

vSphere Availability Guide

Update 1

ESX 4.0

ESXi 4.0

vCenter Server 4.0

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-000265-04

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 © 2009–2011 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/or 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

Updated Information	5
About This Book	7
1 Business Continuity and Minimizing Downtime	9
Reducing Planned Downtime	9
Preventing Unplanned Downtime	10
VMware HA Provides Rapid Recovery from Outages	10
VMware Fault Tolerance Provides Continuous Availability	11
2 Creating and Using VMware HA Clusters	13
How VMware HA Works	13
VMware HA Admission Control	15
Creating a VMware HA Cluster	21
Customizing VMware HA Behavior	25
Best Practices for VMware HA Clusters	27
3 Providing Fault Tolerance for Virtual Machines	31
How Fault Tolerance Works	31
Fault Tolerance Use Cases	32
Fault Tolerance Configuration Requirements	33
Fault Tolerance Interoperability	34
Preparing Your Cluster and Hosts for Fault Tolerance	35
Turning On Fault Tolerance for Virtual Machines	37
Viewing Information About Fault Tolerant Virtual Machines	39
Fault Tolerance Best Practices	40
VMware Fault Tolerance Configuration Recommendations	41
Troubleshooting Fault Tolerance	42
Appendix: Fault Tolerance Error Messages	45
Index	51

Updated Information

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

This table provides the update history of the *vSphere Availability Guide*.

Revision	Description
EN-000265-04	Edited note in “Creating a VMware HA Cluster,” on page 21 to indicate that automatic startup is not supported when used with VMware HA.
EN-000265-03	The section “Failure Detection and Host Network Isolation,” on page 14 is updated to change the sentence “By default, the isolated host leaves its virtual machines powered on, but you can change the host isolation response to Shut Down VM or Power Off VM.” to “By default, the isolated host shuts down its virtual machines, but you can change the host isolation response to Leave powered on or Power off.”
EN-000265-02	Added new information on VMXNET3 driver and Paravirtualized SCSI (PVSCSI) adapter in Table 3-1 in the topic “Other Features Incompatible with Fault Tolerance,” on page 34.
EN-000265-01	Minor updates.
EN-000265-00	Initial release.

About This Book

The *vSphere Availability Guide* describes solutions that provide business continuity, including how to establish VMware[®] High Availability (HA) and VMware Fault Tolerance.

Intended Audience

This book is for anyone who wants to provide business continuity through the VMware HA and Fault Tolerance solutions. The information in this book is for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacenter operations.

Document Feedback

VMware welcomes your suggestions for improving our documentation. If you have comments, send your feedback to docfeedback@vmware.com.

vSphere Documentation

The vSphere[®] documentation consists of the combined VMware vCenter Server and ESX/ESXi documentation set. The *vSphere Availability Guide* covers ESX[®], ESXi, and vCenter[®] Server.

Technical Support and Education Resources

The following technical support resources are available to you. To access the current version of this book and other books, go to <http://www.vmware.com/support/pubs>.

Online and Telephone Support

To use online support to submit technical support requests, view your product and contract information, and register your products, go to <http://www.vmware.com/support>.

Customers with appropriate support contracts should use telephone support for the fastest response on priority 1 issues. Go to http://www.vmware.com/support/phone_support.html.

Support Offerings

To find out how VMware support offerings can help meet your business needs, go to <http://www.vmware.com/support/services>.

VMware Professional Services

VMware Education Services courses offer extensive hands-on labs, case study examples, and course materials designed to be used as on-the-job reference tools. Courses are available onsite, in the classroom, and live online. For onsite pilot programs and implementation best practices, VMware Consulting

Services provides offerings to help you assess, plan, build, and manage your virtual environment. To access information about education classes, certification programs, and consulting services, go to <http://www.vmware.com/services>.

Business Continuity and Minimizing Downtime

1

Downtime, whether planned or unplanned, brings with it considerable costs. However, solutions to ensure higher levels of availability have traditionally been costly, hard to implement, and difficult to manage.

VMware software makes it simpler and less expensive to provide higher levels of availability for important applications. With vSphere, organizations can easily increase the baseline level of availability provided for all applications as well as provide higher levels of availability more easily and cost effectively. With vSphere, you can:

- Provide higher availability independent of hardware, operating system, and applications.
- Eliminate planned downtime for common maintenance operations.
- Provide automatic restart in cases of failure.

vSphere makes it possible to reduce planned downtime, prevent unplanned downtime, and recover rapidly from outages.

This chapter includes the following topics:

- [“Reducing Planned Downtime,”](#) on page 9
- [“Preventing Unplanned Downtime,”](#) on page 10
- [“VMware HA Provides Rapid Recovery from Outages,”](#) on page 10
- [“VMware Fault Tolerance Provides Continuous Availability,”](#) on page 11

Reducing Planned Downtime

Planned downtime typically accounts for over 80% of datacenter downtime. Hardware maintenance, server migration, and firmware updates all require downtime for physical servers. To minimize the impact of this downtime, organizations are forced to delay maintenance until inconvenient and difficult-to-schedule downtime windows.

vSphere makes it possible for organizations to dramatically reduce planned downtime. Because workloads in a vSphere environment can be dynamically moved to different physical servers without downtime or service interruption, server maintenance can be performed without requiring application and service downtime. With vSphere organizations can:

- Eliminate downtime for common maintenance operations.
- Eliminate planned maintenance windows.
- Perform maintenance at any time without disrupting users and services.

The VMotion[®] and Storage VMotion functionality in vSphere makes it possible for organizations to dramatically reduce planned downtime because workloads in a VMware environment can be dynamically moved to different physical servers or to different underlying storage without service interruption. Administrators can perform faster and completely transparent maintenance operations, without being forced to schedule inconvenient maintenance windows.

Preventing Unplanned Downtime

While an ESX/ESXi host provides a robust platform for running applications, an organization must also protect itself from unplanned downtime caused from hardware or application failures. vSphere builds important capabilities into datacenter infrastructure that can help you prevent unplanned downtime.

These vSphere capabilities are part of virtual infrastructure and are transparent to the operating system and applications running in virtual machines. These features can be configured and utilized by all the virtual machines on a physical system, reducing the cost and complexity of providing higher availability. Key fault-tolerance capabilities are built into vSphere:

- Shared storage. Eliminate single points of failure by storing virtual machine files on shared storage, such as Fibre Channel or iSCSI SAN, or NAS. The use of SAN mirroring and replication features can be used to keep updated copies of virtual disk at disaster recovery sites.
- Network interface teaming. Provide tolerance of individual network card failures.
- Storage multipathing. Tolerate storage path failures.

In addition to these capabilities, the VMware HA and Fault Tolerance features can minimize or eliminate unplanned downtime by providing rapid recovery from outages and continuous availability, respectively.

VMware HA Provides Rapid Recovery from Outages

VMware HA leverages multiple ESX/ESXi hosts configured as a cluster to provide rapid recovery from outages and cost-effective high availability for applications running in virtual machines.

VMware HA protects application availability in two ways:

- It protects against a server failure by automatically restarting the virtual machines on other hosts within the cluster.
- It protects against application failure by continuously monitoring a virtual machine and resetting it in the event that a failure is detected.

Unlike other clustering solutions, VMware HA provides the infrastructure to protect all workloads with the infrastructure:

- No special software needs to be installed within the application or virtual machine. All workloads are protected by VMware HA. After VMware HA is configured, no actions are required to protect new virtual machines. They are automatically protected.
- VMware HA can be combined with VMware Distributed Resource Scheduler (DRS) not only to protect against failures but also to provide load balancing across the hosts within a cluster.

VMware HA has a number of advantages over traditional failover solutions:

Minimal setup	After a VMware HA cluster is set up, all virtual machines in the cluster get failover support without additional configuration.
Reduced hardware cost and setup	The virtual machine acts as a portable container for the applications and it can be moved among hosts. Administrators avoid duplicate configurations on multiple machines. When you use VMware HA, you must have sufficient resources to fail over the number of hosts you want to protect with VMware HA. However, the vCenter Server system automatically manages resources and configures clusters.
Increased application availability	Any application running inside a virtual machine has access to increased availability. Because the virtual machine can recover from hardware failure, all applications that start at boot have increased availability without increased computing needs, even if the application is not itself a clustered application. By monitoring and responding to VMware Tools heartbeats and resetting non-responsive virtual machines, it also protects against guest operating system crashes.
DRS and VMotion integration	If a host fails and virtual machines are restarted on other hosts, DRS can provide migration recommendations or migrate virtual machines for balanced resource allocation. If one or both of the source and destination hosts of a migration fail, VMware HA can help recover from that failure.

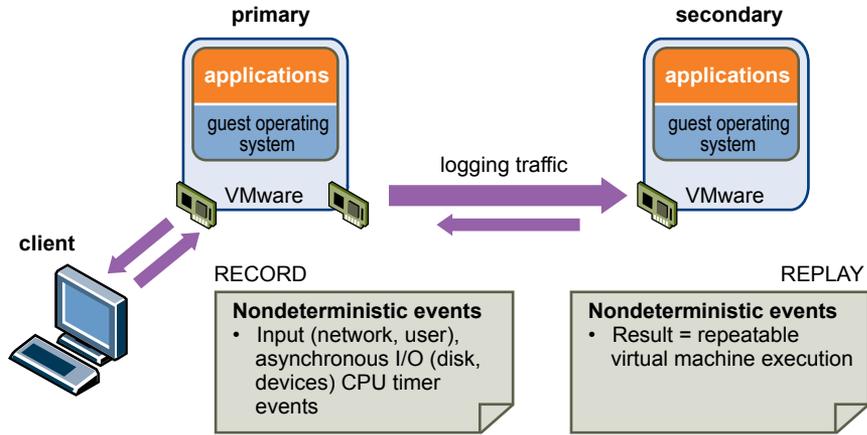
VMware Fault Tolerance Provides Continuous Availability

VMware HA provides a base level of protection for your virtual machines by restarting virtual machines in the event of a host failure. VMware Fault Tolerance provides a higher level of availability, allowing users to protect any virtual machine from a host failure with no loss of data, transactions, or connections.

Fault Tolerance uses the VMware vLockstep technology on the ESX/ESXi host platform to provide continuous availability. This is done by ensuring that the states of the Primary and Secondary VMs are identical at any point in the instruction execution of the virtual machine. vLockstep accomplishes this by having the Primary and Secondary VMs execute identical sequences of x86 instructions. The Primary VM captures all inputs and events -- from the processor to virtual I/O devices -- and replays them on the Secondary VM. The Secondary VM executes the same series of instructions as the Primary VM, while only a single virtual machine image (the Primary VM) is seen executing the workload.

If either the host running the Primary VM or the host running the Secondary VM fails, a transparent failover occurs whereby the host that is still functioning seamlessly becomes the host of the Primary VM. With transparent failover, there is no data loss and network connections are maintained. After a transparent failover occurs, a new Secondary VM is automatically respawned and redundancy is re-established. The entire process is transparent and fully automated and occurs even if vCenter Server is unavailable.

Figure 1-1. Primary VM and Secondary VM in Fault Tolerance Pair



Creating and Using VMware HA Clusters

2

VMware HA clusters enable a collection of ESX/ESXi hosts to work together so that, as a group, they provide higher levels of availability for virtual machines than each ESX/ESXi host could provide individually. When you plan the creation and usage of a new VMware HA cluster, the options you select affect the way that cluster responds to failures of hosts or virtual machines.

