the MAC address of a virtual machine network adapter manually.

If you have an existing virtual machine with a conflicting MAC address, you must provide a unique MAC address in the **Virtual Hardware** settings.

- Power off the virtual machine, configure the adapter to use a manual MAC address, and type the new address.
- If you cannot power the virtual machine off for configuration, re-create the network adapter that is in conflict with enabled manual MAC address assignment and type the new address. In the guest operating system, set the same static IP address to the re-added adapter as before.

For information about configuring the network adapters of virtual machines, see the *vSphere Networking* and *vSphere Virtual Machine Administration* documentation.

- If the vCenter Server instance generates the MAC addresses of virtual machines according to the default allocation, VMware OUI, change the vCenter Server instance ID or use another allocation method to resolve conflicts.

NOTE Changing the vCenter Server instance ID or switching to a different allocation scheme does not resolve MAC address conflicts in existing virtual machines. Only virtual machines created or network adapters added after the change receive addresses according to the new scheme.

For information about MAC address allocation schemes and setup, see the *vSphere Networking* documentation.

Solution	Description
Change the vCenter Server ID	<p>You can keep using the VMware OUI allocation scheme if your deployment contains a small number of vCenter Server instances. According to this scheme, a MAC address has the following format:</p> <p>00:50:56:XX:YY:ZZ</p> <p>where 00:50:56 represents the VMware OUI, XX is calculated as (80 + vCenter Server ID), and YY:ZZ is a random number.</p> <p>To change the vCenter Server ID, configure the vCenter Server unique ID option in the Runtime Settings section from the General settings of the vCenter Server instance and restart it.</p> <p>The VMware OUI allocation works with up to 64 vCenter Server instances and is suitable for small scale deployments.</p>
Switch to prefix-based allocation	<p>You can use a custom OUI. For example, for a 02:12:34 locally administered address range, MAC addresses have the form 02:12:34:XX:YY:ZZ. You can use the fourth octet XX to distribute the OUI address space between the vCenter Server instances. This structure results in 255 address clusters, each cluster managed by a vCenter Server instance, and in about 65000 MAC addresses per vCenter Server. For example, 02:12:34:01:YY:ZZ for vCenter Server A, 02:12:34:02:YY:ZZ for vCenter Server B, and so on.</p> <p>Prefix-based allocation is suitable for deployments of a larger scale.</p> <p>For globally unique MAC addresses, the OUI must be registered in IEEE.</p>

- Configure MAC address allocation.
- Apply the new MAC address allocation scheme to an existing virtual machine in its **Virtual Hardware** settings.
 - Power off a virtual machine, configure the adapter to use a manual MAC address, revert to automatic MAC address allocation, and power on the virtual machine.
 - If the virtual machine is in production and you cannot power it off for configuration, after you change the vCenter Server ID or the address allocation scheme, re-create the network adapter in conflict with enabled automatic MAC address assignment. In the guest operating system, set the same static IP address to the re-added adapter as before.

- Enforce MAC address regeneration when transferring a virtual machine between vCenter Server instances by using the virtual machine files from a datastore.
 - a Power off a virtual machine, remove it from the inventory, and in its configuration file (.vmx), set the ethernetX.addressType parameter to **generated**.
X next to ethernet stands for the sequence number of the virtual NIC in the virtual machine.
 - b Import the virtual machine from one vCenter Server system to another by registering the virtual machine from a datastore in the target vCenter Server.

The virtual machine files can reside in a datastore that is shared between the two vCenter Server instances or can be uploaded to a datastore that is accessible only from the target vCenter Server system.

For information about registering a virtual machine from a datastore, see *vSphere Virtual Machine Administration*.
 - c Power on the virtual machines for the first time.

While the virtual machine is starting up, an information icon appears on the virtual machine in the vSphere Web Client.
 - d Right-click the virtual machine and select **Guest OS > Answer Question**.
 - e Select the **I Copied It** option.

The target vCenter Server re-generates the MAC address of the virtual machine. The new MAC address starts with the VMware OUI 00:0c:29 and is based on the BIOS UUID of the virtual machine. The BIOS UUID of the virtual machine is calculated from the BIOS UUID of the host.
- If the vCenter Server and hosts are version 6.0 and later and the vCenter Server instances are connected in Enhanced Linked Mode, migrate virtual machines by using vMotion across vCenter Server systems.

When a virtual machine is migrated across vCenter Server systems, the source vCenter Server adds the MAC address of the virtual machine to a blacklist and does not assign them to other virtual machines.

Attempt to Power On a Virtual Machine Fails Due to a MAC Address Conflict

After you set a certain static MAC address to a virtual machine adapter, you cannot power on the virtual machine.

Problem

In the vSphere Web Client, after you assign a MAC address within the range 00:50:56:40:YY:ZZ – 00:50:56:7F:YY:ZZ to a virtual machine, attempts to power the virtual machine on fail with a status message that the MAC address is in conflict.

00:50:56:XX:YY:ZZ is not a valid static Ethernet address. It conflicts with VMware reserved MACs for other usage.

Cause

You attempt to assign a MAC address which starts with the VMware OUI 00:50:56 and is within the address range allocated for host VMkernel adapters on the vCenter Server system.

Solution

If you want to preserve the VMware OUI prefix, set a static MAC address within the range 00:50:56:00:00:00 – 00:50:56:3F:FF:FF. Otherwise, set an arbitrary MAC address whose prefix is different from the VMware OUI one. For information about the ranges available for static MAC addresses that have the VMware OUI prefix, see the *vSphere Networking* documentation.

The Conversion to the Enhanced LACP Support Fails

Under certain conditions, the conversion from an existing LACP configuration to the enhanced LACP support on a vSphere Distributed Switch 5.5 and later might fail.

Problem

After you upgrade a vSphere distributed switch to version 5.5 and later, when you initiate the conversion to the enhanced LACP support from an existing LACP configuration, the conversion fails at a certain stage of the process.

Cause

The conversion from an existing LACP configuration to the enhanced LACP support includes several tasks for reconfiguring the distributed switch. The conversion might fail because another user might have reconfigured the distributed switch during the conversion. For example, physical NICs from the hosts might have been reassigned to different uplinks or the teaming and failover configuration of the distributed port groups might have been changed.

