

vSphere Troubleshooting

ESXi 5.5
vCenter Server 5.5

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see <http://www.vmware.com/support/pubs>.

EN-001255-00

vmware[®]

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

<http://www.vmware.com/support/>

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 © 2010–2013 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

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

Contents

About vSphere Troubleshooting	5
1 Troubleshooting Virtual Machines	7
Troubleshooting Fault Tolerant Virtual Machines	7
Troubleshooting USB Passthrough Devices	11
Recover Orphaned Virtual Machines	12
Cannot Consolidate Snapshots for Large-Capacity Disks	13
Virtual Machine Does Not Power On After Cloning or Deploying from Template in the vSphere Web Client	13
2 Troubleshooting Hosts	15
Troubleshooting vSphere HA Host States	15
Troubleshooting Auto Deploy	19
Authentication Token Manipulation Error	24
Active Directory Rule Set Error Causes Host Profile Compliance Failure in the vSphere Web Client	25
Unable to Join Domain when Likewise Resources are Low	26
3 Troubleshooting vCenter Server and the vSphere Web Client	27
Troubleshooting vCenter Server	27
Troubleshooting the vSphere Web Client	29
Linked Mode Troubleshooting	31
Troubleshooting vCenter Server and ESXi Host Certificates	34
Troubleshooting vCenter Server Plug-Ins	35
4 Troubleshooting Availability	37
Troubleshooting vSphere HA Admission Control	37
Troubleshooting Heartbeat Datastores	39
Troubleshooting vSphere HA Failover Protection	40
Troubleshooting vSphere Fault Tolerance in Network Partitions	42
5 Troubleshooting Resource Management	45
DRS Troubleshooting Information	45
Troubleshooting Storage DRS	54
Troubleshooting Storage I/O Control	59
6 Troubleshooting Storage	61
Resolving SAN Storage Display Problems	61
Resolving SAN Performance Problems	63
Virtual Machines with RDMs Need to Ignore SCSI INQUIRY Cache	66
Software iSCSI Adapter Is Enabled When Not Needed	67
Failure to Mount NFS Datastores	67

VMkernel Log Files Contain SCSI Sense Codes 68
Troubleshooting Storage Adapters 68
Checking Metadata Consistency with VOMA 69

7 Troubleshooting Networking 71

Duplicate MAC Addresses of Virtual Machines on the Same Network 71
The Conversion to the Enhanced LACP Support Fails 74
Unable to Remove a Host from a vSphere Distributed Switch 75
Hosts on a vSphere Distributed Switch 5.1 and Later Lose Connectivity to vCenter Server 76
Hosts on vSphere Distributed Switch 5.0 and Earlier Lose Connectivity to vCenter Server 77
Alarm for Loss of Network Redundancy on a Host 78
Virtual Machines Lose Connectivity After Changing the Uplink Failover Order of a Distributed Port Group 79
A Virtual Machine that Runs a VPN Client Causes Denial of Service for Virtual Machines on the Host or Across a vSphere HA Cluster 80
Low Throughput for UDP Workloads on Windows Virtual Machines 82
Virtual Machines on the Same Distributed Port Group and on Different Hosts Cannot Communicate with Each Other 83

8 Troubleshooting Licensing 85

Troubleshooting Host Licensing 85
Troubleshooting License Reporting 86
Unable to Power On a Virtual Machine 89
Unable to Assign a License Key to vCenter Server 89
Unable to Configure or Use a Feature 90

Index 91

About vSphere Troubleshooting

vSphere Troubleshooting describes troubleshooting issues and procedures for 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 datacenter operations.

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 7
- [“Troubleshooting USB Passthrough Devices,”](#) on page 11
- [“Recover Orphaned Virtual Machines,”](#) on page 12
- [“Cannot Consolidate Snapshots for Large-Capacity Disks,”](#) on page 13
- [“Virtual Machine Does Not Power On After Cloning or Deploying from Template in the vSphere Web Client,”](#) on page 13

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, 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, this can effect the performance of the Primary VM.

Evidence of this problem could be if the vLockstep Interval on the Primary VM's Fault Tolerance panel is yellow or red. This means that the Secondary VM is running several seconds behind the Primary VM. In such cases, Fault Tolerance slows down the Primary VM. If the vLockstep Interval remains yellow or red for an extended period of time, this is a strong indication that the Secondary VM is not getting enough CPU resources to keep up with the Primary VM.

Cause

A Secondary VM running on a host that is overcommitted for CPU resources might not get the same amount of CPU 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

To resolve this problem, set an explicit CPU reservation for the Primary VM at a 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 are able to execute at a specified rate. For guidance setting this reservation, view the performance graphs of the virtual machine (prior to Fault Tolerance being enabled) to see how much CPU resources it used under normal conditions.

Virtual Machines with Large Memory Can Prevent Use of Fault Tolerance

You can only enable Fault Tolerance on a virtual machine with a maximum of 64GB of memory.

Problem

Enabling Fault Tolerance on a virtual machine with more than 64GB memory can fail. Migrating a running fault tolerant virtual machine using vMotion also can fail if its memory is greater than 15GB or if memory is changing at a rate faster than vMotion can copy over the network.

Cause

This occurs if, due to the virtual machine's memory size, there is not enough bandwidth to complete the vMotion switchover operation within the default timeout window (8 seconds).

Solution

To resolve this problem, before you enable Fault Tolerance, power off the virtual machine and increase its timeout window by adding the following line to the vmx file of the virtual machine:

```
ft.maxSwitchoverSeconds = "30"
```

where 30 is the timeout window in number in seconds. Enable Fault Tolerance and power the virtual machine back on. This solution should work except under conditions of very high network activity.

NOTE If you increase the timeout to 30 seconds, the fault tolerant virtual machine might become unresponsive for a longer period of time (up to 30 seconds) when enabling FT or when a new Secondary VM is created after a failover.

Secondary VM CPU Usage Appears Excessive

In some cases, you might notice that the CPU usage for a Secondary VM is higher than for its associated Primary VM.

Problem

When the Primary VM is idle, the relative difference between the CPU usage of the Primary and Secondary VMs might seem large.

Cause

Replaying events (such as timer interrupts) on the Secondary VM can be slightly more expensive than recording them on the Primary VM. This additional overhead is small.

Solution

None needed. Examining the actual CPU usage shows that very little CPU resource is being consumed by the Primary VM or the Secondary VM.

Primary VM Suffers Out of Space Error

If the storage system you are using has thin provisioning built in, a Primary VM can crash when it encounters an out of space error.

Problem

When used with a thin provisioned storage system, a Primary VM can crash. The Secondary VM replaces the Primary VM, but the error message "There is no more space for virtual disk <disk_name>" appears.

Cause

If thin provisioning is built into the storage system, it is not possible for ESX/ESXi hosts to know if enough disk space has been allocated for a pair of fault tolerant virtual machines. If the Primary VM asks for extra disk space but there is no space left on the storage, the primary VM crashes.

Solution

The error message gives you the choice of continuing the session by clicking "Retry" or clicking "Cancel" to terminate the session. Ensure that there is sufficient disk space for the fault tolerant virtual machine pair and click "Retry".

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.

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.

USB Passthrough Device Is Nonresponsive

USB devices can become nonresponsive for several reasons, including unsafely interrupting a data transfer or if a guest operating system driver sends an unsupported command to the device.

Problem

The USB device is nonresponsive.

Cause

A data transfer was interrupted or nonsupported devices are being used. For example, if a guest driver sends a `SCSI REPORT LUNS` command to unsupported USB flash drives, the device stops responding to all commands.

Solution

- ◆ Physically detach the USB device from the ESXi host and reattach it.

If the host is not physically accessible, you can shut down the host (not reboot) and leave it turned off for at least 30 seconds to ensure that the host USB bus is completely powered off.

When you turn on the host, the USB device is restored from its nonresponsive state.

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 passed through to 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 name.

Problem

Virtual machines that reside on an ESXi host managed by vCenter Server 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 datacenter on which the virtual machine files are stored.

Solution

- 1 Right-click the virtual machine and select **Migrate**.

A list of available hosts appears.

- 2 Click **Change host** and click **Next**.
- 3 Select the host on which to place the virtual machine.

If no hosts are available, add a host that can access the datastore on which the virtual machine's files are stored.

- 4 Click **Finish** to save your changes.

The virtual machine is connected to the new host and appears in the inventory list.

Cannot Consolidate Snapshots for Large-Capacity Disks

If you take at least one snapshot of a virtual machine that has disks larger than 2TB, and you delete one or more of the snapshots, the snapshot files might not consolidate.

Problem

One of the snapshot files, a redo log or child disk, might remain unconsolidated. An unconsolidated file can result in inefficient use of storage, and the unconsolidated child disk can grow over time because of the guest operating system I/O.

Solution

- Delete the snapshots while the virtual machine is turned off.
Plan for routine virtual machine down time to remove unconsolidated files to avoid this problem.
- Alternatively, you can turn off the virtual machine and forcibly consolidate the unconsolidated disks by using the vSphere Web Client. See the *Virtual Machine Administration* documentation.