Before creating a VMware HA cluster, you should be aware of how VMware HA identifies host failures and isolation and responds to these situations. You also should know how admission control works so that you can choose the policy that best fits your failover needs. After a cluster has been established, you can customize its behavior with advanced attributes and optimize its performance by following recommended best practices.

This chapter includes the following topics:

- [“How VMware HA Works,”](#) on page 13
- [“VMware HA Admission Control,”](#) on page 15
- [“Creating a VMware HA Cluster,”](#) on page 21
- [“Customizing VMware HA Behavior,”](#) on page 25
- [“Best Practices for VMware HA Clusters,”](#) on page 27

How VMware HA Works

VMware HA provides high availability for virtual machines by pooling them and the hosts 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.

Primary and Secondary Hosts in a VMware HA Cluster

When you add a host to a VMware HA cluster, an agent is uploaded to the host and configured to communicate with other agents in the cluster. The first five hosts added to the cluster are designated as primary hosts, and all subsequent hosts are designated as secondary hosts. The primary hosts maintain and replicate all cluster state and are used to initiate failover actions. If a primary host is removed from the cluster, VMware HA promotes another host to primary status.

Any host that joins the cluster must communicate with an existing primary host to complete its configuration (except when you are adding the first host to the cluster). At least one primary host must be functional for VMware HA to operate correctly. If all primary hosts are unavailable (not responding), no hosts can be successfully configured for VMware HA.

One of the primary hosts is also designated as the active primary host and its responsibilities include:

- Deciding where to restart virtual machines.
- Keeping track of failed restart attempts.
- Determining when it is appropriate to keep trying to restart a virtual machine.

If the active primary host fails, another primary host replaces it.

Failure Detection and Host Network Isolation

Agents communicate with each other and monitor the liveness of the hosts in the cluster. This is done through the exchange of heartbeats, by default, every second. If a 15-second period elapses without the receipt of heartbeats from a host, and the host cannot be pinged, it is declared as failed. In the event of a host failure, the virtual machines running on that host are failed over, that is, restarted on the alternate hosts with the most available unreserved capacity (CPU and memory.)

NOTE In the event of a host failure, VMware HA does not fail over any virtual machines to a host that is in maintenance mode, because such a host is not considered when VMware HA computes the current failover level. When a host exits maintenance mode, the VMware HA service is reenabled on that host, so it becomes available for failover again.

Host network isolation occurs when a host is still running, but it can no longer communicate with other hosts in the cluster. With default settings, if a host stops receiving heartbeats from all other hosts in the cluster for more than 12 seconds, it attempts to ping its isolation addresses. If this also fails, the host declares itself as isolated from the network.

When the isolated host's network connection is not restored for 15 seconds or longer, the other hosts in the cluster treat it as failed and attempt to fail over its virtual machines. However, when an isolated host retains access to the shared storage it also retains the disk lock on virtual machine files. To avoid potential data corruption, VMFS disk locking prevents simultaneous write operations to the virtual machine disk files and attempts to fail over the isolated host's virtual machines fail. By default, the isolated host shuts down its virtual machines, but you can change the host isolation response to **Leave powered on** or **Power off**. See [“Virtual Machine Options,”](#) on page 23.

NOTE If you ensure that your network infrastructure is sufficiently redundant and that at least one network path is available at all times, host network isolation should be a rare occurrence.

Using VMware HA and DRS Together

Using VMware HA in conjunction with Distributed Resource Scheduler (DRS) combines automatic failover with load balancing. This combination can result in faster rebalancing of virtual machines after VMware HA has moved virtual machines to different hosts.

When VMware HA performs failover and restarts virtual machines on different hosts, its first priority is the immediate availability of all virtual machines. After the virtual machines have been restarted, those hosts on which they were powered on might be heavily loaded, while other hosts are comparatively lightly loaded. VMware HA uses the CPU and memory reservation to determine failover, while the actual usage might be higher.

In a cluster using DRS and VMware HA with admission control turned on, virtual machines might not be evacuated from hosts entering maintenance mode. This is because of the resources reserved to maintain the failover level. You must manually migrate the virtual machines off of the hosts using VMotion.

When VMware HA admission control is disabled, failover resource constraints are not passed on to DRS and VMware Distributed Power Management (DPM). The constraints are not enforced.

- DRS does evacuate virtual machines from hosts and place the hosts in maintenance mode or standby mode regardless of the impact this might have on failover requirements.
- VMware DPM does power off hosts (place them in standby mode) even if doing so violates failover requirements.

For more information about DRS, see *Resource Management Guide*.

VMware HA Admission Control

vCenter Server uses admission control to ensure that sufficient resources are available in a cluster to provide failover protection and to ensure that virtual machine resource reservations are respected.

Including VMware HA, there are three types of admission control.

Host	Ensures that a host has sufficient resources to satisfy the reservations of all virtual machines running on it.
Resource pool	Ensures that a resource pool has sufficient resources to satisfy the reservations, shares, and limits of all virtual machines associated with it.
VMware HA	Ensures that sufficient resources in the cluster are reserved for virtual machine recovery in the event of host failure.

Admission control imposes constraints on resource usage and any action that would violate these constraints is not permitted. Examples of actions that could be disallowed include:

- Powering on a virtual machine.
- Migrating a virtual machine onto a host or into a cluster or resource pool.
- Increasing the CPU or memory reservation of a virtual machine.

Of the three types of admission control, only VMware HA admission control can be disabled. However, without it there is no assurance that all virtual machines in the cluster can be restarted after a host failure. VMware recommends that you do not disable admission control, but you might need to do so temporarily, for the following reasons:

- If you need to violate the failover constraints when there are not enough resources to support them (for example, if you are placing hosts in standby mode to test them for use with DPM).
- If an automated process needs to take actions that might temporarily violate the failover constraints (for example, as part of an upgrade directed by VMware Update Manager).
- If you need to perform testing or maintenance operations.

Host Failures Cluster Tolerates

You can configure VMware HA to tolerate a specified number of host failures. With the Host Failures Cluster Tolerates admission control policy, VMware HA ensures that a specified number of hosts can fail and sufficient resources remain in the cluster to fail over all the virtual machines from those hosts.

With the Host Failures Cluster Tolerates policy, VMware HA performs admission control in the following way:

- 1 Calculates the slot size.

A slot is a logical representation of the memory and CPU resources that satisfy the requirements for any powered-on virtual machine in the cluster.

- 2 Determines how many slots each host in the cluster can hold.

- 3 Determines the Current Failover Capacity of the cluster.

This is the number of hosts that can fail and still leave enough slots to satisfy all of the powered-on virtual machines.

- 4 Determines whether the Current Failover Capacity is less than the Configured Failover Capacity (provided by the user).

If it is, admission control disallows the operation.

NOTE The maximum Configured Failover Capacity that you can set is four. Each cluster has up to five primary hosts and if all fail simultaneously, failover of all hosts might not be successful.

Slot Size Calculation

Slot size is comprised of two components, CPU and memory. VMware HA calculates these values.

- The CPU component by obtaining the CPU reservation of each powered-on virtual machine and selecting the largest value. If you have not specified a CPU reservation for a virtual machine, it is assigned a default value of 256 MHz (this value can be changed using the `das.vmCpuMinMHz` advanced attribute.)
- The memory component by obtaining the memory reservation (plus memory overhead) of each powered-on virtual machine and selecting the largest value.

If your cluster contains any virtual machines that have much larger reservations than the others, they will distort slot size calculation. To avoid this, you can specify an upper bound for the CPU or memory component of the slot size by using the `das.slotCpuInMHz` or `das.slotMemInMB` advanced attributes, respectively.

When using these advanced attributes, there is a risk of resource fragmentation where virtual machines larger than the slot size are assigned multiple slots. In a cluster that is close to capacity, there might be enough slots in aggregate for a virtual machine to be failed over. However, those slots could be located on multiple hosts and are unusable by a virtual machine assigned multiple slots because a virtual machine can run on only a single ESX/ESXi host at a time.

Using Slots to Compute the Current Failover Capacity

After the slot size is calculated, VMware HA determines each host's CPU and memory resources that are available for virtual machines. These amounts are those contained in the host's root resource pool, not the total physical resources of the host. Resources being used for virtualization purposes are not included. Only hosts that are connected, not in maintenance mode, and have no VMware HA errors are considered.

The maximum number of slots that each host can support is then determined. To do this, the host's CPU resource amount is divided by the CPU component of the slot size and the result is rounded down. The same calculation is made for the host's memory resource amount. These two numbers are compared and the lower is the number of slots that the host can support.

The Current Failover Capacity is computed by determining how many hosts (starting from the largest) can fail and still leave enough slots to satisfy the requirements of all powered-on virtual machines.

Advanced Runtime Info

When you select the Host Failures Cluster Tolerates admission control policy, the **Advanced Runtime Info** link appears in the VMware HA section of the cluster's **Summary** tab in the vSphere Client. Click this link to display the following information about the cluster:

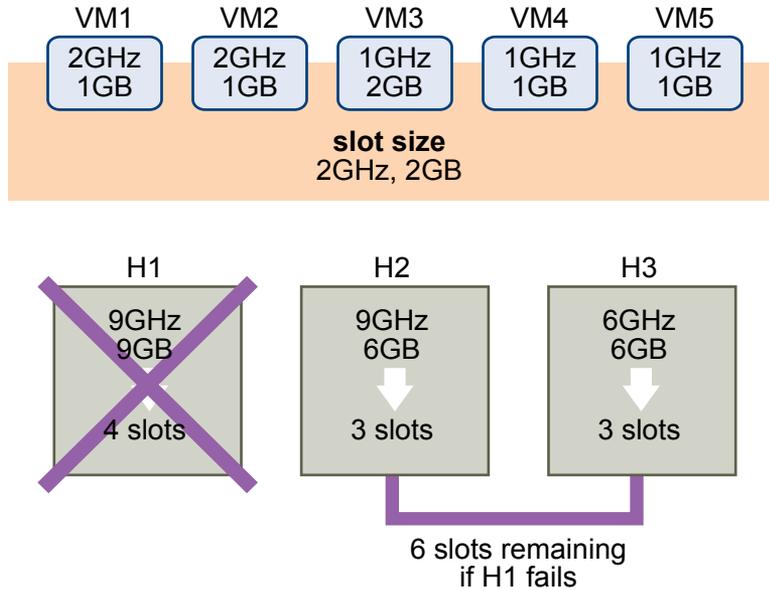
- Slot size.
- Total slots in cluster. The sum of the slots supported by the good hosts in the cluster.
- Used slots. The number of slots assigned to powered-on virtual machines. It can be more than the number of powered-on virtual machines if you have defined an upper bound for the slot size using the advanced options.

- Available slots. The number of slots available to power on additional virtual machines in the cluster. VMware HA automatically reserves the required number of slots for failover. The remaining slots are available to power on new virtual machines.
- Total powered on VMs in cluster.
- Total hosts in cluster.
- Total good hosts in cluster. The number of hosts that are connected, not in maintenance mode, and have no VMware HA errors.