Another reason for the failure might be that some of the hosts have disconnected during the conversion.

Solution

When the conversion to the enhanced LACP support fails on a certain stage, it is completed only partially. You must check the configuration of the distributed switch and the participating hosts to identify the objects with incomplete LACP configuration.

Check the target configuration that must result from each conversion stage in the order that is listed in the table. When you locate the stage where the conversion has failed, complete its target configuration manually and continue with the stages that follow.

Table 8-1. Steps to Complete the Conversion to the Enhanced LACP Manually

Conversion Stage	Target Configuration State	Solution
1. Create a new LAG.	A newly created LAG must be present on the distributed switch.	Check the LACP configuration of the distributed switch and create a new LAG if there is none.
2. Create a an intermediate LACP teaming and failover configuration on the distributed port groups.	The newly created LAG must be standby that lets you migrate physical NICs to the LAG without losing connectivity.	Check the teaming and failover configuration of the distributed port group. Set the new LAG as standby if it is not. If you do not want to use a LAG to handle the traffic for all distributed port groups, revert the teaming and failover configuration to a state where standalone uplinks are active and the LAG is unused .
3. Reassign physical NICs from standalone uplinks to LAG ports.	All physical NICs from the LAG ports must be reassigned from standalone uplinks to the LAG ports	Check whether physical NICs are assigned to the LAG ports. Assign a physical NIC to every LAG port. NOTE The LAG must remain standby in the teaming and failover order of the distributed port groups while you reassign physical NICs to the LAG ports.
4. Create the final LACP teaming and failover configuration on the distributed port groups.	The final LACP teaming and failover configuration is the following. <ul style="list-style-type: none"> ■ Active: only the new LAG ■ Standby: empty ■ Unused: all standalone uplinks 	Check the teaming and failover configuration of the distributed port group. Create a valid LACP teaming and failover configuration for all distributed port groups for which you want to apply LACP.

For example, suppose you verify that a new LAG has been created on the distributed switch and that an intermediate teaming and failover configuration has been created for the distributed port groups. You continue with checking whether there are physical NICs assigned to the LAG ports. You find out that not all hosts have physical NICs assigned to the LAG ports, and you assign the NICs manually. You complete the conversion by creating the final LACP teaming and failover configuration for the distributed port groups.

Unable to Remove a Host from a vSphere Distributed Switch

Under certain conditions, you might be unable to remove a host from the vSphere distributed switch.

Problem

- Attempts to remove a host from a vSphere distributed switch fail, and you receive a notification that resources are still in use. The notification that you receive might look like the following:

The resource '16' is in use.

```
vDS DSwitch port 16 is still on host 10.23.112.2 connected to MyVM nic=4000 type=vmVnic
```

- Attempts to remove a host proxy switch that still exists on the host from a previous networking configuration fail. For example, you moved the host to a different data center or vCenter Server system, or upgraded the ESXi and vCenter Server software, and created new networking configuration. When trying to remove the host proxy switch, the operation fails because resources on the proxy switch are still in use.

Cause

You cannot remove the host from the distributed switch or delete the host proxy switch because of the following reasons.

- There are VMkernel adapters on the switch that are in use.
- There are virtual machine network adapters connected to the switch.

Solution

Problem	Solution
Cannot remove a host from a distributed switch	<ol style="list-style-type: none"> 1 In the vSphere Web Client, navigate to the distributed switch. 2 On the Configure tab, select More > Ports. 3 Locate all ports that are still in use and check which VMkernel or virtual machine network adapters on the host are still attached to the ports. 4 Migrate or delete the VMkernel and virtual machine network adapters that are still connected to the switch. 5 Use the Add and Manage Hosts wizard in the vSphere Web Client to remove the host from the switch. <p>After the host is removed, the host proxy switch is deleted automatically.</p>
Cannot remove a host proxy switch	<ol style="list-style-type: none"> 1 In the vSphere Web Client, navigate to the host. 2 Delete or migrate any VMkernel or virtual machine network adapters that are still connected to the host proxy switch. 3 Delete the host proxy switch from the Networking view on the host.

Hosts on a vSphere Distributed Switch 5.1 and Later Lose Connectivity to vCenter Server

Hosts on a vSphere Distributed Switch 5.1 and later cannot connect to vCenter Server after a port group configuration.

Problem

After you change the networking configuration of a port group on a vSphere Distributed Switch 5.1 and later that contains the VMkernel adapters for the management network, the hosts on the switch lose connectivity to vCenter Server. In the vSphere Web Client the status of the hosts is nonresponsive.

Cause

On a vSphere Distributed Switch 5.1 and later in vCenter Server that has networking rollback disabled, the port group containing the VMkernel adapters for the management network is misconfigured in vCenter Server and the invalid configuration is propagated to the hosts on the switch.

NOTE In vSphere 5.1 and later, networking rollback is enabled by default. However, you can enable or disable rollbacks at the vCenter Server level. For more information see the *vSphere Networking* documentation.

Solution

- 1 From the Direct Console User Interface (DCUI) to an affected host, use the **Restore vDS** option from the **Network Restore Options** menu to configure the uplinks and the ID of the VLAN for the management network.

The DCUI creates a local ephemeral port and applies the VLAN and uplink configuration to the port. The DCUI changes the VMkernel adapter for the management network to use the new host local port to restore connectivity to vCenter Server.

After the host re-connects to vCenter Server, the vSphere Web Client displays a warning that some hosts on the switch have different networking configuration from the configuration stored in vSphere distributed switch.

- 2 In the vSphere Web Client, configure the distributed port group for the management network with correct settings.

Situation	Solution
You have altered the port group configuration only once	You can roll the configuration of the port group back one step. Right-click the port group, click Restore Configuration , and select Restore to previous configuration .
You have backed up a valid configuration of the port group	You can restore the configuration of the port group by using the backup file. Right-click the port group, click Restore Configuration , and select Restore configuration from a file . You can also restore the configuration for the entire switch, including the port group, from a backup file for the switch.
You have performed more than one configuration step and you do not have a backup file	You must provide valid settings for the port group manually.

For information about networking rollback, recovery, and restore, see the *vSphere Networking* documentation.

- 3 Migrate the VMkernel adapter for the management network from the host local ephemeral port to a distributed port on the switch by using the Add and Manage Hosts wizard.