Virtual Machine Does Not Power On After Cloning or Deploying from Template in the vSphere Web Client

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 can select the **Power on this virtual machine after creation** check box on the Ready to Complete page. However, the virtual machine might not automatically power on upon 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 dropdown 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 of 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 **Related Objects** tab and click the **Virtual Machines** tab.
 - c For each virtual machine to move, right-click the virtual machine and select **Migrate**.
 - d Select **Change datastore**.
 - 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 **Manage** tab and click **Settings**.
 - c Under Virtual Machines, select **Virtual Machine Swapfile 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 Click **Selected 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 15
- [“Troubleshooting Auto Deploy,”](#) on page 19
- [“Authentication Token Manipulation Error,”](#) on page 24
- [“Active Directory Rule Set Error Causes Host Profile Compliance Failure in the vSphere Web Client,”](#) on page 25
- [“Unable to Join Domain when Likewise Resources are Low,”](#) on page 26

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 from contacting the master host and 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 agent on the host has failed, and the watchdog process was unable to restart it.

Solution

Determine if vCenter Server is reporting the host as not responding. If so, there is a networking problem or a total cluster failure. After either 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.

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. 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 8192. 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 5.0 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 15 or “vSphere HA Agent is in the Uninitialized State,” on page 16 for possible solutions.
Lack of file space	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 Secondary VM after a failure (see ["Primary VM Remains in the Need Secondary State,"](#) on page 42).

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 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, vSphere HA applies the power-off or shutdown host isolation response to virtual machines running on the host. vSphere HA continues to monitor the virtual machines that are left powered on. While a host is in this state, vSphere HA's ability to restart virtual machines after a failure is affected. vSphere HA only powers off or shuts down a virtual machine if the agent on the host determines that a master host is responsible for the virtual machine.

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 Virtual SAN 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 Virtual SAN 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.

Troubleshooting Auto Deploy

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

Auto Deploy TFTP Timeout Error at Boot Time

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

Problem

A TFTP Timeout error message appears when a host provisioned by 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.

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` 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 Auto Deploy Server

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

Problem

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

Cause

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

Solution

- ◆ Correct the IP address of the 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 Auto Deploy Host

When you run a PowerCLI cmdlet that assigns an image profile that is not 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 Auto Deploy. You get the error if you attempt to write an 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 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 Auto Deploy.

Solution

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

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

If your 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 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
```

Auto Deploy Host Reboots After Five Minutes

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

Problem

A host to be provisioned with 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-ESXiImageProfile 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 temporarily assign an image profile to the host by running the `Apply-EsxImageProfile` cmdlet.

You can permanently assign an image profile to the host as follows.

- 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.

Auto Deploy Host Cannot Contact TFTP Server

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

Problem

When you attempt to boot a host provisioned with 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.

Auto Deploy Host Cannot Retrieve ESXi Image from Auto Deploy Server

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

Problem

When you attempt to boot a host provisioned with 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 Auto Deploy server.

Cause

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

Solution

- 1 Log in to the system on which you installed the Auto Deploy server.
- 2 Check that the 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 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 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`.

Auto Deploy Host Does Not Get a DHCP Assigned Address

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

Problem

When you attempt to boot a host provisioned with Auto Deploy, the host performs a network boot but is not assigned a DHCP address. The 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 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.

Auto Deploy Host Does Not Network Boot

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

Problem

When you attempt to boot a host provisioned with 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.

If you have an EFI host, you must switch the EFI system to BIOS compatibility mode.
- 2 In the BIOS configuration, enable Network Boot in the Boot Device configuration.

Problems if You Upgrade vCenter Server But Do Not Upgrade Auto Deploy Server

When you upgrade vCenter Server, you can upgrade the Auto Deploy Server at the same time. If you postpone the update, problems with the vSphere HA agent might result.

Problem

When you upgrade vCenter Server, vCenter Server replaces the vSphere HA agent (vmware-fdm) version 5.0 with vSphere HA agent version 5.1 or later on each ESXi host. On hosts provisioned with Auto Deploy, the replacement is not permanent because no state is on the host. If vCenter Server is not available, the ESXi hosts do not have the correct vSphere HA agent and cannot join a cluster.

Cause

The Auto Deploy 5.0 server does not automatically upgrade the FDM VIB to version 5.1 or later. Unless you create a new image that includes the VIB, Auto Deploy reverts to the FDM VIB version 5.0 after reboot.

Solution

Upgrade the Auto Deploy server.

If you cannot upgrade the Auto Deploy server, you can use Image Builder PowerCLI cmdlets included with vSphere PowerCLI to create an ESXi 5.0 image profile that includes the new vmware-fdm VIB. You can supply your hosts with that image profile.

- 1 At the PowerCLI prompt, add the ESXi 5.0 software depot and add the software depot that contains the new vmware-fdm VIB.

```
Add-EsxSoftwareDepot
C:\Path\VMware-Esxi-5.0.0-buildnumber-depot.zip
```

```
Add-EsxSoftwareDepot http://vcenter_server/vSphere-HA-depot
```

- 2 Create a rule that assigns the new image profile to your hosts, and add the rule to the ruleset.

```
New-DeployRule -Name "Rule Name"
-Item "ImageName"
-Pattern "my host pattern"
Add-DeployRule -DeployRule "Rule Name"
```

- 3 Perform a test-and-repair compliance operation for the hosts to permanently include the vSphere HA agent on the hosts.

```
$result = Test-DeployRuleSetCompliance Host_List
Repair-DeployRuleSetCompliance -TestResult $result
```

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 error message appears: A general system error occurred: passwd: Authentication token manipulation error.

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

You can also use a passphrase, which is a phrase consisting of at least three words, each of which is 8 to 40 characters long.

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

Active Directory Rule Set Error Causes Host Profile Compliance Failure in the vSphere Web Client

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 **Rules and Profiles > Host Profiles** on the vSphere Web Client Home page.
- 2 Right-click the host profile and select **Edit Host Profile**.
- 3 Click **Next**.
- 4 Select **Security and Services > Firewall Configuration > Firewall Configuration > Ruleset Configuration > activeDirectoryAll**.
- 5 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.
- 6 Click **Finish**.

Unable to Join Domain when Likewise Resources are Low

You are unable to add a host to an Active Directory domain or you are unable to list domain users when you add user permissions if the peak memory reservation for Likewise daemons is exceeded.

Problem

When you try to add a host to an Active Directory domain, the operation fails. Alternatively, the operation succeeds, but you cannot list domain users when you are adding user permissions.

Cause

When there are more than three trusted Active Directory domains, the peak memory reservation for the Likewise plug-in exceeds the default limit specified in the vSphere Web Client. You might be unable to add a host to an Active Directory domain or list domain users when you are adding user permissions until you increase the memory limit for the Likewise plug-in in the system resource pool.

Solution

- 1 Browse to the host in the vSphere Web Client object navigator.
- 2 Select **Manage > Settings > System Resource Allocation**.
- 3 Click **Advanced** to open the list of system resource pools.
- 4 Select **host > vim > vmvisor > plugins > likewise**.
- 5 Click **likewise** and click **Edit**.
- 6 Select **Memory Resources > Limit**, and increase the limit.
- 7 Click **OK**.

Troubleshooting vCenter Server and the vSphere Web Client

3

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 27
- [“Troubleshooting the vSphere Web Client,”](#) on page 29
- [“Linked Mode Troubleshooting,”](#) on page 31
- [“Troubleshooting vCenter Server and ESXi Host Certificates,”](#) on page 34
- [“Troubleshooting vCenter Server Plug-Ins,”](#) on page 35

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.

Configuring Logging for the VMware Inventory Service

Prior to generating a support bundle request, to facilitate better troubleshooting, you should reconfigure the logging level of the VMware Inventory Service to TRACE.

Problem

You might have to change your vCenter Server logging configuration if any of several problems occur when you use the vSphere Web Client. The following problems are possible:

- Loading the inventory tree does not work.
- Client is unable to log into vCenter Server.
- Properties or objects in the client appear out of date or missing.

Solution

- 1 Open `<Inventory Service install location>\lib\server\config\log4j.properties`.
- 2 Change the keys `log4j.logger.com.vmware.vim` and `log4j.appender.LOGFILE.Threshold` to the new log level.

For example, `log4j.logger.com.vmware.vim = TRACE` (or `log4j.appender.LOGFILE.Threshold = TRACE`) sets the Inventory Service logging to trace.

Valid log levels are TRACE, DEBUG, INFO, WARN, ERROR, in increasing order of verbosity.

- 3 Restart the VMware Inventory Service to pick up the new log level.

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?

Error When You Change vCenter Server Appliance Host Name

When you change the host name of the vCenter Server Appliance, a lookup service error appears when you restart the appliance.

Problem

After you change the host name of the vCenter Server Appliance, the following error appears when you restart the appliance: `Failed to connect to VMware Lookup Service. SSL certificate verification failed.`

Cause

You must update the vCenter Server certificate with the new host name.

Solution

- 1 Log in to the vCenter Server Appliance Web interface.
- 2 Click the **Network** tab and click **Address**.
- 3 Change the host name and click **Save Settings**.

You cannot change the host name if the appliance uses DHCP to obtain an address.

- 4 Click the **Admin** tab and click **Toggle certificate setting**.

vCenter Server generates new certificates for systems that use default certificates. For systems that use custom certificates, you must regenerate the certificates manually.

- 5 Click the **System** tab and click **Reboot** to restart the vCenter Server Appliance.

You must restart the appliance, not just the services running on the appliance.

What to do next

If you use custom certificates, manually generate the certificates as described in the *vSphere Security* documentation.

VMware vCenter Management Webservices Service Fails to Start

When you reboot the vCenter Server machine after installing vCenter Server, the VMware VirtualCenter Management Webservices service does not start.

Problem

The VMware VirtualCenter Management Webservices service does not start automatically.

Cause

This problem can occur when vCenter Server and the database are installed on the same machine.

Solution

- ◆ Start the service manually.

Select **Settings > Control Panel > Administrative Tools > Services > VMware VirtualCenter Management Webservices** and start the service. The machine might require several minutes to start the service.

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 releases of vSphere before vSphere 5.1, you log into individual vCenter Server systems with the vSphere Client. Unless you work in Linked Mode, only one instance of vCenter Server appears in the inventory.

In vSphere 5.1 and 5.5, 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 lookup service 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 lookup service as the vSphere Web Client.
The vSphere Web Client discovers only vCenter Server systems that are registered with the same lookup service.

Unable to Start the Virtual Machine Console in the vSphere Web Client

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 7331, 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\vSphere Web Client\
vCenter Server Appliance	/var/lib/vmware/vsphere-client/

Unable to View the Alarm Definitions Tab of a Datacenter

You might be unable to view the alarm definitions for a datacenter object in the vSphere Web Client.

Problem

When you click the **Alarm Definitions** tab of a datacenter, the tab appears darkened by a translucent overlay and no error message appears.

Cause

Inability to view the alarm definitions might be caused by insufficient memory. Problems that occur on the vCenter Server side should result in an error message, but lack of available memory for Adobe Flash Player on the client machine prevents the error notification dialog from appearing.

Solution

- ◆ Verify that your vCenter Server and vSphere Web Client instances are not constrained by insufficient system resources.

For hardware requirements, see *vSphere Installation and Setup*.

Login to vSphere Web Client Fails with a Duplicate Session Error

When you attempt to log in to vCenter Server using the vSphere Web Client, the login fails with a duplicate session error.

Problem

When you attempt to log in using the vSphere Web Client, you see an error message similar to the following:

```
4/19/2012 15:48:45.416 [ERROR] ErrorNotificationManager:
faultCode:Server.Processing.DuplicateSessionDetected faultString:'Detected duplicate HTTP-based
FlexSessions, generally due to the remote host disabling session cookies. Session cookies must
be enabled to manage the client connection correctly.' faultDetail:'null'
```

Cause

This error occurs if the vSphere Web Client service was restarted or upgraded while you had the vSphere Web Client UI open in a Web browser.

Solution

- 1 Close the Web browser completely.
- 2 Relaunch the Web browser.
- 3 Go to the vSphere Web Client URL and log in.

Linked Mode Troubleshooting

If you are having trouble with your Linked Mode group, consider the following points.

When you have multiple vCenter Server instances, each instance must have a working relationship with the domain controller and not conflict with another machine that is in the domain. Conflicts can occur, for example, when you clone a vCenter Server instance that is running in a virtual machine and you do not use sysprep or a similar utility to ensure that the cloned vCenter Server instance has a globally unique identifier (GUID).

If the domain controller is unreachable, vCenter Server might be unable to start. You might be unable to change the Linked Mode configuration of the affected vCenter Server system. If this occurs, resolve the problem with the domain controller and restart vCenter Server. If resolving the problem with the domain controller is impossible, you can restart vCenter Server by removing the vCenter Server system from the domain and isolating the system from its current Linked Mode group.

The DNS name of the machine must match with the actual machine name. Symptoms of machine names not matching the DNS name are data replication problems, ticket errors when trying to search, and missing search results from remote instances.

NOTE Make sure your Windows and network-based firewalls are configured to allow Linked Mode.

Joining a Linked Mode Group

There is correct order of operations for joining a Linked Mode group.

Procedure

- 1 Verify that the vCenter Server domain name matches the machine name. If they do not match, change one or both to make them match.
- 2 Update the URLs to make them compatible with the new domain name and machine name.

If you do not update the URLs, remote instances of vCenter Server cannot reach the vCenter Server system, because the default URL entries are no longer accurate.

- 3 Join the vCenter Server system to a Linked Mode group.

If a vCenter Server instance is no longer reachable by remote instances of vCenter Server, the following symptoms might occur:

- Clients logging in to other vCenter Server systems in the group cannot view the information that belongs to the vCenter Server system on which you changed the domain name because the users cannot log in to the system.
- Any users that are currently logged in to the vCenter Server system might be disconnected.
- Search queries do not return results from the vCenter Server system.

To resolve these problems, make sure that the `Virtualcenter.VimApiUrl` key points to the location where the clients can access the vCenter Server system, and the `Virtualcenter.VimWebServicesUrl` key points to the location where vCenter Server Webservices is installed. For the `Virtualcenter.InstanceName` key, change the value so that the modified name appears in the vCenter Server inventory view.

What to do next

If you cannot join a vCenter Server instance, you can resolve the problem with the following actions:

- Ensure that the machine is grouped into the correct organizational unit in the corresponding domain controller.
- When you install vCenter Server, ensure that the logged in user account has administrator privileges on the machine.
- To resolve trust problems between a machine and the domain controller, remove the machine from the domain and then add it to the domain again.
- To ensure that the Windows policy cache is updated, run the `gpupdate /force` command from the Windows command line. This command performs a group policy update.

If the local host cannot reach the remote host during a join operation, verify the following:

- Remote vCenter Server IP address or fully qualified domain name is correct.

- LDAP port on the remote vCenter Server is correct.
- VMwareVCMSDS service is running.

Configure a Windows Firewall to Allow a Specified Program Access

vCenter Server uses Microsoft ADAM/AD LDS to enable Linked Mode, which uses the Windows RPC port mapper to open RPC ports for replication. When you install vCenter Server in Linked Mode, you must modify the firewall configuration on the local machine .

Incorrect configuration of firewalls can cause licenses and roles to become inconsistent between instances.

Prerequisites

- The Windows version must be earlier than Windows Server 2008. For Windows Server 2008, Windows automatically configures the firewall to permit access.
- No network-based firewalls can exist between vCenter Server Linked Mode instances. For environments with network-based firewalls, see [“Configure Firewall Access by Opening Selected Ports,”](#) on page 33.

Procedure

- 1 Select **Start > Run**.
- 2 Type **firewall.cpl** and click **OK**.
- 3 Make sure that the firewall is set to allow exceptions.
- 4 Click the **Exceptions** tab.
- 5 Click **Add Program**.
- 6 Add an exception for **C:\Windows\ADAM\dsamain.exe** and click **OK**.
- 7 Click **OK**.

Configure Firewall Access by Opening Selected Ports

vCenter Server uses Microsoft ADAM/AD LDS to enable Linked Mode, which uses the Windows RPC port mapper to open RPC ports for replication. When you install vCenter Server in Linked Mode, the firewall configuration on any network-based firewalls must be modified.

Incorrect configuration of firewalls can cause licenses and roles to become inconsistent between instances.

Procedure

- ◆ Configure Windows RPC ports to generically allow selective ports for machine-to-machine RPC communication.

Choose one of the following methods.

- Change the registry settings. See <http://support.microsoft.com/kb/154596/en-us>.
- Use Microsoft's RPCCfg.exe tool. See <http://support.microsoft.com/kb/908472/en-us>.

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 the Database

After you replace default vCenter Server certificates, you might be unable to connect to the vCenter Server database.

Problem

vCenter Server is unable to connect to the vCenter Server database after you replace default vCenter Server certificates, and management web services do not start.

Cause

The database password must be updated in its encrypted form.

Solution

Update the database password by running the following command: `vpxd -P pwd`.

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 vCenter Server Plug-Ins

In cases where vCenter Server plug-ins are not working, you have several options to correct the problem.

vCenter Server plug-ins that run on the Tomcat server have `extension.xml` files, which contain the URL where the corresponding Web application can be accessed. These files are located in `C:\Program Files\VMware\Infrastructure\VirtualCenter Server\extensions`. Extension installers populate these XML files using the DNS name for the machine.

Example from the stats `extension.xml` file: `<url>https://SPULOV-XP-VM12.vmware.com:8443/statsreport/vicr.do</url>`.

vCenter Server, plug-in servers, and the clients that use them must be located on systems under the same domain. If they are not under the same domain, or if the DNS of the plug-in server is changed, the plug-in clients will not be able to access the URL, and the plug-in will not work.

You can edit the XML files manually by replacing the DNS name with an IP address. Reregister the plug-in after you edit its `extension.xml` file.

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 37
- [“Troubleshooting Heartbeat Datastores,”](#) on page 39
- [“Troubleshooting vSphere HA Failover Protection,”](#) on page 40
- [“Troubleshooting vSphere Fault Tolerance in Network Partitions,”](#) on page 42

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.

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 slave host over the management network, the master host uses datastore heartbeating to determine if the slave host might have failed or is in a network partition. If the slave 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 Virtual SAN.

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 or permanent device loss.
- 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 of 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 Failover Protection

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.

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 Virtual SAN 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.

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 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 Resource Management

5

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:

- [“DRS Troubleshooting Information,”](#) on page 45
- [“Troubleshooting Storage DRS,”](#) on page 54
- [“Troubleshooting Storage I/O Control,”](#) on page 59

DRS Troubleshooting Information

This information describes vSphere® Distributed Resource Scheduler (DRS) problems for particular categories: cluster, host, and virtual machine problems.

Cluster Problems

Cluster problems can prevent DRS from performing optimally or from reporting faults.

Load Imbalance on Cluster

A cluster has a load imbalance of resources.

Problem

A cluster might become unbalanced because of uneven resource demands from virtual machines and unequal capacities of hosts.

Cause

The following are possible reasons why the cluster has a load imbalance:

- The migration threshold is too high.
A higher threshold makes the cluster a more likely candidate for load imbalance.
- VM/VM or VM/Host DRS rules prevent virtual machines from being moved.
- DRS is disabled for one or more virtual machines.
- A device is mounted to one or more virtual machines preventing DRS from moving the virtual machine in order to balance the load.