Example: Admission Control Using Host Failures Cluster Tolerates Policy

The way that slot size is calculated and utilized with this admission control policy can be shown with an example. Make the following assumptions about a cluster:

- The cluster is comprised of three hosts, each with a different amount of available CPU and memory resources. The first host (H1) has 9GHz of available CPU resources and 9GB of available memory, while Host 2 (H2) has 9GHz and 6GB and Host 3 (H3) has 6GHz and 6GB.
- There are five powered-on virtual machines in the cluster with differing CPU and memory requirements. VM1 needs 2GHz of CPU resources and 1GB of memory, while VM2 needs 2GHz and 1GB, VM3 needs 1GHz and 2GB, VM4 needs 1GHz and 1GB, and VM5 needs 1GHz and 1GB.
- The Host Failures Cluster Tolerates is set to one.

Figure 2-1. Admission Control Example with Host Failures Cluster Tolerates Policy

- 1 Slot size is calculated by comparing both the CPU and memory requirements of the virtual machines and selecting the largest.

The largest CPU requirement (shared by VM1 and VM2) is 2GHz, while the largest memory requirement (for VM3) is 2GB. Based on this, the slot size is 2GHz CPU and 2GB memory.

- 2 Maximum number of slots that each host can support is determined.

H1 can support four slots. H2 can support three slots (which is the smaller of 9GHz/2GHz and 6GB/2GB) and H3 can also support three.

- 3 Current Failover Capacity is computed.

The largest host is H1 and if it fails, six slots remain in the cluster, which is sufficient for all five of the powered-on virtual machines. If both H1 and H2 fail, only three slots remain, which is insufficient. Therefore, the Current Failover Capacity is one.

The cluster has one available slot (the six slots on H2 and H3 minus the five used slots). VMware HA admission control allows you to power on one additional virtual machine (that does not exceed the slot size).

Percentage of Cluster Resources Reserved

You can configure VMware HA to perform admission control by reserving a specific percentage of cluster resources for recovery from host failures.

With the Percentage of Cluster Resources Reserved admission control policy, VMware HA ensures that a specified percentage of aggregate cluster resources is reserved for failover.

With the Cluster Resources Reserved policy, VMware HA performs admission control.

- 1 Calculates the total resource requirements for all powered-on virtual machines in the cluster.
- 2 Calculates the total host resources available for virtual machines.
- 3 Calculates the Current CPU Failover Capacity and Current Memory Failover Capacity for the cluster.
- 4 Determines if either the Current CPU Failover Capacity or Current Memory Failover Capacity is less than the Configured Failover Capacity (provided by the user).

If so, admission control disallows the operation.

It uses the actual reservations of the virtual machines. If a virtual machine does not have reservations, meaning that the reservation is 0, a default of 0MB memory and 256MHz CPU is applied. This is controlled by the same HA advanced options used for the failover level policy.

Computing the Current Failover Capacity

The total resource requirements for the powered-on virtual machines is comprised of two components, CPU and memory. VMware HA calculates these values.

- The CPU component by summing the CPU reservations of the powered-on virtual machines. If you have not specified a CPU reservation for a virtual machine, it is assigned a default value of 256 MHz (this value can be changed using the `das.vmCpuMinMHz` advanced attribute.)
- The memory component by summing the memory reservation (plus memory overhead) of each powered-on virtual machine.

The total host resources available for virtual machines is calculated by summing the hosts' CPU and memory resources. These amounts are those contained in the host's root resource pool, not the total physical resources of the host. Resources being used for virtualization purposes are not included. Only hosts that are connected, not in maintenance mode, and have no VMware HA errors are considered.

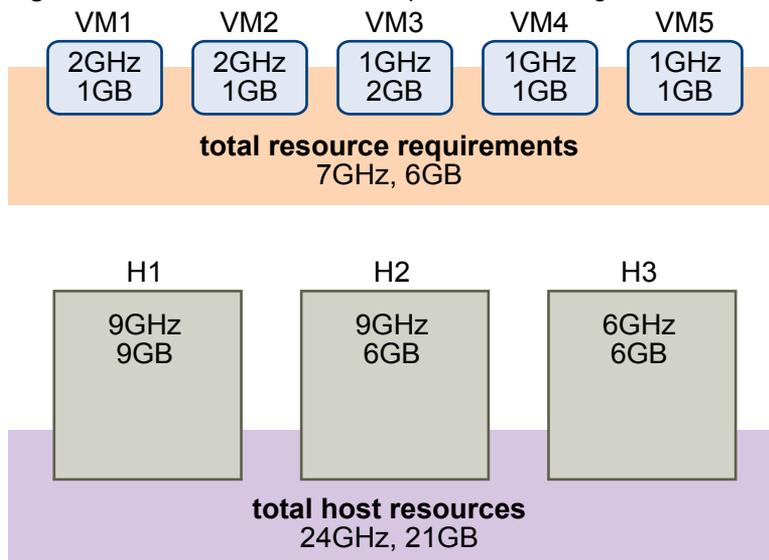
The Current CPU Failover Capacity is computed by subtracting the total CPU resource requirements from the total host CPU resources and dividing the result by the total host CPU resources. The Current Memory Failover Capacity is calculated similarly.

Example: Admission Control Using Percentage of Cluster Resources Reserved Policy

The way that Current Failover Capacity is calculated and utilized with this admission control policy can be shown with an example. Make the following assumptions about a cluster:

- The cluster is comprised of three hosts, each with a different amount of available CPU and memory resources. The first host (H1) has 9GHz of available CPU resources and 9GB of available memory, while Host 2 (H2) has 9GHz and 6GB and Host 3 (H3) has 6GHz and 6GB.
- There are five powered-on virtual machines in the cluster with differing CPU and memory requirements. VM1 needs 2GHz of CPU resources and 1GB of memory, while VM2 needs 2GHz and 1GB, VM3 needs 1GHz and 2GB, VM4 needs 1GHz and 1GB, and VM5 needs 1GHz and 1GB.
- The Configured Failover Capacity is set to 25%.

Figure 2-2. Admission Control Example with Percentage of Cluster Resources Reserved Policy



The total resource requirements for the powered-on virtual machines is 7GHz and 6GB. The total host resources available for virtual machines is 24GHz and 21GB. Based on this, the Current CPU Failover Capacity is 70% $((24\text{GHz} - 7\text{GHz})/24\text{GHz})$. Similarly, the Current Memory Failover Capacity is 71% $((21\text{GB} - 6\text{GB})/21\text{GB})$.

Because the cluster's Configured Failover Capacity is set to 25%, 45% of the cluster's total CPU resources and 46% of the cluster's memory resources are still available to power on additional virtual machines.

Specify a Failover Host

You can configure VMware HA to designate a specific host as the failover host.

With the Specify a Failover Host admission control policy, when a host fails, VMware HA attempts to restart its virtual machines on a specified failover host. If this is not possible, for example the failover host itself has failed or it has insufficient resources, then VMware HA attempts to restart those virtual machines on another host in the cluster.

To ensure that spare capacity is available on the failover host, you are prevented from powering on virtual machines or using VMotion to migrate virtual machines to the failover host. Also, DRS does not use the failover host for load balancing.

The Current Failover Host is displayed in the VMware HA section of the cluster's Summary tab in the vSphere Client. The status icon next to the host can be green, yellow, or red.

- Green. The host is connected, not in maintenance mode, and has no VMware HA errors. Also, no powered-on virtual machines reside on the host.
- Yellow. The host is connected, not in maintenance mode, and has no VMware HA errors. However, powered-on virtual machines reside on the host.
- Red. The host is disconnected, in maintenance mode, or has VMware HA errors.

Choosing an Admission Control Policy

You should choose a VMware HA admission control policy based on your availability needs and the characteristics of your cluster.

When choosing an admission control policy, you should consider a number of factors.

Avoiding Resource Fragmentation

Resource fragmentation occurs when there are enough resources in aggregate for a virtual machine to be failed over. However, those resources are located on multiple hosts and are unusable because a virtual machine can run on one ESX/ESXi host at a time. The Host Failures Cluster Tolerates policy avoids resource fragmentation by defining a slot as the maximum virtual machine reservation. The Percentage of Cluster Resources policy does not address the problem of resource fragmentation. With the Specify a Failover Host policy, resources are not fragmented because a single host is reserved for failover.

Flexibility of Failover Resource Reservation

Admission control policies differ in the granularity of control they give you when reserving cluster resources for failover protection. The Host Failures Cluster Tolerates policy allows you to set the failover level from one to four hosts. The Percentage of Cluster Resources policy allows you to designate up to 50% of cluster resources for failover. The Specify a Failover Host policy only allows you to specify a single failover host.

Heterogeneity of Cluster

Clusters can be heterogeneous in terms of virtual machine resource reservations and host total resource capacities. In a heterogeneous cluster, the Host Failures Cluster Tolerates policy can be too conservative because it only considers the largest virtual machine reservations when defining slot size and assumes the largest hosts fail when computing the Current Failover Capacity. The other two admission control policies are not affected by cluster heterogeneity.

Creating a VMware HA Cluster

VMware HA operates in the context of a cluster of ESX/ESXi hosts. You must create a cluster, populate it with hosts, and configure VMware HA settings before failover protection can be established.

When you create a VMware HA cluster, you must configure a number of settings that determine how the feature works. Before you do this, first identify your cluster's nodes. These are the ESX/ESXi hosts that will provide the resources to support virtual machines and that VMware HA will use for failover protection. Then you should determine how those nodes are to be connected to one another and to the shared storage where your virtual machine data resides. After that networking architecture is in place, you can add the hosts to the cluster and finish configuring VMware HA.

You can enable and configure VMware HA before you add host nodes to the cluster. However, until the hosts are added your cluster is not fully operational and some of the cluster settings are unavailable. For example, the Specify a Failover Host admission control policy is unavailable until there is a host that can be designated as the failover host.

NOTE The Virtual Machine Startup and Shutdown (automatic startup) feature is disabled for all virtual machines residing on hosts that are in (or moved into) a VMware HA cluster. Automatic startup is not supported when used with VMware HA.

Create a VMware HA Cluster

Your cluster can be enabled for VMware HA, and a VMware HA-enabled cluster is a prerequisite for Fault Tolerance. VMware recommends that you first create an empty cluster. After you have planned the resources and networking architecture of your cluster, you can use the vSphere Client to add hosts to the cluster and specify the cluster's VMware HA settings.

Connect vSphere Client to vCenter Server using an account with cluster administrator permissions.

Prerequisites

All virtual machines and their configuration files must reside on shared storage. So that you can power on the virtual machines using different hosts in the cluster, the hosts must be configured to access that shared storage.

Each host in a VMware HA cluster must have a host name assigned and a static IP address associated with each of the virtual NICs.

Hosts must be configured to have access to the virtual machine network.

VMware recommends redundant network connections for VMware HA.

- For ESX, set up redundant service console networking.
- For ESXi, set up redundant VMkernel networking.

For information about setting up network redundancy, see [“Network Path Redundancy,”](#) on page 29.

Procedure

- 1 Select the Hosts & Clusters view.
- 2 Right-click the Datacenter in the Inventory tree and click **New Cluster**.
- 3 Complete the New Cluster wizard.
 - Do not enable VMware HA (or DRS) at this time.
- 4 Click **Finish** to close the wizard and create the cluster.
 - You have created an empty cluster.