Unlike distributed ports, the ephemeral local port of the VMKernel has a non-numeric ID.

For information about handling VMkernel adapters through the Add and Manage Hosts wizard, see the *vSphere Networking* documentation.

- 4 Apply the configuration of the distributed port group and VMkernel adapter from vCenter Server to the host.
 - Push the correct configuration of the distributed port group and VMkernel adapter from vCenter Server to the host.
 - a In the vSphere Web Client, navigate to the host.
 - b On the **Configure** tab, click **Networking**.
 - c From the **Virtual switches** list, select the distributed switch and click **Rectify the state of the selected distributed switch on the host**.
 - Wait until vCenter Server applies the settings within the next 24 hours.

Hosts on vSphere Distributed Switch 5.0 and Earlier Lose Connectivity to vCenter Server

Hosts on a vSphere Distributed Switch 5.0 and earlier cannot connect to vCenter Server after a port group configuration.

Problem

After you change the networking configuration of a port group on a vSphere Distributed Switch 5.0 or earlier that contains the VMkernel adapters for the management network, the hosts on the switch lose connectivity to vCenter Server. In the vSphere Web Client the status of the hosts is nonresponsive.

Cause

On a vSphere Distributed Switch 5.0 and earlier in vCenter Server, the port group containing the VMkernel adapters for the management network is misconfigured in vCenter Server and the invalid configuration is propagated to the hosts on the switch.

Solution

- 1 Connect to an affected host by using the vSphere Client.
- 2 Under **Configuration**, select **Networking**.
- 3 In the vSphere Standard Switch view, create a new standard switch if the host does not have a standard switch suitable for the management network.
 - a Click **Add Networking**.
 - b In the Add Network wizard, under Connection Types select **Virtual Machine**, and click **Next**.
 - c Select **Create a vSphere standard switch**.
 - d Under the **Create a vSphere standard switch** section, select one or more unoccupied physical adapters on the host to carry the management traffic and click **Next**.

If all physical adapters are already busy with traffic from other switches, create the switch without a physical network adapter connected. Later, remove the physical adapter for the management network from the proxy switch of the distributed switch and add it to this standard switch.

- e In the Port Group Properties section, type a network label that identifies the port group that you are creating and optionally a VLAN ID.
 - f Click **Finish**.
- 4 In the vSphere Distributed Switch view, migrate the VMkernel adapter for the network to a standard switch.
 - a Select the vSphere Distributed Switch view, and for the distributed switch, click **Manage Virtual Adapters**.
 - b In the Manage Virtual Adapters wizard, select the VMkernel adapter from the list and click **Migrate**.
 - c Select the newly created or another standard switch to migrate the adapter to, and click **Next**.
 - d Enter a network label that is unique in the scope of the host and optionally a VLAN ID for the management network, and click **Next**.
 - e Review the settings on the target standard switch and click **Finish**.
 - 5 In the vSphere Web Client, configure the distributed port group for the management network with correct settings.
 - 6 Migrate the VMkernel adapter for the management network from the standard switch to a port on the distributed switch by using the Add and Manage Hosts wizard.
For information about the Add and Manage Hosts wizard, see the *vSphere Networking* documentation.
 - 7 If you have moved the physical adapter from the proxy switch to the standard switch, you can reattach it to the distributed switch again by using the Add and Manage Hosts wizard.

Alarm for Loss of Network Redundancy on a Host

An alarm reports a loss of uplink redundancy on a vSphere standard or a distributed switch for a host.

Problem

No redundant physical NICs for a host are connected to a particular standard or a distributed switch, and the following alarm appears:

Host name or IP Network uplink redundancy lost

Cause

Only one physical NIC on the host is connected to a certain standard or a distributed switch. The redundant physical NICs are either down or are not assigned to the switch.

For example, assume that a host in your environment has physical NICs *vmnic0* and *vmnic1* connected to *vSwitch0*, and the physical NIC *vmnic1* goes offline, leaving only *vmnic0* connected to *vSwitch0*. As a result, the uplink redundancy for *vSwitch0* is lost on the host.

Solution

Check which switch has lost uplink redundancy on the host. Connect at least one more physical NIC on the host to this switch and reset the alarm to green. You can use the vSphere Web Client or the ESXi Shell.

If a physical NIC is down, try to bring it back up by using the ESXi Shell on the host.

For information about using the networking commands in the ESXi Shell, see *vSphere Command-Line Interface Reference*. For information about configuring networking on a host in the vSphere Web Client, see *vSphere Networking*.

Virtual Machines Lose Connectivity After Changing the Uplink Failover Order of a Distributed Port Group

Changes in the failover NIC order on a distributed port group cause the virtual machines associated with the group to disconnect from the external network.

Problem

After you rearrange the uplinks in the failover groups for a distributed port group in vCenter Server, for example, by using the vSphere Web Client, some virtual machines in the port group can no longer access the external network.

Cause

After changing the failover order, many reasons might cause virtual machines to lose connectivity to the external network.

- The host that runs the virtual machines does not have physical NICs associated with the uplinks that are set to active or standby. All uplinks that are associated with physical NICs from the host for the port group are moved to unused.
- A Link Aggregation Group (LAG) that has no physical NICs from the host is set as the only active uplink according to the requirements for using LACP in vSphere.
- If the virtual machine traffic is separated in VLANs, the host physical adapters for the active uplinks might be connected to trunk ports on the physical switch that do not handle traffic from these VLANs.
- If the port group is configured with IP hash load balancing policy, an active uplink adapter is connected to a physical switch port that might not be in an EtherChannel.

You can examine the connectivity of the virtual machines in the port group to associated host uplinks and uplink adapters from the central topology diagram of the distributed switch or from the proxy switch diagram for the host.

Solution

- Restore the failover order with the uplink that is associated with a single physical NIC on the host back to active.
- Create a port group with identical settings, make it use the valid uplink number for the host, and migrate the virtual machine networking to the port group.
- Move the NIC to an uplink that participates in the active failover group.

You can use the vSphere Web Client to move the host physical NIC to another uplink.