- Virtual machines are not compatible with the hosts to which DRS would move them. That is, at least one of the hosts in the cluster is incompatible for the virtual machines that would be migrated. For example, if host A's CPU is not vMotion-compatible with host B's CPU, then host A becomes incompatible for powered-on virtual machines running on host B.

- It would be more detrimental for the virtual machine's performance to move it than for it to run where it is currently located. This may occur when loads are unstable or the migration cost is high compared to the benefit gained from moving the virtual machine.
- vMotion is not enabled or set up for the hosts in the cluster.

Solution

Address the problem that is causing the load imbalance.

Cluster is Yellow

The cluster is yellow due to a shortage of resources.

Problem

If the cluster does not have enough resources to satisfy the reservations of all resource pools and virtual machines, but does have enough resources to satisfy the reservations of all running virtual machines, DRS continues to run and the cluster is yellow.

Cause

A cluster can become yellow if the host resources are removed from the cluster (for example, if a host fails).

Solution

Add host resources to the cluster or reduce the resource pool reservations.

Cluster is Red Because of Inconsistent Resource Pool

A DRS cluster becomes red when it is invalid. It may become red because the resource pool tree is not internally consistent.

Problem

If the cluster resource pool tree is not internally consistent (for example, the sum of the children's reservations is greater than the parent pool's nonexpandable reservation), the cluster does not have enough resources to satisfy the reservations of all running virtual machines making the cluster red.

Cause

This can occur if vCenter Server is unavailable or if resource pool settings are changed while a virtual machine is in a failover state.

Solution

Revert the associated changes or otherwise revise the resource pool settings.

Cluster is Red Because Failover Capacity is Violated

A DRS cluster becomes red when it is invalid. It may become red because failover capacity is violated.

Problem

The cluster attempts to failover virtual machines in case of host failure, but is not guaranteed to have enough resources available to failover all virtual machines covered by the failover requirements.

Cause

If a cluster enabled for HA loses so many resources that it can no longer fulfill its failover requirements, a message appears and the cluster's status changes to red.

Solution

Review the list of configuration issues in the yellow box at the top of the cluster Summary page and address the issue that is causing the problem.

No Hosts are Powered Off When Total Cluster Load is Low

Hosts are not powered off when the total cluster load is low.

Problem

Hosts are not powered off when the total cluster load is low because extra capacity is needed for HA failover reservations.

Cause

Hosts might not be powered off for the following reasons:

- The `MinPoweredOn{Cpu|Memory}Capacity` advanced options settings need to be met.
- Virtual machines cannot be consolidated onto fewer hosts due to their resource reservations, VM/Host DRS rules, VM/VM DRS rules, not being DRS-enabled, or not being compatible with the hosts having available capacity.
- Loads are unstable.
- DRS migration threshold is at the highest setting and only allows mandatory moves.
- vMotion is unable to run because it is not configured.
- DPM is disabled on the hosts that might be powered off.
- Hosts are not compatible for virtual machines to be moved to another host.
- Host does not have Wake On LAN, IPMI, or iLO technology. Either one is required for DPM to enter a host in standby.

Solution

Address the issue that prevents hosts from being powered off when the total cluster load is low.

Hosts are Powered Off When Total Cluster Load is High

Hosts are powered off when total cluster load is high.

Problem

DRS determined that virtual machines could be run on a fewer number of hosts without degrading the host or virtual machine performance. DRS is also constrained from moving the virtual machines running on the highly-utilized hosts to the hosts scheduled for power-off.

Cause

This occurs when the total cluster load is too high.

Solution

Reduce the cluster load.

DRS Seldom or Never Performs vMotion Migrations

DRS seldom or never performs vMotion migrations.

Problem

DRS does not perform vMotion migrations.

Cause

DRS never performs vMotion migrations when one or more of the following issues is present on the cluster.

- DRS is disabled on the cluster.
- The hosts do not have shared storage.
- The hosts in the cluster do not contain a vMotion network.
- DRS is manual and no one has approved the migration.

DRS seldom performs vMotion when one or more of the following issues is present on the cluster:

- Loads are unstable, or vMotion takes a long time, or both. A move is not appropriate.
- DRS seldom or never migrates virtual machines.
- DRS migration threshold is set too high.

DRS moves virtual machines for the following reasons:

- Evacuation of host that a user requested enter maintenance or standby mode.
- VM/Host DRS rules or VM/VM DRS rules.
- Reservation violations.
- Load imbalance.
- Power management.

Solution

Address the issues that are causing DRS to avoid performing vMotion migrations.

Host Problems

Host problems might cause DRS to not perform as expected.

DRS Recommends Host be Powered On to Increase Capacity When Total Cluster Load Is Low

The host should be powered on to help provide more capacity for the cluster or help hosts that are overloaded.

Problem

DRS recommends that the host be powered on to increase capacity when the total cluster load is low.

Cause

The recommendation might be made because:

- The cluster is a DRS-HA cluster. Additional powered-on hosts are needed in order to provide more failover capability.
- Some hosts are overloaded and virtual machines on currently powered-on hosts can be moved to hosts in standby mode to balance the load.

- The capacity is needed to meet the `MinPoweredOn{Cpu|Memory}Capacity` advanced options.

Solution

Power on the host.

Total Cluster Load Is High

The total cluster load is high.

Problem

When the total cluster load is high, DRS does not power on the host.

Cause

The following are possible reasons why DRS does not power on the host:

- VM/VM DRS rules or VM/Host DRS rules prevent the virtual machine from being moved to this host.
- Virtual machines are pinned to their current hosts, hence DRS cannot move these virtual machines to hosts in standby mode to balance the load.
- DRS or DPM is in manual mode and the recommendations were not applied.
- No virtual machines on highly utilized hosts will be moved to that host.
- DPM is disabled on the host because of a user setting or host previously failing to successfully exit standby.

Solution

Address that issue that prevents DRS from powering on the host.

Total Cluster Load Is Low

The total cluster load is low.

Problem

When the total cluster load is low, DRS does not power off the host.

Cause

The following are possible reasons why DRS does not power off the host:

- Distributed Power Management (DPM) detected better candidates to power off.
- vSphere HA needs extra capacity for failover.
- The load is not low enough to trigger the host to power off.
- DPM projects that the load will increase.
- DPM is not enabled for the host.
- DPM threshold is set too high.
- While DPM is enabled for the host, no suitable power-on mechanism is present for the host.
- DRS cannot evacuate the host.
- The DRS migration threshold is at the highest setting and only performs mandatory moves.

Solution

Address the issue that is preventing DRS from powering off the host.

DRS Does Not Evacuate a Host Requested to Enter Maintenance or Standby Mode

DRS does not evacuate a host requested to enter maintenance mode or standby mode.

Problem

When you attempt to put a host into maintenance or standby mode, DRS does not evacuate the host as expected.

Cause

vSphere HA is enabled and evacuating this host might violate HA failover capacity.

Solution

There is no solution. If appropriate, disable vSphere HA before you attempt to put the host into maintenance mode or standby mode.

DRS Does Not Move Any Virtual Machines onto a Host

DRS does not move any virtual machines onto a host.

Problem

DRS does not recommend migration of virtual machine to a host that has been added to a DRS-enabled cluster.

Cause

After a host has been added to a DRS-enabled cluster, the virtual machines deployed to the host become part of the cluster. DRS can recommend migration of some virtual machines to this host just added to the cluster. If that does not occur, there may be problems with vMotion, host compatibility, or affinity rules. The following are possible reasons:

- vMotion is not configured or enabled on this host.
- Virtual machines on other hosts are not compatible with this host.
- The host does not have sufficient resources for any virtual machine.
- Moving any virtual machines to this host would violate a VM/VM DRS rule or VM/Host DRS rule.
- This host is reserved for HA failover capacity.
- A device is mounted to the virtual machine.
- The vMotion threshold is too high.
- DRS is disabled for the virtual machines, hence the virtual machine could not be moved onto the destination host.

Solution

Address the issue that prevents DRS from moving virtual machines onto a host.

DRS Does Not Move Any Virtual Machines from a Host

DRS does not move any virtual machines from a host.

Problem

Virtual machines are not moved from this host.

Cause

This may be because of problems with vMotion, DRS, or host compatibility. The following are the possible reasons:

- vMotion is not configured or enabled on this host.
- DRS is disabled for the virtual machines on this host.
- Virtual machines on this host are not compatible with any other hosts.
- No other hosts have sufficient resources for any virtual machines on this host.
- Moving any virtual machines from this host would violate a VM/VM DRS rule or VM/Host DRS rule.
- DRS is disabled for one or more virtual machines on the host.
- A device is mounted to the virtual machine.

Solution

Address the issues that are preventing DRS from moving virtual machines from the host.

Virtual Machine Problems

Virtual machine problems might cause DRS to not perform as expected.