- 5 Based on your plan for the resources and networking architecture of the cluster, use the vSphere Client to add hosts to the cluster.
 - 6 Right-click the cluster and click **Edit Settings**.
The cluster's Settings dialog box is where you can modify the VMware HA (and other) settings for the cluster.
 - 7 On the Cluster Features page, select **Turn On VMware HA**.
 - 8 Configure the VMware HA settings as appropriate for your cluster.
 - Host Monitoring Status
 - Admission Control
 - Virtual Machine Options
 - VM Monitoring
 - 9 Click **OK** to close the cluster's Settings dialog box.
- A configured VMware HA cluster, populated with hosts, is created.

Cluster Features

The first panel in the New Cluster wizard allows you to specify basic options for the cluster.

In this panel you can specify the cluster name and choose one or both cluster features.

Name	Specifies the name of the cluster. This name appears in the vSphere Client inventory panel. You must specify a name to continue with cluster creation.
Turn On VMware HA	If this check box is selected, virtual machines are restarted on another host in the cluster if a host fails. You must turn on VMware HA to enable VMware Fault Tolerance on any virtual machine in the cluster.
Turn On VMware DRS	If this check box is selected, DRS balances the load of virtual machines across the cluster. DRS also places and migrates virtual machines when they are protected with HA.

You can change any of these cluster features at a later time.

Host Monitoring Status

After you create a cluster, enable Host Monitoring so that VMware HA can monitor heartbeats sent by ESX/ESXi hosts in the cluster.

If **Enable Host Monitoring** is selected, each ESX/ESXi host in the cluster is checked to ensure it is running. If a host failure occurs, virtual machines are restarted on another host. Host Monitoring is also required for the VMware Fault Tolerance recovery process to work properly.

If you need to perform network maintenance that might trigger host isolation responses, VMware recommends that you first suspend VMware HA by disabling Host Monitoring. After the maintenance is complete, reenabling Host Monitoring.

Enabling or Disabling Admission Control

The New Cluster wizard allows you to enable or disable admission control for the VMware HA cluster and choose a policy for how it is enforced.

You can enable or disable admission control for the HA cluster.

Prevent VMs from being powered on if they violate availability constraints Enables admission control and enforces availability constraints and preserves failover capacity. Any operation on a virtual machine that decreases the unreserved resources in the cluster and violates availability constraints is not permitted.

Allow VMs to be powered on even if they violate availability constraints Disables admission control. If you select this option, virtual machines can, for example, be powered on even if that causes insufficient failover capacity. When this is done, no warnings are presented, and the cluster does not turn red. If a cluster has insufficient failover capacity, VMware HA can still perform failovers and it uses the VM Restart Priority setting to determine which virtual machines to power on first.

VMware HA provides three policies for enforcing admission control, if it is enabled.

- Host failures cluster tolerates
- Percentage of cluster resources reserved as failover spare capacity
- Specify a failover host

NOTE See [“Choosing an Admission Control Policy,”](#) on page 20 for more information about how VMware HA admission control works.

Virtual Machine Options

Default virtual machine settings control the order in which virtual machines are restarted and how VMware HA responds if hosts lose network connectivity with other hosts. These settings apply to all virtual machines in the cluster in the case of a host failure or isolation. You can configure exceptions for each virtual machine.

VM Restart Priority

VM restart priority determines the relative order in which virtual machines are restarted after a host failure. Such virtual machines are restarted sequentially on new hosts, with the highest priority virtual machines first and continuing to those with lower priority until all virtual machines are restarted or no more cluster resources are available. If the number of hosts failures or virtual machines restarts exceeds what admission control permits, the virtual machines with lower priority might not be restarted until more resources become available. Virtual machines are restarted on the failover host, if one is specified, or on the host with the highest percentage of available resources.

The values for this setting are: Disabled, Low, Medium (the default), and High. If Disabled is selected, VMware HA is disabled for the virtual machine, meaning that it is not restarted on other ESX/ESXi hosts if its ESX/ESXi host fails. If Disabled is selected, this does not affect virtual machine monitoring, which means that if a virtual machine fails on a host that is functioning properly, that virtual machine is reset on that same host. You can change this property for individual virtual machines.

The restart priority settings for virtual machines vary depending on user needs. VMware recommends that you assign higher restart priority to the virtual machines that provide the most important services.

For example, in the case of a multitier application you might rank assignments according to functions hosted on the virtual machines.

- High. Database servers that will provide data for applications.
- Medium. Application servers that consume data in the database and provide results on web pages.
- Low. Web servers that receive user requests, pass queries to application servers, and return results to users.

Host Isolation Response

Host isolation response determines what happens when a host in a VMware HA cluster loses its service console networks (or VMkernel networks, in ESXi) connection but continues running. Host isolation responses require that Host Monitoring Status is enabled. If it is disabled, host isolation responses are also suspended. A host determines that it is isolated when it stops receiving heartbeats from all other hosts and it is unable to ping its isolation addresses. When this occurs, the host executes its isolation response. The responses are: Leave powered on, Power off, and Shut down (the default). You can customize this property for individual virtual machines.

To use the Shut down VM setting, you must install VMware Tools in the guest operating system of the virtual machine. Shutting down the virtual machine provides the advantage of preserving its state. This is better than powering it off, which does not flush most recent changes to disk or commit transactions. Virtual machines that are shut down will take longer to fail over while the shutdown completes. Virtual Machines that have not shut down in 300 seconds, or the time specified in the advanced attribute `das.isolationShutdownTimeout` seconds, are powered off.

NOTE After you create a VMware HA cluster, you can override the default cluster settings for Restart Priority and Isolation Response for specific virtual machines. Such overrides are useful for virtual machines that are used for special tasks. For example, virtual machines that provide infrastructure services like DNS or DHCP might need to be powered on before other virtual machines in the cluster.

VM Monitoring

VM Monitoring restarts individual virtual machines if their VMware Tools heartbeats are not received within a set time. You can configure the degree to which VMware HA is sensitive to such non-responsiveness.

If you select **Enable VM Monitoring**, the VM Monitoring service (using VMware Tools) evaluates whether each virtual machine in the cluster is running by checking for regular heartbeats from the VMware Tools process running inside the guest. If no heartbeats are received, this is most likely because the guest operating system has failed or VMware Tools is not being allocated any time to complete tasks. In such a case, the VM Monitoring service determines that the virtual machine has failed and the virtual machine is rebooted to restore service.

You can also configure the level of monitoring sensitivity. Highly sensitive monitoring results in a more rapid conclusion that a failure has occurred. While unlikely, highly sensitive monitoring might lead to falsely identifying failures when the virtual machine in question is actually still working, but heartbeats have not been received due to factors such as resource constraints. Low sensitivity monitoring results in longer interruptions in service between actual failures and virtual machines being reset. Select an option that is an effective compromise for your needs.

After failures are detected, VMware HA resets virtual machines. This helps ensure that services remain available. To avoid resetting virtual machines repeatedly for nontransient errors, by default virtual machines will be reset only three times during a certain configurable time interval. After virtual machines have been reset three times, VMware HA makes no further attempts to reset the virtual machines after any subsequent failures until after the specified time has elapsed. You can configure the number of resets using the **Maximum per-VM resets** custom setting.

Occasionally, virtual machines that are still functioning properly stop sending heartbeats. To avoid unnecessarily resetting such virtual machines, the VM Monitoring service also monitors a virtual machine's I/O activity. If no heartbeats are received within the failure interval, the I/O stats interval (a cluster-level attribute) is checked. The I/O stats interval determines if any disk or network activity has occurred for the virtual machine during the previous two minutes (120 seconds). If not, the virtual machine is reset. This default value (120 seconds) can be changed using the advanced attribute `das.iostatsInterval`.

NOTE The VM Monitoring settings cannot be configured through advanced attributes. Modify settings in the VM Monitoring page of the cluster's Settings dialog box.

The default settings for VM Monitoring sensitivity are described in the table.

Table 2-1. VM Monitoring Settings

Setting	Failure Interval (seconds)	Reset Period
High	30	1 hour
Medium	60	24 hours
Low	120	7 days

You can specify custom values for both VM Monitoring sensitivity and the I/O stats interval, as described in [“Customizing VMware HA Behavior,”](#) on page 25.

Customizing VMware HA Behavior

After you have established a cluster, you can modify the specific attributes that affect how VMware HA behaves. You can also change the cluster default settings inherited by individual virtual machines.

This section guides you through setting advanced attributes for VMware HA and lists a few attributes you might want to set. Because these attributes affect the functioning of HA, change them with caution. Review the advanced settings you can use to optimize the VMware HA clusters in your environment.

Table 2-2. VMware HA Attributes

Attribute	Description
<code>das.isolationaddress[...]</code>	Sets the address to ping to determine if a host is isolated from the network. This address is pinged only when heartbeats are not received from any other host in the cluster. If not specified, the default gateway of the console network is used. This default gateway has to be a reliable address that is available, so that the host can determine if it is isolated from the network. You can specify multiple isolation addresses (up to 10) for the cluster: <code>das.isolationaddressX</code> , where <code>X = 1-10</code> . Typically you should specify one per service console. Specifying too many addresses makes isolation detection take too long and can affect VMware HA behavior.
<code>das.usedefaultisolationaddress</code>	By default, VMware HA uses the default gateway of the console network as an isolation address. This attribute specifies whether or not this default is used (<code>true false</code>).
<code>das.failedetectiontime</code>	Changes the default failure detection time for host monitoring. The default is 15000 milliseconds (15 seconds). This is the time period, when a host has received no heartbeats from another host, that it waits before declaring that host as failed.
<code>das.failedetectioninterval</code>	Changes the heartbeat interval among VMware HA hosts. By default, this occurs every 1000 milliseconds (1 second).

Table 2-2. VMware HA Attributes (Continued)

Attribute	Description
das.defaultfailoverhost	Defines the host that VMware HA tries to fail virtual machines over to. Use this option only if the VMware HA admission control policy is failover level or cluster resource percentage. If this option is used with the failover host admission control policy, it takes precedence over the failover host named in the policy. You can define only one failover host.
das.isolationShutdownTimeout	The period of time the system waits for a virtual machine to shut down before powering it off. This only applies if the host's isolation response is Shut down VM. Default value is 300 seconds.
das.slotMemInMB	Defines the maximum bound on the memory slot size. If this option is used, the slot size is the smaller of this value or the maximum memory reservation plus memory overhead of any powered-on virtual machine in the cluster.
das.slotCpuInMHz	Defines the maximum bound on the CPU slot size. If this option is used, the slot size is the smaller of this value or the maximum CPU reservation of any powered-on virtual machine in the cluster.
das.vmMemoryMinMB	Defines the default memory resource value assigned to a virtual machine if its memory reservation is not specified or zero. This is used for the Host Failures Cluster Tolerates admission control policy. If no value is specified, the default is 0 MB.
das.vmCpuMinMHz	Defines the default CPU resource value assigned to a virtual machine if its CPU reservation is not specified or zero. This is used for the Host Failures Cluster Tolerates admission control policy. If no value is specified, the default is 256MHz.
das.iostatsInterval	Changes the default I/O stats interval for VM monitoring sensitivity. The default is 120 (seconds). Can be set to any value greater than, or equal to 0. Setting to 0 disables the check.

NOTE If you change the value of any of the following advanced attributes, you must disable and then re-enable VMware HA before your changes take effect.

- das.isolationaddress[...]
- das.usedefaultisolationaddress
- das.failedetectiontime
- das.failedetectioninterval
- das.isolationShutdownTimeout

Set Advanced VMware HA Options

To customize VMware HA behavior, set advanced VMware HA options.