- Use the Add and Manage Hosts wizard on the distributed switch.
 - a Navigate to the distributed switch in the vSphere Web Client.
 - b From the **Actions** menu select **Add and Manage Hosts**.
 - c On the **Select task** page, select the **Manage host networking** option and select the host.
 - d To assign the NIC of the host to an active uplink, navigate to the **Manage physical network adapters** page and associate the NIC to the switch uplink.
- Move the NIC at the level of the host.
 - a Navigate to the host in the vSphere Web Client, and on the **Configure** tab, expand the **Networking** menu.
 - b Select **Virtual Switches** and select the distributed proxy switch.
 - c Click **Manage the physical network adapters connected to the selected switch**, and move the NIC to the active uplink

Unable to Add a Physical Adapter to a vSphere Distributed Switch That Has Network I/O Control Enabled

You might be unable to add a physical adapter with low speed, for example, 1 Gbps, to a vSphere Distributed Switch that has vSphere Network I/O Control version 3 configured.

Problem

You try to add a physical adapter with low speed, for example, 1 Gbps, to a vSphere Distributed Switch that is connected to physical adapters with high speed, for example, 10 Gbps. Network I/O Control version 3 is enabled on the switch and bandwidth reservations exist for one or more system traffic types, such as vSphere management traffic, vSphere vMotion traffic, vSphere NFS traffic, and so on. The task for adding the physical adapter fails with a status message that a parameter is incorrect.

A specified parameter was not correct: spec.host[.].backing.pnicSpec[.]

Cause

Network I/O Control aligns the bandwidth that is available for reservation to the 10-Gbps speed of the individual physical adapters that are already connected to the distributed switch. After you reserve a part of this bandwidth, adding a physical adapter whose speed is less than 10 Gbps might not meet the potential needs of a system traffic type.

For information about Network I/O Control version 3, see the *vSphere Networking* documentation.

Solution

- 1 In the vSphere Web Client, navigate to the host.
- 2 On the **Configure** tab, expand the **System** group of settings.
- 3 Select **Advanced System Settings** and click **Edit**.
- 4 Type the physical adapters that you want to use outside the scope of Network I/O Control as a comma-separated list for the `Net.IOControlPnicOptOut` parameter.

For example: `vmnic2,vmnic3`

- 5 Click **OK** to apply the changes.
- 6 In the vSphere Web Client, add the physical adapter to the distributed switch.

Troubleshooting SR-IOV Enabled Workloads

Under certain conditions, you might experience connectivity or power-on problems with virtual machines that use SR-IOV to send data to physical network adapters.

SR-IOV Enabled Workload Cannot Communicate After You Change Its MAC Address

After you change the MAC address in the guest operating system of an SR-IOV enabled virtual machine, the virtual machine loses connectivity.

Problem

When you connect the network adapter of a virtual machine to an SR-IOV virtual function (VF), you create a passthrough network adapter for the virtual machine. After the (VF) driver in the guest operating system modifies the MAC address for the passthrough network adapter, the guest operating system shows that the change is successful but the VM network adapter loses connectivity. Although the guest operating system shows that the new MAC address is enabled, a log message in the `/var/log/vmkernel.log` file indicates that the operation has failed.

Requested mac address change to *new MAC address* on port *VM NIC port number*, disallowed by vswitch policy.

where

- *new MAC address* is the MAC address in the guest operation system.
- *VM NIC port number* is the port number of the VM network adapter in hexadecimal format.

Cause

The default security policy on the port group to which the passthrough network adapter is connected does not allow changes in the MAC address in the guest operating system. As a result, the networking interface in the guest operating system cannot acquire an IP address and loses connectivity.

Solution

- ◆ In the guest operating system, reset the interface to cause the passthrough network adapter to regain its valid MAC address. If the interface is configured to use DHCP for address assignment, the interface acquires an IP address automatically.

For example, on a Linux virtual machine run the `ifconfig` console command.

```
ifconfig ethX down
ifconfig ethX up
```

where *X* in `ethX` represents the sequence number of the virtual machine network adapter in the guest operating system.

A Virtual Machine that Runs a VPN Client Causes Denial of Service for Virtual Machines on the Host or Across a vSphere HA Cluster

A virtual machine sending Bridge Protocol Data Unit (BPDU) frames, for example, a VPN client, causes some virtual machines connected to the same port group to lose connectivity. The transmission of BPDU frames might also break the connection of the host or of the parent vSphere HA cluster.

Problem

A virtual machine that is expected to send BPDU frames causes the traffic to the external network of the virtual machines in the same port group to be blocked.

If the virtual machine runs on a host that is a part of a vSphere HA cluster, and the host becomes network-isolated under certain conditions, you observe Denial of Service (DoS) on the hosts in the cluster.

Cause

As a best practice, a physical switch port that is connected to an ESXi host has the Port Fast and BPDU guard enabled to enforce the boundary of the Spanning Tree Protocol (STP). A standard or distributed switch does not support STP, and it does not send any BPDU frames to the switch port. However, if any BPDU frame from a compromised virtual machine arrives at a physical switch port facing an ESXi host, the BPDU guard feature disables the port to stop the frames from affecting the Spanning Tree Topology of the network.

In certain cases a virtual machine is expected to send BPDU frames, for example, when deploying VPN that is connected through a Windows bridge device or through a bridge function. If the physical switch port paired with the physical adapter that handles the traffic from this virtual machine has the BPDU guard on, the port is error-disabled, and the virtual machines and VMkernel adapters using the host physical adapter cannot communicate with the external network anymore.

If the teaming and failover policy of the port group contains more active uplinks, the BPDU traffic is moved to the adapter for the next active uplink. The new physical switch port becomes disabled, and more workloads become unable to exchange packets with the network. Eventually, almost all entities on the ESXi host might become unreachable.

If the virtual machine runs on a host that is a part of a vSphere HA cluster, and the host becomes network-isolated because most of the physical switch ports connected to it are disabled, the active master host in the cluster moves the BPDU sender virtual machine to another host. The virtual machine starts disabling the physical switch ports connected to the new host. The migration across the vSphere HA cluster eventually leads to accumulated DoS across the entire cluster.

Solution

- If the VPN software must continue its work on the virtual machine, allow the traffic out of the virtual machine and configure the physical switch port individually to pass the BPDU frames.