Insufficient CPU or Memory Resources

The virtual machine does not receive enough CPU or memory resources.

Problem

In some cases, the virtual machine's demand is greater than its resource entitlement. When this occurs, the virtual machine doesn't receive enough CPU or memory resources.

Cause

The following sections describe the factors that influence the entitlement for a virtual machine.

- | | |
|--|---|
| Cluster is Yellow or Red | If the cluster is yellow or red, the capacity is insufficient to meet the resource reservations configured for all virtual machines and resource pools in the cluster. The particular virtual machine might be one that is not receiving its reservation. Check the status of the cluster (red or yellow) and resolve the situation. |
| Resource Limit is Too Restrictive | The virtual machine, its parent resource pool, or its resource pool ancestors might have a configured resource limit that is too restrictive. Check whether demand is equal to or greater than any configured limits. |
| Cluster is Overloaded | The cluster on which the virtual machine is running might have insufficient resources. Also, the virtual machine's share value is such that other virtual machines are granted proportionally more of the resources. To determine the demand is larger than the capacity, check the cluster statistics. |
| Host is Overloaded | To determine if the host's resources are oversubscribed, check the host statistics. If they are oversubscribed, consider why DRS is not moving any of the virtual machines now running on the host to other hosts. This condition might exist for the following reasons: <ul style="list-style-type: none"> ■ The VM/VM DRS rules and VM/Host DRS rules require the current virtual machine-to-host mapping. If such rules are configured in the cluster, consider disabling one or more of them. Then run DRS and check whether the situation is corrected. |

- DRS cannot move this virtual machine or enough of the other virtual machines to other hosts to free up capacity. DRS will not move a virtual machine for any of the following reasons:
 - DRS is disabled for the virtual machine.
 - A host device is mounted to the virtual machine.
 - Either of its resource reservations is so large that the virtual machine cannot run on any other host in the cluster.
 - The virtual machine is not compatible with any other host in the cluster.

Check whether any of these conditions exist for the virtual machine. If none exist, the conditions might exist for other virtual machines in the cluster. If this is the case, DRS cannot balance the cluster to address the virtual machine's demand.

- Decrease the DRS migration threshold setting and check whether the situation is resolved.
- Increase the virtual machine's reservation.

Solution

Address the problem that is causing the virtual machine to not receive enough CPU or memory resources.

VM/VM DRS Rule or VM/Host DRS Rule Violated

DRS rules specify which host a virtual machine must or must not reside on, or which virtual machines must be or must not be on the same host.

Problem

A VM/VM DRS rule or a VM/Host DRS rule is violated.

Cause

VM/VM DRS rules specify that selected virtual machines should be placed on the same host (affinity) or that virtual machines be placed on different hosts (anti-affinity). VM/Host DRS rules specify that selected virtual machines should be placed on specified hosts (affinity) or that selected virtual machines should not be placed on specified hosts (anti-affinity).

When a VM/VM DRS rule or VM/Host DRS rule is violated, it might be because DRS cannot move some or all of the virtual machines in the rule. The reservation of the virtual machine or other virtual machines in the affinity rule, or their parent resource pools, might prevent DRS from locating all virtual machines on the same host.

Solution

- Check the DRS faults panel for faults associated with affinity rules.
- Compute the sum of the reservations of all the virtual machines in the affinity rule. If that value is greater than the available capacity on any host, the rule cannot be satisfied.
- Compute the sum of the reservations of their parent resource pools. If that value is greater than the available capacity of any host, the rule cannot be satisfied if the resources are obtained from a single host.

Virtual Machine Power On Operation Fails

An error message appears stating that the virtual machine fails to power on.

Problem

The virtual machine fails to power on.

Cause

The virtual machine might fail to power on because of insufficient resources or because there are no compatible hosts for the virtual machine.

Solution

If the cluster does not have sufficient resources to power on a single virtual machine or any of the virtual machines in a group power-on attempt, check the resources required by the virtual machine against those available in the cluster or its parent resource pool. If necessary, reduce the reservations of the virtual machine to be powered-on, reduce the reservations of its sibling virtual machines, or increase the resources available in the cluster or its parent resource pool.

DRS Does Not Move the Virtual Machine

DRS does not move the virtual machine when it is initially powered on despite insufficient resources on the host.

Problem

When you power on a virtual machine, DRS does not migrate it as expected when there are not enough resources on the host where the virtual machine is registered.

Cause

The following are possible reasons why DRS does not move the virtual machine.

- DRS is disabled on the virtual machine.
- The virtual machine has a device mounted.
- The virtual machine is not compatible with any other hosts.
- No other hosts have a sufficient number of physical CPUs or capacity for each CPU for the virtual machine.
- No other hosts have sufficient CPU or memory resources to satisfy the reservations and required memory of this virtual machine.
- Moving the virtual machine will violate an affinity or anti-affinity rule.
- The DRS automation level of the virtual machine is manual and the user does not approve the migration recommendation.
- DRS will not move fault tolerance-enabled virtual machines.

Solution

Address the issue that prevents DRS from moving the virtual machine.

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 in the vSphere Web Client

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 54 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 datacenters.

Storage DRS is not supported on datastores that are shared across multiple datacenters. This configuration can occur when a host in one datacenter mounts a datastore in another datacenter, or when a host using the datastore is moved to a different datacenter. When a datastore is shared across multiple datacenters, 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 datacenters.

- 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 datacenter. Move the hosts to the same datacenter or unmount the datastore from hosts that reside in other datacenters.
- 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 in the vSphere Web Client

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 in the vSphere Web Client

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 in the vSphere Web Client

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 your hosts in the SAN environment. For information about setting up the SAN storage and working with datastores, see the *vSphere Storage* documentation.

This chapter includes the following topics:

- [“Resolving SAN Storage Display Problems,”](#) on page 61
- [“Resolving SAN Performance Problems,”](#) on page 63
- [“Virtual Machines with RDMs Need to Ignore SCSI INQUIRY Cache,”](#) on page 66
- [“Software iSCSI Adapter Is Enabled When Not Needed,”](#) on page 67
- [“Failure to Mount NFS Datastores,”](#) on page 67
- [“VMkernel Log Files Contain SCSI Sense Codes,”](#) on page 68
- [“Troubleshooting Storage Adapters,”](#) on page 68
- [“Checking Metadata Consistency with VOMA,”](#) on page 69

Resolving SAN Storage Display Problems

When you use the vSphere Web Client to display 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 6-1. 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.

Table 6-1. Troubleshooting Fibre Channel LUN Display (Continued)

Troubleshooting Task	Description
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 6-2. 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.	<p>If the expected LUNs do not appear after rescan, access control might not be configured correctly on the storage system side:</p> <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.

Table 6-2. Troubleshooting iSCSI LUN Display (Continued)

Troubleshooting Task	Description
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.	<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 LUN masking. ■ Reconnect a cable. ■ Make a change to a host in a cluster. ■ Change CHAP settings or add new discovery addresses. <p>For information, see the <i>vSphere Storage</i> documentation.</p>

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

Operations that require getting a file lock or a metadata lock in VMFS result in short-lived SCSI reservations. SCSI reservations lock an 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 ESXi hosts use the SCSI reservations mechanism only when storage devices do not support the hardware acceleration. 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.

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.

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 Adjust the `Disk.SchedNumReqOutstanding` parameter, so that it matches the queue depth value.

Adjust Queue Depth for QLogic and Emulex HBAs

If you are not satisfied with your host's performance, change the maximum queue depth for the QLogic or Emulex HBA.

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 one of the following commands:
 - For QLogic: `esxcli --server=server_name system module list |grep qla`
 - For Emulex: `esxcli --server=server_name system module list |grep lpfc`
- 2 Adjust the queue depth for the appropriate module.

NOTE The examples show the QLogic `qla2xxx` and Emulex `lpfc820` modules. Use the appropriate module based on the outcome of the previous step.

- For QLogic:


```
esxcli --server=server_name system module parameters set -m qla2xxx -p ql2xmaxqdepth=value
```
 - For Emulex:


```
esxcli --server=server_name system module parameters set -m lpfc820 -p lpfc0_lun_queue_depth=value
```
- 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 your QLogic or Emulex module, such as `lpfc820` or `qla2xxx`.

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
iscsivmk_LunQDepth=value
```

The `iscsivmk_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.



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

Change Maximum Outstanding Disk Requests in the vSphere Web Client

If you adjusted the LUN queue depth, change the `Disk.SchedNumReqOutstanding` parameter, so that its value matches the queue depth. The parameter controls the maximum number of outstanding 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.

Procedure

- 1 Browse to the host in the vSphere Web Client object navigator.
- 2 Click the **Manage** tab and click **Settings**.
- 3 Under System, click **Advanced System Settings**.
- 4 Scroll to **Disk.SchedNumReqOutstanding** and click the **Edit** icon.
- 5 Change the parameter value to the number of your choice and click **OK**.

Virtual Machines with RDMS Need to Ignore SCSI INQUIRY Cache

Storage vendors might require that virtual machines with RDMS (hardware version 8) 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.

Because this parameter is configurable only on virtual machines with hardware version 8, upgrade the virtual machine prior to adding the parameter.

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 file names 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 file names 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.

VMkernel Log Files Contain SCSI Sense Codes

Certain VMkernel messages related to storage might contain SCSI Sense codes.

Problem

When you analyze ESXi host's `/var/log/vmkernel` log files, you encounter events or error messages that contain SCSI Sense codes.

Solution

Ability to interpret the SCSI Sense codes can help you better understand problems in your storage environment. Because the SCSI Sense code values are assigned by the T10 committee, you need to consult the T10 standards documentation to determine the meaning of the codes. This topic explains how to use the T10 documentation to interpret the SCSI Sense codes.

Example: Interpreting SCSI Sense Codes

The following is an example of a SCSI error message that appears in the ESXi log file:

```
2011-04-04T21:07:30.257Z cpu2:2050)ScsiDeviceIO: 2315: Cmd(0x4124003edb00) 0x12, CmdSN 0x51 to dev "naa.600508XXXXXXXXXXXX" failed H:0x0 D:0x2 P:0x0 Valid sense data: 0x5 0x25 0x0
```

In this example, SCSI Sense codes are represented by two fields, `H:0x0 D:0x2 P:0x0` and `0x5 0x25 0x0`.

The first field, `H:0x0 D:0x2 P:0x0`, is a combination of SCSI Status codes for the three components in your storage environment, the host, the device, and the plug-in. The SCSI Status code is used to determine the success or failure of a SCSI command. To interpret each SCSI Status code, see the <http://www.t10.org/lists/2status.htm>.

NOTE Hexadecimal numbers in the T10 documentation use the `NNNh` format, while SCSI Sense codes in the ESXi log files follow the `0xNNN` format. For example, `0x2 = 02h`.

You will get the following interpretation for the status field of the above example: `H:0x0 D:0x2 P:0x0 = H(host):GOOD D(device):CHECK CONDITION P(plug-in):GOOD`.

The second field in a typical SCSI error message provides more detailed information about the error. It is a combination of Sense Key (`sense`), Additional Sense Code (`asc`), and Additional Sense Code Qualifier (`ascq`) parameters.

For example, the `0x5 0x25 0x0` field from the above error message can be represented as `sense=5 asc=25 ascq=0`.

To interpret Sense Keys, see <http://www.t10.org/lists/2sensekey.htm>.

To determine the meaning of the Additional Sense Code (`asc`) and Additional Sense Code Qualifier (`ascq`), use the two codes together. See <http://www.t10.org/lists/2asc.htm> for details.

You should get the following interpretation for the `0x5 0x25 0x0` field:

```
sense=5 (ILLEGAL REQUEST), ASC=25 ASCQ=0 (LOGICAL UNIT NOT SUPPORTED)
```

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 adapter problems and to find appropriate solutions.

Table 6-3. `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.

Checking Metadata Consistency with VOMA

Use VMware Ondisk Metadata Analyser (VOMA) when you experience problems with your VMFS datastore and need to check metadata consistency of VMFS or logical volume backing the VMFS volume.

Problem

The following examples show circumstances in which you might need to perform a metadata check:

- You experience SAN outages.
- After you rebuild RAID or perform a disk replacement.
- You see metadata errors in the `vmkernel.log` file.
- You are unable to access files on the VMFS datastore that are not in use by any other host.

Solution

To check metadata consistency, run VOMA from the CLI of an ESXi host version 5.1 or later. VOMA can check both the logical volume and the VMFS for metadata inconsistencies. You can use VOMA on VMFS3 and VMFS5 datastores. VOMA runs in a read-only mode and serves only to identify problems. VOMA does not fix errors that it detects. Consult VMware Support to resolve errors reported by VOMA.

Follow these guidelines when you use the VOMA tool:

- Make sure that the VMFS datastore you analyze does not span multiple extents. You can run VOMA only against a single-extent datastore.
- Power off any virtual machines that are running or migrate them to a different datastore.

Follow these steps when you use the VOMA tool to check VMFS metadata consistency.

- 1 Obtain the name and partition number of the device that backs the VMFS datastore that you need to check.

#esxcli storage vmfs extent list

The Device Name and Partition columns in the output identify the device. For example:

```
Volume Name  XXXXXXXX  Device Name                               Partition
1TB_VMFS5   XXXXXXXX  naa.600508e000000000b367477b3be3d703    3
```

- 2 Run VOMA to check for VMFS errors.

Provide the absolute path to the device partition that backs the VMFS datastore, and provide a partition number with the device name. For example:

```
# voma -m vmfs -f check -d /vmfs/devices/disks/naa.600508e000000000b367477b3be3d703:3
```

The output lists possible errors. For example, the following output indicates that the heartbeat address is invalid.

```
XXXXXXXXXXXXXXXXXXXXXXXXX
Phase 2: Checking VMFS heartbeat region
ON-DISK ERROR: Invalid HB address
Phase 3: Checking all file descriptors.
Phase 4: Checking pathname and connectivity.
Phase 5: Checking resource reference counts.
```

```
Total Errors Found:          1
```

The VOMA tool uses the following options.

Table 6-4. VOMA Command Options

Command Option	Description
-m --module	The module to run: <code>vmfs</code> or <code>lvm</code> . If you specify <code>vmfs</code> , minimal checks are performed for LVM as well. The default module is <code>vmfs</code> .
-f --func	Functions to be performed: <code>query</code> - list functions supported by module. <code>check</code> - 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, <code>/vmfs/devices/disks/naa.00000000000000000000000000000000:1</code> .
-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.

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:

- [“Duplicate MAC Addresses of Virtual Machines on the Same Network,”](#) on page 71
- [“The Conversion to the Enhanced LACP Support Fails,”](#) on page 74
- [“Unable to Remove a Host from a vSphere Distributed Switch,”](#) on page 75
- [“Hosts on a vSphere Distributed Switch 5.1 and Later Lose Connectivity to vCenter Server,”](#) on page 76
- [“Hosts on vSphere Distributed Switch 5.0 and Earlier Lose Connectivity to vCenter Server,”](#) on page 77
- [“Alarm for Loss of Network Redundancy on a Host,”](#) on page 78
- [“Virtual Machines Lose Connectivity After Changing the Uplink Failover Order of a Distributed Port Group,”](#) on page 79
- [“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 80
- [“Low Throughput for UDP Workloads on Windows Virtual Machines,”](#) on page 82
- [“Virtual Machines on the Same Distributed Port Group and on Different Hosts Cannot Communicate with Each Other,”](#) on page 83

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 from one vCenter Server instance to another in the same network, 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 in 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 **All vCenter Actions > 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.

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 might fail.

Problem

After you upgrade a vSphere distributed switch to version 5.5, 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 7-1. 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 datacenter 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 Select Manage > 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.

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.

- 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 All vCenter Actions > 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 All vCenter Actions > 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.

- 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.

- 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. Navigate to the host, and under **Manage**, click **Networking**. From **Virtual switches** select the distributed switch and click **Rectify**.
 - 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

- Connect to an affected host by using the vSphere Client.
- 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**, and click **Next**.
 - d In the Port Group Properties section, type a network label that identifies the port group that you are creating and optionally a VLAN ID.
 - e Click **Finish**.
 - f In the vSphere Standard Switch view, click **Properties** for the created switch.
 - g Click **Network Adapters**, click **Add**, and select an unoccupied physical adapter to carry the management traffic.

If all physical adapters are already busy with traffic from other switches, remove the physical adapter for the management network from the proxy switch of the distributed switch, and add it to this standard switch.
 - h Click the **Ports** tab and provide a valid configuration of the VMkernel adapter port group.
 - i Click **Close**.
- 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 Type a network label and optionally a VLAN ID for the management network, and click **Next**.
- 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 Select the **Manage host networking** option and select the host.
 - d To assign the NIC of the host to an active uplink, select the **Manage physical adapters** option and associate the NIC to the switch uplink in the Manage physical adapters page.
- Move the NIC at the level of the host.
 - a Navigate to the host in the vSphere Web Client, and under **Manage**, click **Networking**.
 - b Select **Virtual Switches** and select the distributed proxy switch.
 - c Click **Manage the physical adapters**, and move the NIC to the active uplink

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	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. 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.
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	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. 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.
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 **Manage** 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> <code>16000</code>
Disable coalescing for low throughput or latency-sensitive workloads.	<code>ethernetX.coalescingScheme</code>	<code>disabled</code>
Revert to the coalescing algorithm from earlier ESXi releases. <small>NOTE The ability to revert to the earlier algorithm will not be available in releases later than vSphere 5.5.</small>	<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 <code>esxcli system settings advanced set</code> Command	Value
Set a default interrupt rate higher than the expected packet rate.	<code>Net.CoalesceRBCRate</code>	<code>/Net/CoalesceRBCRate</code>	For example, set it to 16000 if 15000 interrupts are expected per second.
Disable coalescing for low throughput or latency-sensitive workloads.	<code>Net.CoalesceDefaultOn</code>	<code>/Net/CoalesceDefaultOn</code>	Set it to 0.
Revert to the coalescing scheme from earlier ESXi releases. NOTE The ability to revert to the earlier algorithm will not be available in releases later than vSphere 5.5.	<code>Net.CoalesceVersion</code>	<code>/Net/CoalesceVersion</code>	Set it to 1.