Prerequisites

A VMware HA cluster for which to modify settings.

Cluster administrator privileges.

Procedure

- 1 In the cluster's Settings dialog box, select **VMware HA**.
- 2 Click the **Advanced Options** button to open the Advanced Options (HA) dialog box.
- 3 Enter each advanced attribute you want to change in a text box in the **Option** column and enter a value in the **Value** column.
- 4 Click **OK**.

The cluster uses options you added or modified.

Customize VMware HA Behavior for an Individual Virtual Machine

Each virtual machine in a VMware HA cluster is assigned the cluster default settings for VM Restart Priority, Host Isolation Response, and VM Monitoring. You can specify specific behavior for each virtual machine by changing these defaults. If the virtual machine leaves the cluster, these settings are lost.

Cluster administrator privileges.

Procedure

- 1 Select the cluster and select **Edit Settings** from the right-click menu.
- 2 Select **Virtual Machine Options** under VMware HA.
- 3 In the Virtual Machine Settings pane, select a virtual machine and customize its **VM Restart Priority** or **Host Isolation Response** setting.
- 4 Select **VM Monitoring** under VMware HA.
- 5 In the Virtual Machine Settings pane, select a virtual machine and customize its **VM Monitoring** setting.
- 6 Click **OK**.

The virtual machine's behavior now differs from the cluster defaults for each setting you changed.

Best Practices for VMware HA Clusters

To help ensure optimal VMware HA cluster performance, VMware recommends that you follow certain best practices. Also, networking configuration and redundancy are important when designing and implementing your cluster.

Setting Alarms to Monitor Cluster Changes

When VMware HA or Fault Tolerance take action to maintain availability, for example, a virtual machine failover, you might need to be notified about such changes. You can configure alarms in vCenter Server to be triggered when these actions are taken, and have alerts, such as emails, sent to a specified set of administrators.

Monitoring Cluster Validity

A valid cluster is one in which the admission control policy has not been violated.

A cluster enabled for VMware HA becomes invalid (red) when the number of virtual machines powered on exceeds the failover requirements, that is, the current failover capacity is smaller than configured failover capacity. If admission control is disabled, clusters do not become invalid.

The cluster's Summary page in the vSphere Client displays a list of configuration issues for clusters. The list explains what has caused the cluster to become invalid or over-committed (yellow).

DRS behavior is not affected if a cluster is red because of a VMware HA issue.

Networking Best Practices

VMware recommends some best practices for the configuration of host NICs and network topology for VMware HA. This includes recommendations not only for your ESX/ESXi hosts, but also for cabling, switches, routers, and firewalls.

Network Configuration and Maintenance

The following network maintenance suggestions can help you avoid the accidental detection of failed hosts and network isolation due to dropped VMware HA heartbeats.

- When making changes to the network(s) that your clustered ESX/ESXi hosts are on, VMware recommends that you suspend the Host Monitoring feature. Changing your network hardware or networking settings can interrupt the heartbeats that VMware HA uses to detect host failures, and this might result in unwanted attempts to fail over virtual machines.
- When you change the networking configuration on the ESX/ESXi hosts themselves, for example, adding port groups, or removing vSwitches, VMware recommends that in addition to suspending Host Monitoring, you place the host in maintenance mode.

NOTE Because networking is a vital component of VMware HA, if network maintenance needs to be performed the VMware HA administrator should be informed.

Networks Used for VMware HA Communications

To identify which network operations might disrupt the functioning of VMware HA, you should be aware of which network(s) are being used for heart beating and other VMware HA communications.

- On ESX hosts in the cluster, VMware HA communications travel over all networks that are designated as service console networks. VMkernel networks are not used by these hosts for VMware HA communications.
- On ESXi hosts in the cluster, VMware HA communications, by default, travel over VMkernel networks, except those marked for use with VMotion. If there is only one VMkernel network, VMware HA shares it with VMotion, if necessary. With ESXi 4.0, you must also explicitly enable the Management Network checkbox for VMware HA to use this network.

Cluster-Wide Networking Considerations

For VMware HA to function, all hosts in the cluster must have compatible networks. The first node added to the cluster dictates the networks that all subsequent hosts allowed into the cluster must also have. Networks are considered compatible if the combination of the IP address and subnet mask result in a network that matches another host's. If you attempt to add a host with too few, or too many, networks, or if the host being added has incompatible networks, the configuration task fails, and the Task Details pane specifies this incompatibility.

For example, if the first host you add to the cluster has two networks being used for VMware HA communications, 10.10.135.0/255.255.255.0 and 10.17.142.0/255.255.255.0, all subsequent hosts must have the same two networks configured and used for VMware HA communications.

Network Isolation Addresses

A network isolation address is an IP address that is pinged to determine if a host is isolated from the network. This address is pinged only when a host has stopped receiving heartbeats from all other hosts in the cluster. If a host can ping its network isolation address, the host is not network isolated, and the other hosts in the cluster have failed. However, if the host cannot ping its isolation address, it is likely that the host has become isolated from the network and no failover action is taken.

By default, the network isolation address is the default gateway for the host. There is only one default gateway specified, regardless of how many service console networks have been defined, so you should use the `das.isolationaddress[...]` advanced attribute to add isolation addresses for additional networks. For example, `das.isolationAddress2` to add an isolation address for your second network, `das.isolationAddress3` for the third, up to a maximum of `das.isolationAddress9` for the ninth.

When you specify additional isolation address, VMware recommends that you increase the setting for the `das.failedetectiontime` advanced attribute to 20000 milliseconds (20 seconds) or greater. A node that is isolated from the network needs time to release its virtual machine's VMFS locks if the host isolation response is to fail over the virtual machines (not to leave them powered on.) This must happen before the other nodes declare the node as failed, so that they can power on the virtual machines, without getting an error that the virtual machines are still locked by the isolated node.

For more information on VMware HA advanced attributes, see [“Customizing VMware HA Behavior,”](#) on page 25.

Other Networking Considerations

Configuring Switches. If the physical network switches that connect your servers support the PortFast (or an equivalent) setting, enable it. This setting prevents a host from incorrectly determining that a network is isolated during the execution of lengthy spanning tree algorithms.

Host Firewalls. On ESX/ESXi hosts, VMware HA needs and automatically opens the following firewall ports.

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

Port Group Names and Network Labels. Use consistent port group names and network labels on VLANs for public networks. Port group names are used to reconfigure access to the network by virtual machines. If you use inconsistent names between the original server and the failover server, virtual machines are disconnected from their networks after failover. Network labels are used by virtual machines to reestablish network connectivity upon restart.

Network Path Redundancy

Network path redundancy between cluster nodes is important for VMware HA reliability. A single service console network ends up being a single point of failure and can result in failovers although only the network has failed.

If you have only one service console network, any failure between the host and the cluster can cause an unnecessary (or false) failover situation. Possible failures include NIC failures, network cable failures, network cable removal, and switch resets. Consider these possible sources of failure between hosts and try to minimize them, typically by providing network redundancy.

You can implement network redundancy at the NIC level with NIC teaming, or at the service console (or VMkernel port on ESXi) level. In most implementations, NIC teaming provides sufficient redundancy, but you can use or add service console (or VMkernel port) redundancy if required. Redundant service console networking on ESX (or VMkernel networking) allows the reliable detection of failures and prevents isolation conditions from occurring, because heartbeats can be sent over multiple networks.

Configure the fewest possible number of hardware segments between the servers in a cluster. The goal being to limit single points of failure. Additionally, routes with too many hops can cause networking packet delays for heartbeats, and increase the possible points of failure.

Network Redundancy Using NIC Teaming

Using a team of two NICs connected to separate physical switches improves the reliability of a service console (or, in ESXi, VMkernel) network. Because servers connected through two NICs (and through separate switches) have two independent paths for sending and receiving heartbeats, the cluster is more resilient. To configure a NIC team for the service console, configure the vNICs in vSwitch configuration for Active or Standby configuration. The recommended parameter settings for the vNICs are:

- Default load balancing = route based on originating port ID
- Failback = No

After you have added a NIC to a host in your VMware HA cluster, you must reconfigure VMware HA on that host.

Network Redundancy Using a Secondary Network

As an alternative to NIC teaming for providing redundancy for heartbeats, you can create a secondary service console (or VMkernel port for ESXi), which is attached to a separate virtual switch. The primary service console is used for network and management purposes. When the secondary service console network is created, VMware HA sends heartbeats over both the primary and secondary service consoles. If one path fails, VMware HA can still send and receive heartbeats over the other path.

Providing Fault Tolerance for Virtual Machines

3

You can enable VMware Fault Tolerance for your virtual machines to ensure business continuity with higher levels of availability and data protection than is offered by VMware HA.

Fault Tolerance is built on the ESX/ESXi host platform (using the VMware vLockstep functionality) and it provides continuous availability by having identical virtual machines run in virtual lockstep on separate hosts.

To obtain the optimal results from Fault Tolerance you should be familiar with how it works, how to enable it for your cluster and virtual machines, the best practices for its usage, and troubleshooting tips.

This chapter includes the following topics:

- [“How Fault Tolerance Works,”](#) on page 31
- [“Fault Tolerance Use Cases,”](#) on page 32
- [“Fault Tolerance Configuration Requirements,”](#) on page 33
- [“Fault Tolerance Interoperability,”](#) on page 34
- [“Preparing Your Cluster and Hosts for Fault Tolerance,”](#) on page 35
- [“Turning On Fault Tolerance for Virtual Machines,”](#) on page 37
- [“Viewing Information About Fault Tolerant Virtual Machines,”](#) on page 39
- [“Fault Tolerance Best Practices,”](#) on page 40
- [“VMware Fault Tolerance Configuration Recommendations,”](#) on page 41
- [“Troubleshooting Fault Tolerance,”](#) on page 42

How Fault Tolerance Works

VMware Fault Tolerance provides continuous availability for virtual machines by creating and maintaining a Secondary VM that is identical to, and continuously available to replace, the Primary VM in the event of a failover situation.

You can enable Fault Tolerance for most mission critical virtual machines. A duplicate virtual machine, called the Secondary VM, is created and runs in virtual lockstep with the Primary VM. VMware vLockstep captures inputs and events that occur on the Primary VM and sends them to the Secondary VM, which is running on another host. Using this information, the Secondary VM's execution is identical to that of the Primary VM. Because the Secondary VM is in virtual lockstep with the Primary VM, it can take over execution at any point without interruption, thereby providing fault tolerant protection.

The Primary and Secondary VMs continuously exchange heartbeats. This allows the virtual machine pair to monitor the status of one another to ensure that Fault Tolerance is continually maintained. A transparent failover occurs if the host running the Primary VM fails, in which case the Secondary VM is immediately activated to replace the Primary VM. A new Secondary VM is started and Fault Tolerance redundancy is reestablished within a few seconds. If the host running the Secondary VM fails, it is also immediately replaced. In either case, users experience no interruption in service and no loss of data.

A fault tolerant virtual machine and its secondary copy are not allowed to run on the same host. Fault Tolerance uses anti-affinity rules, which ensure that the two instances of the fault tolerant virtual machine are never on the same host. This ensures that a host failure cannot result in the loss of both virtual machines.

Fault Tolerance avoids "split-brain" situations, which can lead to two active copies of a virtual machine after recovery from a failure. Atomic file locking on shared storage is used to coordinate failover so that only one side continues running as the Primary VM and a new Secondary VM is respawned automatically.