Network Device	Configuration
Distributed or standard switch	<p>Set the Forged Transmit security property on the port group to Accept to allow BPDU frames to leave the host and reach the physical switch port.</p> <p>You can isolate the settings and the physical adapter for the VPN traffic by placing the virtual machine in a separate port group and assigning the physical adapter to the group.</p> <p>CAUTION Setting the Forged Transmit security property to Accept to enable a host to send BPDU frames carries a security risk because a compromised virtual machine can perform spoofing attacks.</p>
Physical switch	<ul style="list-style-type: none"> ■ Keep the Port Fast enabled. ■ Enable the BPDU filter on the individual port. When a BPDU frame arrives at the port, it is filtered out. <p>NOTE Do not enable the BPDU filter globally. If the BPDU filter is enabled globally, the Port Fast mode becomes disabled and all physical switch ports perform the full set of STP functions.</p>

- To deploy a bridge device between two virtual machine NICs connected to the same Layer 2 network, allow the BPDU traffic out of the virtual machines and deactivate Port Fast and BPDU loop prevention features.

Network Device	Configuration
Distributed or standard switch	<p>Set the Forged Transmit property of the security policy on the port groups to Accept to allow BPDU frames to leave the host and reach the physical switch port.</p> <p>You can isolate the settings and one or more physical adapters for the bridge traffic by placing the virtual machine in a separate port group and assigning the physical adapters to the group.</p> <p>CAUTION Setting the Forged Transmit security property to Accept to enable bridge deployment carries a security risk because a compromised virtual machine can perform spoofing attacks.</p>
Physical switch	<ul style="list-style-type: none"> ■ Disable Port Fast on the ports to the virtual bridge device to run STP on them. ■ Disable BPDU guard and filter on the ports facing the bridge device.

- Protect the environment from DoS attacks in any case by activating the BPDU filter on the ESXi host or on the physical switch.
 - On a host running ESXi 4.1 Update 3, ESXi 5.0 Patch 04 and later 5.0 releases, and ESXi 5.1 Patch 01 and later, enable the Guest BPDU filter in one of the following ways and reboot the host:
 - In the Advanced System Settings table on the **Configure** tab for the host in the vSphere Web Client, set the Net.BlockGuestBPDU property to **1**.
 - In an ESXi Shell to the host, type the following vCLI command:


```
esxcli system settings advanced set -o /Net/BlockGuestBPDU -i 1
```
 - On a host that does not have the Guest BPDU filter implemented enable the BPDU filter on the physical switch port to the virtual bridge device.

Network Device	Configuration
Distributed or standard switch	Set the Forged Transmit property of the security policy on the port group to Reject .
Physical switch	<ul style="list-style-type: none"> ■ Keep the Port Fast configuration. ■ Enable the BPDU filter on the individual physical switch port. When a BPDU frame arrives at the physical port, it is filtered out. <p>NOTE Do not enable the BPDU filter globally. If the BPDU filter is enabled globally, the Port Fast mode becomes disabled and all physical switch ports perform the full set of STP functions.</p>

Low Throughput for UDP Workloads on Windows Virtual Machines

When a Windows virtual machine in vSphere 5.1 and later transmits large UDP packets, the throughput is lower than expected or is oscillating even when other traffic is negligible.

Problem

When a Windows virtual machine transmits UDP packets larger than 1024 bytes, you experience lower than expected or oscillating throughput even when other traffic is negligible. In case of a video streaming server, video playback pauses.

Cause

For every UDP packet larger than 1024 bytes, the Windows network stack waits for a transmit completion interrupt before sending the next packet. Unlike for earlier releases, vSphere 5.1 and later releases do not provide a transparent workaround of the situation.

Solution

- Increase the threshold in bytes at which Windows changes its behavior for UDP packets by modifying the registry of the Windows guest OS.
 - a Locate the `HKLM\System\CurrentControlSet\Services\Afd\Parameters` registry key.
 - b Add a value with the name `FastSendDatagramThreshold` of type `DWORD` equal to 1500.

For information about fixing this issue in the Windows registry, see <http://support.microsoft.com/kb/235257>.

- Modify the coalescing settings of the virtual machine NIC.

If the Windows virtual machine has a VMXNET3 vNIC adapter, configure one of the following parameters in the `.vmx` file of the virtual machine. Use the vSphere Web Client, or directly modify the `.vmx` file.

Action	Parameter	Value
Increase the interrupt rate of the virtual machine to a higher rate than expected packet rate. For example, if the expected packet rate is 15000 interrupts per second, set the interrupt rate to 16000 interrupts per second. Set the <code>ethernetX.coalescingScheme</code> parameter to rbc and the <code>ethernetX.coalescingParams</code> parameter to 16000 . The default interrupt rate is 4000 interrupts per second.	<code>ethernetX.coalescingScheme</code> <code>ethernetX.coalescingParams</code>	<code>rbc</code> 16000
Disable coalescing for low throughput or latency-sensitive workloads. For information about configuring low-latency workloads, see Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs .	<code>ethernetX.coalescingScheme</code>	<code>disabled</code>
Revert to the coalescing algorithm from earlier ESXi releases. NOTE The ability to revert to the earlier algorithm will not be available in later vSphere releases.	<code>ethernetX.coalescingScheme</code>	<code>calibrate</code>

X next to `ethernet` stands for the sequence number of the vNIC in the virtual machine.

For more information about configuring parameters in the `.vmx` file, see the *vSphere Virtual Machine Administration* documentation.

- Modify ESXi host coalescing settings.

This approach affects all virtual machines and all virtual machine NICs on the host.

You can edit the advanced system settings list for the host in the vSphere Web Client, or by using a vCLI console command on the host from the ESXi Shell.

Action	Parameter in the vSphere Web Client	Parameter for the esxcli system settings advanced set Command	Value
Set a default interrupt rate higher than the expected packet rate. For example, set the interrupt rate to 16000 if 15000 interrupts are expected per second.	Net.CoalesceScheme Net.CoalesceParams	/Net/CoalesceScheme /Net/CoalesceParams	rbc 16000
Disable coalescing for low throughput or latency-sensitive workloads. For information about configuring low-latency workloads, see Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs .	Net.CoalesceDefaultOn	/Net/CoalesceDefaultOn	0
Revert to the coalescing scheme from earlier ESXi releases. NOTE The ability to revert to the earlier algorithm will not be available in later vSphere releases.	Net.CoalesceScheme	/Net/CoalesceScheme	calibrate

For information about configuring a host from the vSphere Web Client, see the *vCenter Server and Host Management* documentation. For information about setting host properties by using a vCLI command, refer to the *vSphere Command-Line Interface Reference* documentation.