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.

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. The troubleshooting information also provides solutions to problems that you might have accessing and using the licensing reporting function.

This chapter includes the following topics:

- [“Troubleshooting Host Licensing,”](#) on page 85
- [“Troubleshooting License Reporting,”](#) on page 86
- [“Unable to Power On a Virtual Machine,”](#) on page 89
- [“Unable to Assign a License Key to vCenter Server,”](#) on page 89
- [“Unable to Configure or Use a Feature,”](#) on page 90

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 Key to an ESXi Host

Under certain conditions, you might not be able to assign a license key to an ESXi host asset.

Problem

You try to assign a license key 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 key 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 Essentials license key with capacity for two processors. You try to assign the key to a host that has four processors. You cannot assign the key, because the required license usage for the host is four processors.
- The features on the host do not match the license edition. For example, you might configure vMotion and DRS on a cluster of hosts while you are using evaluation mode. Later, you try to assign Standard license keys to the hosts. This operation fails because the Standard edition does not include vMotion and DRS.
- You do not apply the correct license key.

- The host is connected to a vCenter Server system that is assigned a license key that restricts the edition of the license that you want to assign. For example, vCenter Server is licensed with vCenter Server Standard, and the license key is for vSphere Essentials.

Solution

- Assign a license key with larger capacity.
- Upgrade the license edition to match the resources and features on the host, or disable the features and resources that do not match the license edition.
- Assign a correct license key. To license ESXi hosts, you must assign a vSphere license key.
- Assign a license key whose edition is compatible with the license edition of vCenter Server. For example, if vCenter Server is licensed with vCenter Server Standard, you need a vSphere Standard license key.

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 or all ESXi hosts disconnect from vCenter Server and you receive a licensing-related error message.
- 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 is expired or the host license is expired.
- The 60-day evaluation period of vCenter Server is expired or the vCenter Server license is expired.

Solution

- Obtain a vSphere license key and assign it to the ESXi host.
- Obtain a vCenter Server license key and assign it to vCenter Server. If the vCenter Server system is managing ESXi 3.5, it must have access to a license server. You can download the VMware License Server from the VMware Web site.

NOTE When you assign a license key to an ESXi host and vCenter Server, the license edition must be compatible with all of the features you configured. If the license edition and the configured features are incompatible, you cannot assign the license key.

Troubleshooting License Reporting

The troubleshooting license reporting topics provide solutions to problems that you might have in accessing the license reporting function, viewing the license usage for products, or exporting a license report.

License Reporting Interface Does Not Appear in the vSphere Web Client

The license reporting function is not available in the vSphere Web Client.

Problem

In the vSphere Web Client, when you navigate to **Administration > Licensing > License Reports**, the license reporting interface does not load and an error message appears.

Cause

The license reporting interface might not load because of the following reasons.

- VMware VirtualCenter Management Webservices service is not running on the selected vCenter Server instance.
- The vCenter Inventory Service is not running on the selected vCenter Server instance.
- A license service is not running on the selected vCenter Server instance.

Solution**Table 8-1.** Enabling the License Reporting

Cause	Solution
VMware VirtualCenter Management Webservices service is not running on vCenter Server.	Verify that the VMware VirtualCenter Management Webservices service is running on the vCenter Server system. Navigate to the vCenter Server system in the inventory and select Monitor > Service Health . The page displays the following message if the VMware VirtualCenter Management Webservices service is not started. <code>Could not get vCenter Health status</code> If you see this error message, start the VMware VirtualCenter Management Webservices service.
vCenter Inventory Service is not running on vCenter Server.	Verify that the vCenter Inventory Service is running on the vCenter Server system. Navigate to the vCenter Server system in the inventory and select Monitor > Service Health . Start the vCenter Inventory Service if it is not running.
A license service is not running on vCenter Server.	Verify that all license services are running on the vCenter Server system. Navigate to the vCenter Server system in the inventory and select Monitor > Service Health . If a license service is not available, perform the actions recommended in the error message that appears. Restarting the VMware VirtualCenter Management Webservices service might also help.

Unable to View License Use Data in the vSphere Web Client

You might be unable to view license use data in the vSphere Web Client.

Problem

- You navigate to **Administration > Licensing > License Reports**. When you try to view the license use for products, one of the following error messages appears:

The licensing service for vCenter Server is unavailable.

Licensing usage data for <vCenter Server instance> is missing for the selected time period.

- When you try to view details for a license key, the following error message appears:

Licensing usage data for <license key> is missing for the selected time period.

Cause

- VMware VirtualCenter Management Webservices service is not running on the selected vCenter Server instance.
- The vCenter Inventory Service is not running on the selected vCenter Server instance.

- A licensing service is not available for the selected vCenter Server instance.
- No license keys are assigned to assets for the selected vCenter Server instance and time period.

Solution

Table 8-2. Enabling the License Reporting

Cause	Solution
VMware VirtualCenter Management Webservices service is not running on vCenter Server.	Verify that the VMware VirtualCenter Management Webservices service is running on the selected vCenter Server instance. Navigate to the vCenter Server instance and select Monitor > Service Health . The page displays the following message if the VMware VirtualCenter Management Webservices service is not started. <code>Could not get vCenter Health status</code> If you see this error message, start the VMware VirtualCenter Management Webservices service.
vCenter Inventory Service is not running on vCenter Server.	Verify that vCenter Inventory Service is running on the selected vCenter Server instance. Navigate to the vCenter Server instance and select Monitor > Service Health . Start the vCenter Inventory Service if it is not running.
A license service is not running on vCenter Server.	Verify that all license services are running on the selected vCenter Server instance. Navigate to the vCenter Server instance and select Monitor > Service Health . If a license service is not available, perform the actions recommended in the error message that appears. Restarting the VMware VirtualCenter Management Webservices service might also help.
No license usage data is available for the selected vCenter Server instance and time period.	Select a time period and a vCenter Server instance for which licenses keys are assigned to assets.

Unable to Export a Licensing Report in the vSphere Web Client

You cannot export a licensing report from the vSphere Web Client.

Problem

- When you click **Export** for a licensing report in the **License Reports** option, the following error message appears:

```
Cannot export licensing usage data.
An integrity problem with the license data detected in the database of vCenter Server.
```

- When you try to export the license usage for all vCenter Server instances in the Linked Mode group, the following error message appears:

```
vCenter Server instances included in the generated export file:
vCenter Server instance 1
vCenter Server instance 2
...
Unable to export license usage data for:
vCenter Server instance 1
vCenter Server instance 2
...
```

Cause

- The license usage data stored in the vCenter Server database has been modified. Do not modify licensing records in vCenter Server database.

- The vCenter Server instances for which you cannot export license usage are not running.
- The vCenter Server instances for which you cannot export license usage are isolated from the Linked Mode group.

Solution

- If the license usage data in the vCenter Server database has been modified, no solution is available. You cannot export licensing reports for this vCenter Server or Linked Mode group within this time period.
- If the vCenter Server instances for which you cannot export a licensing report are not running, restart them if possible and try to export a licensing report again.
- If the vCenter Server instances for which you cannot export a licensing report are isolated, troubleshoot the cause and try to export the licensing report again. The isolated vCenter Server instances might not run or the connection with them might be down.

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.
- The edition of the license key does not match the configured features and resources on the host.

Solution

Table 8-3. Power on a Virtual Machine

Cause	Solution
The evaluation period of the host is expired.	Obtain a vSphere license key and assign the key to the ESXi host. Verify that the edition of the license key matches the configured features and resources on the host. If they do not match, you cannot assign the license key.
The license of the host is expired.	
The edition of the license key does not match the configured features and resources on the host.	

Unable to Assign a License Key to vCenter Server

You cannot assign a license key to a vCenter Server system.

Problem

You try to assign a license key to a vCenter Server system, but the operation is unsuccessful and you receive an error message.

Cause

You might be unable to assign a license key to a vCenter Server system because of the following reasons.

- The license edition does not match the currently configured resources and features on vCenter Server. For example, while in evaluation mode, you add the vCenter Server system to a Linked Mode group. Then you try to assign a Foundation or an Essentials license key to the vCenter Server system. The operation is unsuccessful because the Foundation and Essentials license editions do not support the Linked Mode feature.
- You assign an incorrect license key.

Solution

- Upgrade the license edition to match the currently configured features and resources on vCenter Server. For example, you need a Standard or an Enterprise license key to be able to license a vCenter Server system that is in Linked Mode.
- Assign a correct license key. To license vCenter Server, you need a vCenter Server license key.