NOTE The anti-affinity check is performed when the Primary VM is powered on. It is possible that the Primary and Secondary VMs can be on the same host when they are both in a powered-off state. This is normal behavior and when the Primary VM is powered on, the Secondary VM is started on a different host at that time.

Fault Tolerance Use Cases

Several typical situations can benefit from the use of VMware Fault Tolerance.

Fault Tolerance provides a higher level of business continuity than VMware HA. When a Secondary VM is called upon to replace its Primary VM counterpart, the Secondary VM immediately takes over the Primary VM's role with the entire state of the virtual machine preserved. Applications are already running, and data stored in memory does not need to be re-entered or reloaded. This differs from a failover provided by VMware HA, which restarts the virtual machines affected by a failure.

This higher level of continuity and the added protection of state information and data informs the scenarios when you might want to deploy Fault Tolerance.

- Applications that need to be available at all times, especially those that have long-lasting client connections that users want to maintain during hardware failure.
- Custom applications that have no other way of doing clustering.
- Cases where high availability might be provided through custom clustering solutions, which are too complicated to configure and maintain.

On-Demand Fault Tolerance

Another key use case for protecting a virtual machine with Fault Tolerance can be described as On-Demand Fault Tolerance. In this case, a virtual machine is adequately protected with VMware HA during normal operation. During certain critical periods, you might want to enhance the protection of the virtual machine. For example, you might be executing a quarter-end report which, if interrupted, might delay the availability of mission critical information. With VMware Fault Tolerance, you can protect this virtual machine prior to running this report and then turn off or disable Fault Tolerance after the report has been produced. You can use On-Demand Fault Tolerance to protect the virtual machine during a critical time period and return the resources to normal during non-critical operation.

Fault Tolerance Configuration Requirements

For VMware Fault Tolerance (FT) to perform as expected, the configuration of your cluster, hosts, and virtual machines must meet specific requirements.

Cluster Prerequisites

Unlike VMware HA which, by default, protects every virtual machine in the cluster, VMware Fault Tolerance is enabled on individual virtual machines. For a cluster to support VMware Fault Tolerance, the following prerequisites must be met:

- VMware HA must be enabled on the cluster. Host Monitoring should also be enabled. If it is not, when Fault Tolerance uses a Secondary VM to replace a Primary VM no new Secondary VM is created and redundancy is not restored.
- Host certificate checking must be enabled for all hosts that will be used for Fault Tolerance. See [“Enable Host Certificate Checking,”](#) on page 36.
- Each host must have a vMotion and a Fault Tolerance Logging NIC configured. See [“Configure Networking for Host Machines,”](#) on page 36.
- At least two hosts must have processors from the same compatible processor group. While Fault Tolerance supports heterogeneous clusters (a mix of processor groups), you get the maximum flexibility if all hosts are compatible. See the VMware knowledge base article at <http://kb.vmware.com/kb/1008027> for information on supported processors.
- All hosts must have the same ESX/ESXi version and patch level.
- All hosts must have access to the virtual machines' datastores and networks.

To confirm the compatibility of the hosts in the cluster to support Fault Tolerance, run profile compliance checks.

NOTE VMware HA includes the resource usage of Fault Tolerance Secondary VMs when it performs admission control calculations. For the Host Failures Cluster Tolerates policy, a Secondary VM is assigned a slot, and for the Percentage of Cluster Resources policy, the Secondary VM's resource usage is accounted for when computing the usable capacity of the cluster. See [“VMware HA Admission Control,”](#) on page 15.

Host Prerequisites

A host can support fault tolerant virtual machines if it meets the following requirements.

- A host must have processors from the FT-compatible processor group. See the VMware knowledge base article at <http://kb.vmware.com/kb/1008027>.
- A host must be certified by the OEM as FT-capable. Refer to the current Hardware Compatibility List (HCL) for a list of FT-supported servers (see <http://www.vmware.com/resources/compatibility/search.php>).
- The host configuration must have Hardware Virtualization (HV) enabled in the BIOS. Some hardware manufacturers ship their products with HV disabled. The process for enabling HV varies among BIOSes. See the documentation for your hosts' BIOSes for details on how to enable HV. If HV is not enabled, attempts to power on a fault tolerant virtual machine produce an error and the virtual machine does not power on.

Review the Host Configuration Section of [“Fault Tolerance Best Practices,”](#) on page 40 to select host options that best support VMware Fault Tolerance.

Virtual Machine Requirements

Before Fault Tolerance can be turned on, a virtual machine must meet minimum requirements.

- Virtual machine files must be stored on shared storage. Acceptable shared storage solutions include Fibre Channel, (hardware and software) iSCSI, NFS, and NAS.
- Virtual machines must be stored in virtual RDM or virtual machine disk (VMDK) files that are thick provisioned with the Cluster Features option. If a virtual machine is stored in a VMDK file that is thin provisioned or thick provisioned without clustering features enabled and an attempt is made to enable Fault Tolerance, a message appears indicating that the VMDK file must be converted. Users can accept this automatic conversion (which requires the virtual machine to be powered off), allowing the disk to be converted and the virtual machine to be protected with Fault Tolerance. The amount of time needed for this conversion process can vary depending on the size of the disk and the host's processor type.
- Virtual machines must be running on one of the supported guest operating systems. See the VMware knowledge base article at <http://kb.vmware.com/kb/1008027> for more information.

Fault Tolerance Interoperability

Before configuring VMware Fault Tolerance, you should be aware of the features and products Fault Tolerance cannot interoperate with.

Fault Tolerance and vSphere Features Not Supported

The following vSphere features are not supported for fault tolerant virtual machines.

- Snapshots. Snapshots must be removed or committed before Fault Tolerance can be enabled on a virtual machine. In addition, it is not possible to take snapshots of virtual machines on which Fault Tolerance is enabled.
- Storage VMotion. You cannot invoke Storage VMotion for virtual machines with Fault Tolerance turned on. To migrate the storage, you should temporarily turn off Fault Tolerance, and perform the storage VMotion action. When this is complete, you can turn Fault Tolerance back on.
- DRS features. A fault tolerant virtual machine is automatically configured as DRS-disabled. DRS does initially place a Secondary VM, however, DRS does not make recommendations or load balance Primary or Secondary VMs when load balancing the cluster. The Primary and Secondary VMs can be manually migrated during normal operation.

Other Features Incompatible with Fault Tolerance

For a virtual machine to be compatible with Fault Tolerance, the Virtual Machine must not use the following features or devices.

Table 3-1. Features and Devices Incompatible with Fault Tolerance and Corrective Actions

Incompatible Feature or Device	Corrective Action
Symmetric multiprocessor (SMP) virtual machines. Only virtual machines that support a single vCPU are compatible with Fault Tolerance.	Reconfigure the virtual machine as a single vCPU. Many workloads have good performance configured as a single vCPU.
Physical Raw Disk mapping (RDM). If you want to use Raw Disk Mapping (RDM) for your virtual disks, only virtual RDMs are supported.	Reconfigure virtual machines with physical RDM-backed virtual devices to use virtual RDMs instead.
CD-ROM or floppy virtual devices backed by a physical or remote device.	Remove the CD-ROM or floppy virtual device or reconfigure the backing with an ISO installed on shared storage.
Paravirtualized guests.	If paravirtualization is not required, reconfigure the virtual machine without a VMI ROM.

Table 3-1. Features and Devices Incompatible with Fault Tolerance and Corrective Actions (Continued)

Incompatible Feature or Device	Corrective Action
USB and sound devices.	Remove these devices from the virtual machine.
N_Port ID Virtualization (NPIV).	Disable the NPIV configuration of the virtual machine.
NIC passthrough.	This feature is not supported by Fault Tolerance so it must be turned off.
Network interfaces for legacy network hardware.	While some legacy drivers are not supported, Fault Tolerance does support the VMXNET2 driver. You might need to install VMware tools to access the VMXNET2 driver instead of vlane in certain guest operating systems.
Virtual disks backed with thin-provisioned storage or thick-provisioned disks that do not have clustering features enabled.	When you turn on Fault Tolerance, the conversion to the appropriate disk format is performed by default. The virtual machine must be in a powered-off state to take this action.
Hot-plugging devices.	<p>The hot plug feature is automatically disabled for fault tolerant virtual machines. To hot plug devices (either adding or removing), you must momentarily turn off Fault Tolerance, perform the hot plug, and then turn on Fault Tolerance.</p> <p>NOTE When using Fault Tolerance, changing the settings of a virtual network card while a virtual machine is running is a hot-plug operation, since it requires "unplugging" the network card and then "plugging" it in again. For example, with a virtual network card for a running virtual machine, if you change the network that the virtual NIC is connected to, FT must be turned off first.</p>
Extended Page Tables/Rapid Virtualization Indexing (EPT/RVI).	EPT/RVI is automatically disabled for virtual machines with Fault Tolerance turned on.
VMXNET3 driver.	Remove the VMXNET3 driver and replace it with a supported driver. A few examples of the supported drivers are the e1000 driver, the VMXNET2 driver, and the VMXNET driver.
Paravirtualized SCSI (PVSCSI) adapter.	Remove the PVSCSI adapter and replace it with a supported adapter. A few examples of the supported adapters are the LSILogic adapter and the BusLogic driver.

Preparing Your Cluster and Hosts for Fault Tolerance

To enable VMware Fault Tolerance for your cluster, you must meet the feature's prerequisites and you must perform certain configuration steps on your hosts. After those steps are accomplished and your cluster has been created, you can also check that your configuration complies with the requirements for enabling Fault Tolerance.

The tasks you should complete before attempting to enable Fault Tolerance for your cluster include:

- Enable host certificate checking (if you are upgrading from a previous version of Virtual Infrastructure)
- Configure networking for each host
- Create the VMware HA cluster, add hosts, and check compliance

After your cluster and hosts are prepared for Fault Tolerance, you are ready to turn on Fault Tolerance for your virtual machines. See [“Turn On Fault Tolerance for Virtual Machines,”](#) on page 38.

Enable Host Certificate Checking

Using host certificate checking, you can configure ESX/ESXi hosts to verify each other's identities, thereby helping to ensure a more secure environment. This is required for ESX/ESXi hosts on which fault tolerant virtual machines reside. If you installed VMware vCenter Server version 4.0, this enablement is performed automatically. If you upgraded from a previous version, you must perform the procedure manually. During this procedure, you will be presented with the list of hosts and their certificates for verification. You can verify the host certificate before committing the certificate checking enablement. Hosts not verified in this step must be manually verified and reconnected.

Procedure

- 1 Connect vSphere Client to vCenter Server.
- 2 Select **Administration** and select **vCenter Server Settings**.
The **vCenter Server Settings** window appears.
- 3 Click **SSL Settings** in the left pane.
- 4 Select the **Check host certificates** box.
- 5 Click **OK**.

Configure Networking for Host Machines

On each host that you intend to add to a VMware HA cluster, you must configure two different networking switches so that the host can also support VMware Fault Tolerance.

Prerequisites

Multiple gigabit Network Interface Cards (NICs) are required. For each host supporting Fault Tolerance, you need a total of two VMkernel gigabit NICs: one dedicated to Fault Tolerance logging and one dedicated to VMotion. The VMotion and FT logging NICs must be on different subnets. Additional NICs are recommended for virtual machine and management network traffic.