Virtual Machines on the Same Distributed Port Group and on Different Hosts Cannot Communicate with Each Other

Under certain conditions, the virtual machines that are on the same distributed port group but on different hosts cannot communicate with each other.

Problem

Virtual machines that reside on different hosts and on the same port group are unable to communicate. Pings from one virtual machine to another have no effect. You cannot migrate the virtual machines between the hosts by using vMotion.

Cause

- There are no physical NICs on some of the hosts assigned to active or standby uplinks in the teaming and failover order of the distributed port group.
- The physical NICs on the hosts that are assigned to the active or standby uplinks reside in different VLANs on the physical switch. The physical NICs in different VLANs cannot see each other and thus cannot communicate with each other.

Solution

- In the topology of the distributed switch, check which host does not have physical NICs assigned to an active or standby uplink on the distributed port group. Assign at least one physical NIC on that host to an active uplink on the port group.
- In the topology of the distributed switch, check the VLAN IDs of the physical NICs that are assigned to the active uplinks on the distributed port group. On all hosts, assign physical NICs that are from the same VLAN to an active uplink on the distributed port group.
- To verify that there is no problem at the physical layer, migrate the virtual machines to the same host and check the communication between them. Verify that inbound and outbound ICMP traffic is enabled in the guest OS. By default ICMP traffic is disabled in Windows Server 2008 and Windows Server 2012.

Attempt to Power On a Migrated vApp Fails Because the Associated Protocol Profile Is Missing

You cannot power on a vApp or virtual machine that you transferred to a data center or a vCenter Server system because a network protocol profile is missing.

Problem

After you cold migrate a vApp or a virtual machine to another data center or vCenter Server system, an attempt to power it on fails. An error message states that a property cannot be initialized or allocated because the network of the vApp or virtual machine does not have an associated network protocol profile.

Cannot initialize property '*property*'. Network '*port group*' has no associated network protocol profile.

Cannot allocate IP address for property '*property*'. Network '*port group*' has no associated network protocol profile.

Cause

By using the OVF environment, the vApp or virtual machine retrieves network settings from a network protocol profile that is associated with the port group of the vApp or virtual machine.

vCenter Server creates such a network protocol profile for you when you install the OVF of a vApp and associates the profile with the port group that you specify during the installation.

The mapping between the protocol profile and port group is valid only in the scope of a data center. When you move the vApp, the protocol profile is not transferred to the target data center because of the following reasons:

- The network settings of the protocol profile might not be valid in the network environment of the target data center.
- A port group that has the same name and is associated with another protocol profile might already exist in the target data center, and vApps and virtual machines might be connected to this group. Replacing the protocol profiles for the port group might affect the connectivity of these vApp and virtual machines.

Solution

- Create a network protocol profile on the target data center or vCenter Server system with the required network settings and associate the protocol profile with the port group to which the vApp or virtual machine is connected. For example, this approach is suitable if the vApp or virtual machine is a vCenter Server extension that uses the vCenter Extension vService.

For information about providing network settings to a vApp or virtual machine from a network protocol profile, see the *vSphere Networking* documentation.

- Use the vSphere Web Client to export the OVF file of the vApp or virtual machine from the source data center or vCenter Server system and deploy it on the target data center or vCenter Server system.

When you use the vSphere Web Client to deploy the OVF file, the target vCenter Server system creates the network protocol profile for the vApp.

For information about managing OVF files in the vSphere Web Client, see the *vSphere Virtual Machine Administration* documentation.

Networking Configuration Operation Is Rolled Back and a Host Is Disconnected from vCenter Server

When you attempt to add or configure networking on a vSphere Distributed Switch on a host, the operation is rolled back and the host is disconnected from vCenter Server.

Problem

In vSphere 5.1 or later, an attempt to perform a networking configuration operation on a vSphere Distributed Switch on a host, such as creating a virtual machine adapter or a port group, causes the host to disconnect from vCenter Server and results in the error message *Transaction has rolled back on the host*.

Cause

Under stressful conditions on a host, that is, if many concurrent networking operations compete for limited resources, the time to perform some of the operations might exceed the default timeout for rollback of network configuration operations on the distributed switch. As a result, these operations are rolled back.

For example, such a condition might come up when you create a VMkernel adapter on a host that has a very high number of switch ports or virtual adapters, all of which consume system resources on the host.

The default timeout to roll an operation back is 30 seconds.

Solution

- Use the vSphere Web Client to increase the timeout for rollback on vCenter Server.

If you encounter the same problem again, increase the rollback timeout with 60 seconds incrementally until the operation has enough time to succeed.

- a On the **Configure** tab of a vCenter Server instance, expand **Settings**.
- b Select **Advanced Settings** and click **Edit**.
- c If the property is not present, add the `config.vpxd.network.rollbackTimeout` parameter to the settings.
- d Type a new value, in seconds, for the `config.vpxd.network.rollbackTimeout` parameter
- e Click **OK**.
- f Restart the vCenter Server system to apply the changes.

- Increase the timeout for rollback by editing the `vpxd.cfg` configuration file.

If you encounter the same problem again, increase the rollback timeout with 60 seconds incrementally until the operation has enough time to succeed.

- a On a vCenter Server instance, navigate to the directory that contains the `vpxd.cfg` configuration file.
 - On a Windows Server operating system, navigate to *vCenter Server home directory*\Application Data\VMware\VMware VirtualCenter.
 - On the vCenter Server Appliance, navigate to `/etc/vmware-vpx`.
- b Open the `vpxd.cfg` file for editing.

- c Under the <network> section, increase the timeout, in the <rollbackTimeout> element.

```
<config>
  <vpxd>
    <network>
      <rollbackTimeout>60</rollbackTimeout>
    </network>
  </vpxd>
</config>
```

- d Save and close the file.
- e Restart the vCenter Server system to apply the changes.

Troubleshooting Licensing

The troubleshooting licensing topics provide solutions to problems that you might encounter as a result of an incorrect or incompatible license setup in vSphere.