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

If you downgrade your license from evaluation mode to a license that does not support the features that you configured while you used evaluation mode, you receive a warning message about license downgrade.

Solution

Check the licensed features on the 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

- Active Directory **25**
- address already in use, Jetty server **30**
- Advanced Runtime Info **39**
- advanced settings,
 - Disk.SchedNumReqOutstanding **66**
- alarm definitions **31**
- authentication **62**
- Auto Deploy
 - coredump **20**
 - DHCP address **22**
 - failing to complete boot **21**
 - failure to boot **22**
 - image profile warning **20**
 - network boot problem **23**
 - redirection problem **19**
 - TFTP server **21**
 - timeout error **19**
 - troubleshooting **19**
 - wrong image **19**
- Auto Deploy upgrade **24**

C

- cannot add host to domain **26**
- cannot apply Storage DRS recommendations,
 - troubleshooting **58**
- cannot enable Storage I/O Control **60**
- certificates
 - hosts **34**
 - vCenter Server **34**
- CHAP authentication **62**
- cluster issues
 - cluster load high **47**
 - DRS does not vMotion **48**
 - host not powered off **47**
 - hosts power off **47**
 - load imbalance **45**
 - low cluster load **47**
 - red cluster because failover capacity violated **46**
 - red cluster because inconsistent resource pool **46**
 - yellow cluster **46**
- cluster problems **45**
- compliance failure, host profiles **25**

- config.vpxd.das.electionWaitTimeSec **42**
- converting to the enhanced LACP fails **74**

D

- datastore clusters, maintenance mode **55**
- datastore heartbeating **39**
- datastores
 - maintenance mode **55**
 - performance charts troubleshooting **59**
- Denial of Service
 - virtual machine, VPN **80**
 - See also DoS
- Disk.SchedNumReqOutstanding **66**
- distributed switch, cannot remove host **75**
- distributed port group, virtual machines cannot communicate **83**
- DNS **31**
- domain
 - add host **26**
 - join **26**
- domain controller **31**
- duplicate session error, vSphere Web Client **31**

E

- Export license usage **88**
- extensions, troubleshooting **35**

F

- failed conversion to the enhanced LACP **74**
- failover, lost virtual machine connectivity **79**
- Fault Tolerance
 - logging **10**
 - troubleshooting **7–10**
- Feature **89, 90**
- firewall
 - network-based **33**
 - Windows **33**
- flash player, insufficient memory **31**
- ft.maxSwitchoverSeconds **9**

G

- gpupdate /force command **31**
- group policy update **31**
- GUID **31**

H

- Hardware Virtualization (HV) **7, 8**
- HBAs, queue depth **65**
- host certificates **34**
- Host Failures Cluster Tolerates admission control policy **37–39**
- host isolation response **18**
- host issues
 - high cluster load **49**
 - host not powered off **49**
 - host not powered on **49**
 - low cluster load **48, 49**
 - maintenance mode **50**
 - standby mode **50**
 - virtual machines not moved by DRS **50**
- host problems **48**
- host profiles, compliance failure **25**
- hosts, no connection with vCenter Server **76, 77**

I

- iBFT **67**
- introduction **45**
- IPv4 **18**
- IPv6 **18**

J

- Jetty server, address already in use **30**

L

- license capacity **89**
- License capacity **85**
- License edition **85, 89, 90**
- license key **89**
- License reporting **87**
- License usage **87**
- Licensing reporting **86**
- Linked Mode
 - reachability **31**
 - troubleshooting **31–33**
- login, vSphere Web Client **31**
- lookup service error, vCenter Server Appliance **28**
- loss of uplink redundancy **78**
- lost virtual machine connectivity, failover **79**
- low throughput, Windows virtual machine **82**
- LUN not visible, SP visibility **61, 62**
- LUN queue depth **64**
- LUN thrashing **64**

M

- maintenance mode, datastores **55**
- maximum HBA queue depth **65**
- metadata consistency, checking with VOMA **69**

- monitoring, Storage I/O Control **59**

N

- network partition **39, 40, 42**
- NFS datastores **67**
- no uplink redundancy **78**
- non-ASCII characters **67**

O

- orphaned virtual machines, recovering **12**
- outstanding disk requests **66**
- overcommitted host **8**

P

- password requirements **24**
- path thrashing **64**
- Percentage of Cluster Resources Reserved admission control policy **38**
- performance, problems **63**
- performance charts for datastores, troubleshooting viewing performance charts **59**
- plug-ins, troubleshooting **35**
- Primary VM **42, 43**
- problems
 - performance **63**
 - visibility **61, 62**

Q

- queue depth **65**

R

- recommendations for Storage DRS, troubleshooting **58**
- Red Cluster **37**
- registry settings **33**
- remove datastore **40**
- RPCCfg.exe **33**

S

- SCSI INQUIRY **66**
- SCSI reservations, reducing **63**
- SCSI Sense codes **68**
- SDK **31**
- Secondary VM **42, 43**
- slot size **37–39**
- snapshots
 - consolidating **13**
 - troubleshooting **13**
- software iSCSI adapters, queue depth **65**
- SP visibility, LUN not visible **61, 62**
- SQL compatibility mode **28**
- SSL certificates, troubleshooting in vSphere HA **35**

storage adapters, troubleshooting **68**
 storage devices, display problems **61**
 Storage DRS
 affinity rules **57**
 cannot apply recommendations **58**
 deleting affinity rules **57**
 disabled **54**
 disabling **57**
 faults **57**
 OVF templates **57**
 placement **58**
 recommendations **58**
 rule violation **57**
 troubleshooting **54**
 Storage I/O Control
 monitoring **59**
 troubleshooting **59, 60**

T

TFTP server, Auto Deploy **21**
 timeout error, Auto Deploy **19**
 Tomcat service, vCenter Server upgrade
 failure **28**
 TRACE logging **27**
 troubleshooting
 certificates **34**
 extensions **35**
 Linked Mode **31, 32**
 plug-ins **35**
 USB devices **11**
 vCenter Server **27**
 vCenter Server Appliance **28**
 vSphere Web Client **27, 29**
 troubleshooting licensing **85**
 troubleshooting Fault Tolerance **7**
 Troubleshooting host licensing **85**
 Troubleshooting Hosts **15**
 troubleshooting license reporting **86**
 troubleshooting storage **61**
 Troubleshooting Virtual Machines **7**

U

unmount datastore **40**
 uplink redundancy lost **78**
 URLs, configuring **32**
 USB devices, error messages **11**
 USB passthrough
 restarting the USB arbitrator **12**
 troubleshooting **11**
 troubleshooting device connections **12**

V

vCenter Server
 configuring URLs **32**

 no connection with hosts **76, 77**
 troubleshooting **27**
 troubleshooting certificates **34**
 vCenter Server Appliance, lookup service
 error **28**
 vCenter Server certificates **34**
 vCenter Server license **86, 89**
 vCenter Server system does not appear **29**
 vCenter Server upgrade **24**
 vCenter Server upgrade fails, Tomcat service **28**
 vCenterServer.VimApiUrl **32**
 vCenterServer.VimWebServicesUrl **32**
 VDS, cannot remove host **75**
 virtual machine
 BPDU **80**
 bridge **80**
 Denial of Service **80**
 duplicate MAC addresses **71**
 low throughput **82**
 MAC address conflict **71**
 UDP **82**
 VPN **80**
 Windows **82**
 virtual machine console does not launch,
 vSphere Web Client **30**
 virtual machine issues
 affinity rules **52**
 anti-affinity rules **52**
 CPU **51**
 memory resources **51**
 power on failure **53**
 virtual machine not moved by DRS **53**
 virtual machine problems **51**
 Virtual Machine Protection State **40**
 virtual machines **13**
 virtual machines, orphaned **12**
 Virtual SAN **18, 39, 41**
 visibility problems **61, 62**
 VM-Host affinity rules **41**
 VMFS, checking metadata consistency **69**
 VMware Inventory Service **27**
 VMware Ondisk Metadata Analyser, See VOMA
 VMware vCenter Management Webservices **29**
 vmware-fdm **24**
 VOMA **69**
 vSphere distributed switch
 hosts not responding **76, 77**
 lost virtual machine connectivity **79**
 vSphere DRS **41**
 vSphere Fault Tolerance **42, 43**

- vSphere HA
 - Denial of Service **80**
 - troubleshooting SSL certificates **35**
- vSphere HA admission control **37**
- vSphere HA Admission Control **37**
- vSphere HA agent **15–18**
- vSphere HA cluster **42**
- vSphere HA failovers **40**
- vSphere HA host state
 - Agent Uninitialized **16**
 - Agent Unreachable **15**
 - Host Failed **17**
 - Initialization Error **16**
 - Network Isolated **18**
 - Network Partitioned **18**
 - Uninitialization Error **17**
- vSphere HA restart failures **41**
- vSphere license **86, 89**
- vSphere Web Client
 - duplicate session error **31**
 - failure to log in after upgrade **31**
 - login **31**
 - troubleshooting **27, 29**
 - virtual machine console does not launch **30**
- VWS **31**

W

- webclient.properties file **30**