Procedure

- 1 Connect vSphere Client to vCenter Server.
- 2 In the vCenter Server inventory, select the host and click the **Configuration** tab.
- 3 Select **Networking** under **Hardware**, and click the **Add Networking** link.
The Add Network wizard appears.
- 4 Select **VMkernel** under **Connection Types** and click **Next**.
- 5 Select **Create a virtual switch** and click **Next**.
- 6 Provide a label for the switch, and select either **Use this port group for VMotion** or **Use this port group for fault tolerance logging**.
- 7 Click **Next**.
- 8 Provide an IP address and subnet mask and click **Next**.
- 9 Click **Finish**.

To enable Fault Tolerance for a host, VMware recommends that you complete this procedure twice, once for each port group option to ensure that sufficient bandwidth is available for Fault Tolerance logging. Select one option, finish this procedure, and repeat the procedure a second time, selecting the other port group option.

After you have created both a VMotion and Fault Tolerance logging virtual switch, you should add the host to the cluster and complete any steps needed to turn on Fault Tolerance.

What to do next

To confirm that you successfully enabled both VMotion and Fault Tolerance on the host, view its Summary tab in the vSphere Client. In the General pane, the fields **VMotion Enabled** and **Fault Tolerance Enabled** should show yes.

NOTE If you configure networking to support Fault Tolerance but subsequently disable it, pairs of fault tolerant virtual machines that are already powered on remain so. However, if a failover situation occurs, when the Primary VM is replaced by its Secondary VM a new Secondary VM is not started, causing the new Primary VM to run in a Not Protected state.

Create VMware HA Cluster and Check Compliance

VMware Fault Tolerance is used in the context of a VMware HA cluster. After you have configured networking on each host, create the VMware HA cluster and add the hosts to it. You can check to see if the cluster is configured correctly and complies with the requirements for the successful enablement of Fault Tolerance.

Procedure

- 1 Connect vSphere Client to vCenter Server.
- 2 In the vCenter Server inventory, select the cluster and click the **Profile Compliance** tab.
- 3 Click **Check Compliance Now** to run the compliance tests.

To view the tests that are run, click **Description**.

The results of the compliance test appear at the bottom of the screen. A host is labeled as either Compliant or Noncompliant.

NOTE For a detailed discussion of how to create a VMware HA cluster, see [Chapter 2, “Creating and Using VMware HA Clusters,”](#) on page 13.

Turning On Fault Tolerance for Virtual Machines

After you have taken all of the required steps for enabling VMware Fault Tolerance for your cluster, you can turn on the feature for individual virtual machines.

The option to turn on Fault Tolerance is unavailable (grayed out) if any of these conditions apply:

- The virtual machine resides on a host that does not have a license for the feature.
- The virtual machine resides on a host that is in maintenance mode or standby mode.
- The virtual machine is disconnected or orphaned (its .vmx file cannot be accessed).
- The user does not have permission to turn the feature on.

If the option to turn on Fault Tolerance is available, this task still must be validated and can fail if certain requirements are not met.

Validation Checks for Turning On Fault Tolerance

A number of validation checks are performed on a virtual machine before Fault Tolerance can be turned on.

- SSL certificate checking must be enabled in the vCenter Server settings.
- The host must be in a VMware HA cluster or a mixed VMware HA and DRS cluster.
- The host must have ESX/ESXi 4.0 or greater installed.

- The virtual machine must not have multiple vCPUs.
- The virtual machine must not have snapshots.
- The virtual machine must not be a template.
- The virtual machine must not have VMware HA disabled.

A number of additional validation checks are performed for powered-on virtual machines (or those being powered on).

- The BIOS of the hosts where the fault tolerant virtual machines reside must have Hardware Virtualization (HV) enabled.
- The host that supports the Primary VM must have a processor that supports Fault Tolerance.
- The host that supports the Secondary VM must have a processor that supports Fault Tolerance and is the same CPU family or model as the host that supports the Primary VM.
- The combination of the virtual machine's guest operating system and processor must be supported by Fault Tolerance (for example, 32-bit Solaris on AMD-based processors is not currently supported).
- The configuration of the virtual machine must be valid for use with Fault Tolerance (for example, it must not contain any unsupported devices).

When your effort to turn on Fault Tolerance for a virtual machine passes the validation checks, the Secondary VM is created and the entire state of the Primary VM is copied. The placement and immediate status of the Secondary VM depends upon whether the Primary VM was powered-on or powered-off when you turned on Fault Tolerance.

If the Primary VM is powered on:

- The Secondary VM is created, placed on a separate compatible host, and powered on if it passes admission control.
- The Fault Tolerance Status displayed on the virtual machine's Summary tab in the vSphere Client is **Protected**.

If the Primary VM is powered off:

- The Secondary VM is immediately created and registered to a host in the cluster (it might be re-registered to a more appropriate host when it is powered on.)
- The Secondary VM is not powered on until after the Primary VM is powered on.
- The Fault Tolerance Status displayed on the virtual machine's Summary tab in the vSphere Client is **Not Protected, VM not Running**.
- When you attempt to power on the Primary VM after Fault Tolerance has been turned on, the additional validation checks listed above are performed. To power on properly, the virtual machine must not use paravirtualization (VMI).

After these checks are passed, the Primary and Secondary VMs are powered on, placed on separate, compatible hosts and the Fault Tolerance Status displayed on the virtual machine's Summary tab in the vSphere Client is **Protected**.

Turn On Fault Tolerance for Virtual Machines

You can turn on VMware Fault Tolerance through the vSphere Client.

NOTE When Fault Tolerance is turned on, vCenter Server unsets the virtual machine's memory limit and sets the memory reservation to the memory size of the virtual machine. While Fault Tolerance remains turned on, you cannot change the memory reservation, size, limit, or shares. When Fault Tolerance is turned off, any parameters that were changed are not reverted to their original values.

Connect vSphere Client to vCenter Server using an account with cluster administrator permissions.

Procedure

- 1 Select the Hosts & Clusters view.
- 2 Right-click a virtual machine and select **Fault Tolerance > Turn On Fault Tolerance**.

The specified virtual machine is designated as a Primary VM and a Secondary VM is established on another host. The Primary VM is now fault tolerant.

Viewing Information About Fault Tolerant Virtual Machines

You can view fault tolerant virtual machines in the vCenter Server inventory using the vSphere Client.

NOTE You cannot disable Fault Tolerance from the Secondary VM.

A VMware Fault Tolerance section (pane) is provided in the **Summary** tab for the Primary VM and includes information about the virtual machine.

- Fault Tolerance Status** Indicates the Fault Tolerance status of the virtual machine.
- Protected. Indicates that the Primary and Secondary VMs are powered on and running as expected.
 - Not Protected. Indicates that the Secondary VM is not running. Possible reasons are listed in the table.

Table 3-2. Reasons for Primary VM Not Protected Status

Reason for Not Protected Status	Description
Starting	Fault Tolerance is in the process of starting the Secondary VM. This message is only visible for a short period of time.
Need Secondary VM	The Primary VM is running without a Secondary VM, so the Primary VM is currently not protected. This generally occurs when there is no compatible host in the cluster available for the Secondary VM. Correct this by bringing a compatible host online. If there is a compatible host online in the cluster, further investigation might be required. Under certain circumstances, disabling Fault Tolerance and then re-enabling it corrects this problem.
Disabled	Fault Tolerance is currently disabled (no Secondary VM is running). This happens when Fault Tolerance is disabled by the user or when vCenter Server disables Fault Tolerance after being unable to power on the Secondary VM.
VM not Running	Fault Tolerance is enabled but the virtual machine is powered off. Power on the virtual machine to reach Protected state.

- Secondary location** Displays the ESX/ESXi host on which the Secondary VM is hosted.

Total Secondary CPU	Indicates the CPU usage of the Secondary VM, displayed in MHz.
Total Secondary Memory	Indicates the memory usage of the Secondary VM, displayed in MB.
vLockstep Interval	The time interval (displayed in seconds) needed for the Secondary VM to match the current execution state of the Primary VM. Typically, this interval is less than one-half of one second.
Log Bandwidth	The amount of network capacity being used for sending VMware Fault Tolerance log information from the host running the Primary VM to the host running the Secondary VM.

Fault Tolerance Best Practices

To help ensure optimal Fault Tolerance results, VMware recommends that you follow certain best practices.

Host Configuration

Observe the following best practices when configuring your hosts.

- Hosts running the Primary and Secondary VMs should operate at approximately the same processor frequencies, otherwise the Secondary VM might be restarted more frequently. Platform power management features which do not adjust based on workload (for example, power capping and enforced low frequency modes to save power) can cause processor frequencies to vary greatly. If Secondary VMs are being restarted on a regular basis, disable all power management modes on the hosts running fault tolerant virtual machines or ensure that all hosts are running in the same power management modes.
- Apply the same instruction set extension configuration (enabled or disabled) to all hosts. The process for enabling or disabling instruction sets varies among BIOSes. See the documentation for your hosts' BIOSes for details on how to configure instruction sets.

Homogeneous Clusters

VMware Fault Tolerance can function in clusters with non-uniform hosts, but it works best in clusters with compatible nodes. When constructing your cluster, all hosts should have the following:

- Processors from the same compatible processor group.
- Common access to datastores used by the virtual machines.
- The same virtual machine network configuration.
- The same ESX/ESXi version.
- The same BIOS settings for all hosts.

Run **Check Compliance** to identify incompatibilities and correct them.

Performance

To increase the bandwidth available for the logging traffic between Primary and Secondary VMs use a 10Gbit NIC rather than 1Gbit NIC, and enable the use of jumbo frames.

Store ISOs on Shared Storage for Continuous Access

ISOs that are accessed by virtual machines with Fault Tolerance enabled should be stored on shared storage that is accessible to both instances of the fault tolerant virtual machine. If this configuration is used, the CD-ROM in the virtual machine continues operating normally, even if there is a failover.

For virtual machines with Fault Tolerance enabled, you might use ISO images that are accessible only to the Primary VM. In such a case, the Primary VM is able to access the ISO, but if a failover occurs, the CD-ROM reports errors as if there is no media. This situation might be acceptable if the CD-ROM is being used for a temporary, non-critical operation such as an installation.

Upgrade Hosts Used for Fault Tolerance

When you upgrade hosts that contain fault tolerant virtual machines, ensure that the Primary and Secondary VMs continue to run on hosts with the same ESX/ESXi version and patch level.

Prerequisites

Cluster administrator privileges.

Sets of four or more ESX/ESXi hosts that are hosting fault tolerant virtual machines which are powered on. If the virtual machines are powered off, the Primary and Secondary VMs can be relocated to hosts with different builds.

NOTE This upgrade procedure is for a minimum four-node cluster. The same instructions can be followed for a smaller cluster, though the unprotected interval will be slightly longer.

Procedure

- 1 Using VMotion, migrate the fault tolerant virtual machines off of two hosts.
- 2 Upgrade the two evacuated hosts to the same ESX/ESXi build.
- 3 Disable Fault Tolerance on the Primary VM.
- 4 Using VMotion, move the disabled Primary VM to one of the upgraded hosts.
- 5 Re-enable Fault Tolerance on the Primary VM that was moved.
- 6 Repeat [Step 1](#) to [Step 5](#) for as many fault tolerant virtual machine pairs as can be accommodated on the upgraded hosts.
- 7 Using VMotion, redistribute the fault tolerant virtual machines.

All ESX/ESXi hosts in a cluster are upgraded.