This chapter includes the following topics:

- [“Troubleshooting Host Licensing,”](#) on page 93
- [“Unable to Power On a Virtual Machine,”](#) on page 94
- [“Unable to Configure or Use a Feature,”](#) on page 94

Troubleshooting Host Licensing

You might encounter different problems that result from an incompatible or incorrect license configuration of ESXi hosts.

Unable to Assign a License to an ESXi Host

Under certain conditions, you might be unable to assign a license to an ESXi host.

Problem

You try to assign a license to an ESXi host, but you cannot perform the operation and you receive an error message.

Cause

You might be unable to assign a license to an ESXi host because of the following reasons:

- The calculated license usage for the host exceeds the license capacity. For example, you have a vSphere license key with capacity for two CPUs. You try to assign the key to a host that has four CPUs. You cannot assign the license, because the required license usage for the host is greater than the license capacity.
- The features on the host do not match the license edition. For example, you might configure hosts with vSphere Distributed Switch and vSphere DRS while in evaluation mode. Later, you try to assign vSphere Standard license to the hosts. This operation fails because the vSphere Standard edition does not include vSphere Distributed Switch and vSphere DRS.
- The host is connected to a vCenter Server system that is assigned a license that restricts the edition of the license that you want to assign.

Solution

- Assign a license with larger capacity.

- Upgrade the license edition to match the resources and features on the host, or disable the features that do not match the license edition.
- Assign a vSphere license whose edition is compatible with the license edition of vCenter Server.

ESXi Host Disconnects from vCenter Server

An ESXi host might disconnect from vCenter Server or all ESXi hosts might disconnect from vCenter Server at the same time.

Problem

An ESXi host disconnects from vCenter Server when the host evaluation period or license expires. All ESXi hosts disconnect from vCenter Server when the evaluation period or the license of vCenter Server expire. You receive a licensing-related error message both when a single host disconnects and when all hosts disconnect. You cannot add hosts to the vCenter Server inventory. The hosts and the virtual machines on the hosts continue to run.

Cause

- The 60-day evaluation period of the host has expired or the host license has expired.
- The 60-day evaluation period of vCenter Server is expired or the vCenter Server license is expired.

Solution

- Assign a vSphere license to the ESXi host and try to reconnect it to vCenter Server.
- Assign a vCenter Server license to the vCenter Server system.

Unable to Power On a Virtual Machine

You try to power on a virtual machine, but the operation is unsuccessful and you receive an error message.

Problem

You cannot power on a virtual machine on an ESXi host.

Cause

You might be unable to power on a virtual machine because of the following reasons.

- The 60-day evaluation period of the host is expired.
- The license of the host is expired.

Solution

Table 9-1. Power on a Virtual Machine

Cause	Solution
The evaluation period of the host is expired	Assign a vSphere license to the ESXi host
The license of the host is expired	Assign a vSphere license to the ESXi host

Unable to Configure or Use a Feature

You cannot use a feature or change its configuration.

Problem

You cannot use or configure a feature and a licensing-related error message appears.

Cause

The ESXi host or the vCenter Server system is assigned a license that does not support the features that you want to configure.

Solution

Check the licensed features on the ESXi host and on the vCenter Server system. Upgrade the edition of the license assigned to the host or vCenter Server if they do not include the features that you try to configure or use.

Index

A

- access to the FT metadata datastore, losing **15**
- Active Directory **32**
- address already in use, Jetty server **37**
- advanced settings,
 - Disk.SchedNumReqOutstanding **61**
- Advanced Runtime Info **41**
- All Paths Down (APD) **45**
- authentication **57**
- auto-partitioning, prevent flash disk formatting **67**

C

- cannot apply Storage DRS recommendations, troubleshooting **52**
- cannot enable Storage I/O Control **53**
- certificates
 - hosts **37**
 - vCenter Server **37**
- CHAP authentication **57**
- common logs **11**
- compliance failure, host profiles **32**
- config.vpxd.das.electionWaitTimeSec **25**
- converting to the enhanced LACP fails **77**
- custom reverse proxy **32**

D

- database corruption, vSphere Auto Deploy **30**
- datastore clusters, maintenance mode **48**
- datastore heartbeating **41**
- datastores
 - maintenance mode **48**
 - performance charts troubleshooting **53**
- Denial of Service
 - virtual machine, VPN **84**
 - See also* DoS
- disable_action_OnRetryErrors parameter **65**
- Disk.SchedNumReqOutstanding **61**
- distributed switch
 - cannot remove host **78**
 - physical network adapter **83**
- distributed port group, virtual machines cannot communicate **88**
- downloading VIBs, using custom vCenter Server reverse proxy **32**
- DRS does not place FT VM **16**

E

- enable_action_OnRetryErrors parameter **65**
- Enhanced vMotion Compatibility **16**
- ESXi logs **11**
- EVC **16**

F

- failed conversion to the enhanced LACP **77**
- failover, lost virtual machine connectivity **82**
- failure of vSphere Fault Tolerance for powered-on VM **16**
- Fault Tolerance
 - logging **17**
 - troubleshooting **13, 14**
- Feature **94**
- flash devices, marking **68**
- flash disk formatting, prevent during auto-partitioning **67**
- flash drives **66**

H

- Hardware Virtualization (HV) **13, 14**
- HBAs, queue depth **60**
- host profiles, compliance failure **32**
- host certificates **37**
- Host Failures Cluster Tolerates admission control policy **39–41**
- host isolation response **24**
- hosts, no connection with vCenter Server **79, 80**
- hosts overloaded with FT VMs **15**

I

- I/O filters
 - installation failures **71**
 - installing **72**
 - troubleshooting **71**
- iBFT **62**
- identifying symptoms, troubleshooting **9**
- imbalanced placement of FT VMs **15**
- IPv4 **24**
- IPv6 **24**

J

- Jetty server, address already in use **37**

L

license capacity **94**
 License capacity **93**
 License edition **93, 94**
 license key **94**
 local flash disks, undetectable **67**
 local flash disks are unavailable **66**
 logging **11**
 losing access to the FT metadata datastore **15**
 loss of uplink redundancy **81**
 lost virtual machine connectivity, failover **82**
 low throughput, Windows virtual machine **86**
 LUN not visible, SP visibility **56, 57**
 LUN queue depth **59**
 LUN thrashing **59**

M

maintenance mode, datastores **48**
 management node logs **11**
 marking, flash devices **68**
 maximum HBA queue depth **60**
 metadata consistency, checking with VOMA **64**
 monitoring, Storage I/O Control **52, 53**

N

Network I/O Control, physical network adapter,
 failing to add **83**
 network latency for FT VMs **15**
 network partition **41, 42, 44**
 network protocol profiles, powering on a vApp or
 virtual machine fails **89**
 networking
 host disconnected **90**
 transaction rolled backed **90**
 NFS datastores **63**
 no uplink redundancy **81**
 non-ASCII characters **63**

O

orphaned virtual machines, recovering **19**
 outstanding disk requests **61**
 overcommitted host **14**

P

password requirements **31**
 path thrashing **59**
 Percentage of Cluster Resources Reserved
 admission control policy **40**
 performance, problems **57**
 performance charts for datastores,
 troubleshooting viewing performance
 charts **53**
 Permanent Device Loss (PDL) **45**
 Platform Services Controller logs **11**

prevent flash disk formatting during auto-
 partitioning **67**

Primary VM **44, 45**

problem space, troubleshooting **10**

problems
 performance **57**
 visibility **56, 57**

Q

queue depth **60, 61**

R

recommendations for Storage DRS,
 troubleshooting **52**

Red Cluster **39**

remove datastore **42**

reverse proxy, custom **32**

S

SCSI INQUIRY **62**

SCSI reservations, reducing **58**

Secondary VM **44, 45**

slot size **39–41**

software iSCSI adapters, queue depth **61**

SP visibility, LUN not visible **56, 57**

SQL compatibility mode **36**

SR-IOV

MAC address, lost connectivity **84**

virtual machines, troubleshooting **84**

virtual machine, lost connectivity **84**

SSDs **66**

SSL certificates, troubleshooting in vSphere
 HA **38**

storage adapters, troubleshooting **63**

storage devices, display problems **56**

storage devices, marking as local **68**

Storage DRS

affinity rules **51**

cannot apply recommendations **52**

deleting affinity rules **51**

disabled **47**

disabling **50**

faults **51**

OVF templates **50**

placement **51**

recommendations **51**

rule violation **51**

troubleshooting **47**

Storage I/O Control

monitoring **52, 53**

troubleshooting **52, 53**

T

- testing solutions **10**
- TFTP server, vSphere Auto Deploy **28**
- timeout error, vSphere Auto Deploy **25**
- Tomcat service, vCenter Server upgrade failure **35**
- troubleshooting
 - USB devices **18**
 - vCenter Server **35**
 - vSphere Web Client **35, 36**
 - VVols datastore **69**
- troubleshooting guidelines **9**
- troubleshooting licensing **93**
- troubleshooting overview **9**
- troubleshooting problem space **10**
- troubleshooting Fault Tolerance **13**
- Troubleshooting host licensing **93**
- Troubleshooting Hosts **21**
- troubleshooting storage **55**
- troubleshooting Virtual Volumes
 - OVF deployment failures **70**
 - VM migration failures **70**
- Troubleshooting Virtual Machines **13**

U

- unmount datastore **42**
- updated information **7**
- uplink redundancy lost **81**
- USB devices, error messages **18**
- USB passthrough
 - restarting the USB arbitrator **18**
 - troubleshooting **18**
 - troubleshooting device connections **18**

V

- VAIO filters, troubleshooting **71**
- vCenter Server
 - custom reverse proxy **32**
 - no connection with hosts **79, 80**
 - troubleshooting **35**
- vCenter Server certificates **37, 38**
- vCenter Server license **94**
- vCenter Server system does not appear **36**
- vCenter Server upgrade fails, Tomcat service **35**
- VDS, cannot remove host **78**
- VIB download failure **32**
- Virtual Volumes
 - troubleshooting **69**
 - VM migration failures **70**
- virtual datastore, troubleshooting **69**
- virtual datastores
 - OVF deployment failures **70**
 - VM migration failures **70**
- virtual flash, local flash disks unavailable **66**

- virtual machine
 - BPDU **84**
 - bridge **84**
 - Denial of Service **84**
 - duplicate MAC addresses **74**
 - low throughput **86**
 - MAC address **74**
 - MAC address conflict **74**
 - MAC address, out of range **76**
 - SR-IOV, lost connectivity **84**
 - UDP **86**
 - VPN **84**
 - Windows **86**
- virtual machine console does not launch, vSphere Web Client **37**
- Virtual Machine Protection State **43**
- virtual machines, SR-IOV, troubleshooting **84**
- virtual machines, orphaned **19**
- visibility problems **56, 57**
- VM Component Protection **45**
- VM-Host affinity rules **43**
- VMCP **45**
- VMFS, checking metadata consistency **64**
- VMware On-disk Metadata Analyzer, See VOMA
- VOMA **64**
- vSAN **24, 41, 43**
- vSphere distributed switch
 - hosts not responding **79, 80**
 - lost virtual machine connectivity **82**
- vSphere license **94**
- vSphere Auto Deploy
 - coredump **27**
 - DHCP address **29**
 - failing to complete boot **27**
 - failure to boot **28**
 - image profile warning **26**
 - network boot problem **30**
 - redirection problem **26**
 - TFTP server **28**
 - timeout error **25**
 - troubleshooting **25**
 - wrong image **26**
- vSphere DRS **43**
- vSphere Fault Tolerance **44, 45**
- vSphere HA
 - Denial of Service **84**
 - troubleshooting SSL certificates **38**
- vSphere HA admission control **39**
- vSphere HA Admission Control **39**
- vSphere HA agent **21–24**
- vSphere HA cluster **44**
- vSphere HA failovers **42**

- vSphere HA host state
 - Agent Uninitialized **22**
 - Agent Unreachable **21**
 - Host Failed **23**
 - Initialization Error **22**
 - Network Isolated **24**
 - Network Partitioned **24**
 - Uninitialization Error **23**
- vSphere HA restart failures **43**
- vSphere Web Client
 - troubleshooting **35, 36**
 - virtual machine console does not launch **37**
- VVols **70**
- VVols datastores **70**
- VVols functionality, and esxcli commands **69**

W

- webclient.properties file **37**