VMware Fault Tolerance Configuration Recommendations

VMware recommends that you observe certain guidelines when configuring Fault Tolerance.

- In addition to non-fault tolerant virtual machines, you should have no more than four fault tolerant virtual machines (primaries or secondaries) on any single host. The number of fault tolerant virtual machines that you can safely run on each host is based on the sizes and workloads of the ESX/ESXi host and virtual machines, all of which can vary.
- If you are using NFS to access shared storage, use dedicated NAS hardware with at least a 1Gbit NIC to obtain the network performance required for Fault Tolerance to work properly.
- Ensure that a resource pool containing fault tolerant virtual machines has excess memory above the memory size of the virtual machines. Fault tolerant virtual machines use their full memory reservation. Without this excess in the resource pool, there might not be any memory available to use as overhead memory.
- VMware recommends that you use a maximum of 16 virtual disks per fault tolerant virtual machine.
- To ensure redundancy and maximum Fault Tolerance protection, VMware recommends that you have a minimum of three hosts in the cluster. In a failover situation, this provides a host that can accommodate the new Secondary VM that is created.

Troubleshooting Fault Tolerance

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 topics.

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

You can use the information provided in the appendix *Fault Tolerance Error Messages* to help you troubleshoot Fault Tolerance. The topic 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.

Unexpected Virtual Machine Failovers

You might need to troubleshoot VMware Fault Tolerance by determining the reason for unexpected virtual machine failovers. This type of failover is when your Primary or Secondary VM has failed over and redundancy is reestablished, even though its ESX/ESXi host has not crashed. In such cases, virtual machine execution is not interrupted, but redundancy is temporarily lost.

Partial Hardware Failure Related to Storage

This problem can arise when access to storage is slow or completely 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 issues.

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.

Other Fault Tolerance Troubleshooting Issues

You might need to troubleshoot issues that are adversely affecting the functioning of your fault tolerant virtual machines.

Hardware Virtualization Must Be Enabled

When attempting to power on a virtual machine with VMware Fault Tolerance enabled, an error message might appear. This is often the result of Hardware Virtualization (HV) not being available on the ESX/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 ESX/ESXi server hardware or because HV is not enabled in the BIOS.

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

Compatible Secondary Hosts Must Be Available

After powering on a virtual machine with Fault Tolerance enabled, an error message might appear in the Recent Task Pane:

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

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. 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. 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 frequently must slow down to allow the Secondary VM to keep up, effectively reducing its execution speed to the slower speed of the Secondary VM.

Further 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.

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.

Very Large Virtual Machines Can Prevent Use of Fault Tolerance

Enabling Fault Tolerance or migrating a running fault tolerant virtual machine using VMotion can fail if the virtual machine is too large (greater than 15GB) or if memory is changing at a rate faster than VMotion can copy over the network. 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).

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. This is because 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. When the Primary VM is idle, this relative difference between the Primary and Secondary VMs might seem large, but examining the actual CPU usage shows that very little CPU resource is being consumed by the Primary VM or the Secondary VM.

Appendix: Fault Tolerance Error Messages

You might encounter error messages when trying to use VMware Fault Tolerance (FT). The table lists some of these error messages. For each error message there is a description and information about resolving the error, if applicable.

Table A-1. Fault Tolerance Error Messages

Error Message	Description and Solution
This host contains virtual machines (VMs) with Fault Tolerance turned On; therefore, this host cannot be moved out of its current cluster. To move the host to another cluster, first migrate the VMs with Fault Tolerance turned On to a different host	This host cannot be moved out of the cluster because it contains virtual machines with FT turned on. To move the host to another cluster, first migrate the fault tolerant virtual machines to a different host.
Cannot add a host with virtual machines that have Fault Tolerance turned On to a non-HA enabled cluster	FT requires the cluster to be enabled for VMware HA. Edit your cluster settings and turn on VMware HA.
Cannot add a host with virtual machines that have Fault Tolerance turned On as a stand-alone host	FT cannot be enabled on a stand-alone host. While the host is in the VMware HA-enabled cluster, right-click each virtual machine on the host and select Turn Off Fault Tolerance. Once FT is disabled, the host can be made into a stand-alone host.
Fault Tolerance is enabled on one or more VMs on this host and must be disabled to move the host out of the current cluster	This host cannot be moved out of the cluster until FT is turned off. To turn off FT, right-click the fault tolerant virtual machines and select Turn Off Fault Tolerance.
Fault Tolerance is enabled on VM {vmName}. Disable Fault Tolerance to move the VM from the current [Resource pool, Cluster]	To move the virtual machine to another cluster or to a standalone host, first turn off FT.
The host {hostName} has VMs with Fault Tolerance turned On. Before disconnecting the host, the host should be put into maintenance mode or turn Off Fault Tolerance protection on these VMs	This host cannot be disconnected until it is placed in maintenance mode or until FT is turned off. To turn off FT, right-click the fault tolerant virtual machines and select Turn Off Fault Tolerance.
Virtual machines in the same Fault Tolerance pair cannot be on the same host	You have attempted to VMotion a Secondary VM to the same host a Primary VM is on. A Primary VM and its Secondary VM cannot reside on the same host. Select a different destination host for the Secondary VM.

Table A-1. Fault Tolerance Error Messages (Continued)

Error Message	Description and Solution
The unused disk blocks of the virtual machine's disks have not been scrubbed on the file system. This is needed to support features like Fault Tolerance	You have attempted to turn on FT on a powered-on virtual machine which has thick formatted disks with the property of being lazy-zeroed. FT cannot be enabled on such a virtual machine while it is powered on. Power off the virtual machine, then turn on FT and power the virtual machine back on. This changes the disk format of the virtual machine when it is powered back on. Turning on FT could take some time to complete if the virtual disk is large.
The disk blocks of the virtual machine's disks have not been fully provisioned on the file system. This is needed to support features like Fault Tolerance	You have attempted to turn on FT on a powered-on virtual machine with thin provisioned disks. FT cannot be enabled on such a virtual machine while it is powered on. Power off the virtual machine, then turn on FT and power the virtual machine back on. This changes the disk format of the virtual machine when it is powered back on. Turning on FT could take some time to complete if the virtual disk is large.
Unsupported virtual machine configuration for Fault Tolerance	The virtual machine has a virtual device that does not support FT. The specific reason for the incompatibility (for example, multiple vCPUs) is specified in the sub-fault of this message. This error also occurs when you attempt to reconfigure a fault tolerant virtual machine with an unsupported operation, for example, extend disk.
There are configuration issues for the Fault Tolerance operation. Refer to the errors and warnings list for details	There are FT operation issues. To troubleshoot this issue, in the vSphere Client select the failed FT operation in either the Recent Tasks pane or the Tasks & Events tab and click the View details link that appears in the Details column.
This operation is not supported on a Secondary VM of a Fault Tolerant pair	An unsupported operation was performed directly on the Secondary VM. Typically this operation would come from an API. FT does not allow direct interaction with the Secondary VM (except for relocating or migrating it to a different host). Most operations must be performed on the Primary VM.
The Secondary VM with instanceUuid '{instanceUuid}' has already been enabled	An attempt was made to enable FT for a virtual machine on which FT was already enabled. Typically, such an operation would come from an API.
The Secondary VM with instanceUuid '{instanceUuid}' has already been disabled	An attempt was made to disable FT for a Secondary VM on which FT was already disabled. Typically, such an operation would come from an API.
Cannot power On the Fault Tolerance Secondary VM for virtual machine {vmName}. Refer to the errors list for details	An attempt to power on the Secondary VM failed. To troubleshoot this issue, in the vSphere Client select the failed FT operation in either the Recent Tasks pane or the Tasks & Events tab and click the View details link that appears in the Details column.
Host {hostName} does not support virtual machines with Fault Tolerance turned on. This VMware product does not support Fault Tolerance	The product you are using is not compatible with Fault Tolerance. To use the product you must turn Fault Tolerance off. This error message primarily appears when vCenter Server is managing a host with an earlier version of ESX/ESXi or if you are using VMware Server.
Host {hostName} does not support virtual machines with Fault Tolerance turned on. This product supports Fault Tolerance, but the host processor does not	This hosts' processor does not support Fault Tolerance. Use a host with supported hardware to use FT. See the VMware knowledge base article at http://kb.vmware.com/kb/1008027 for information on supported processors.
Host {hostName} has some Fault Tolerance issues for virtual machine {vmName}. Refer to the errors list for details	vCenter Server has detected FT issues on the host. To troubleshoot this issue, in the vSphere Client select the failed FT operation in either the Recent Tasks pane or the Tasks & Events tab and click the View details link that appears in the Details column.
No suitable host can be found to place the Fault Tolerance Secondary VM for virtual machine {vmName}	FT requires that the hosts for the Primary and Secondary VMs use the same CPU model or family and have the same ESX/ESXi host version and patch level. Enable FT on a virtual machine registered to a host with a matching CPU model or family within the cluster. If no such hosts exist, you must add one.

Table A-1. Fault Tolerance Error Messages (Continued)

Error Message	Description and Solution
Operation to power On the Fault Tolerance Secondary VM for {vmName} could not be completed within {timeout} seconds	The attempt to start the Secondary VM by copying the state of the Primary VM failed with a timeout. Default timeout is 300 seconds. Determine what is preventing the Secondary VM from powering on. Check if the FT logging NIC on the Primary VM's host and those tried for the Secondary VM is being shared with other network traffic. You can reduce traffic on the logging NIC of the Primary and Secondary VMs by moving virtual machines with high network traffic to another host.
The Fault Tolerance Secondary VM was not powered On because the Fault Tolerance Primary VM could not be powered On	The Secondary VM was not powered on due to a failure to power on the Primary VM. This error displays when the vSphere Client is used to attempt to power on a Primary VM or if an SDK client invokes the <code>vim.Datacenter.PowerOnVM()</code> API. You must address the issue that prevented the Primary VM from powering on because vCenter Server attempts to power on the Secondary VM only after it has powered on the Primary VM.
DRS Disabled is the only supported DRS behavior for Fault Tolerance virtual machine {vmName}	An SDK client attempted to set a DRS automation level override for a Primary or Secondary VM. vCenter Server blocks all such attempts to change the DRS automation level of fault tolerant virtual machines.
Host CPU is incompatible with the virtual machine's requirements mismatch detected for these features: CPU does not match	FT requires that the hosts for the Primary and Secondary VMs use the same CPU model, family, and stepping. Enable FT on a virtual machine registered to a host with a matching CPU model, family, and stepping within the cluster. If no such hosts exist, you must add one. This error also occurs when you attempt to migrate a fault tolerant virtual machine to a different host.
Record/Replay is not supported for Guest OS XP/PRO on this CPU	This error occurs when you attempt to power on an FT virtual machine that does not meet all of the configuration requirements for FT. See "Turning On Fault Tolerance for Virtual Machines," on page 37.
The Fault Tolerance configuration of the entity {entityName} has an issue: HA is not enabled on the virtual machine	This virtual machine is on a host that is not in a VMware HA cluster or it has had VMware HA disabled. Fault Tolerance requires VMware HA.
The Fault Tolerance configuration of the entity {entityName} has an issue: Secondary VM already exists	The Primary VM already has a Secondary VM. Do not attempt to create multiple Secondary VMs for the same Primary VM.
The Fault Tolerance configuration of the entity {entityName} has an issue: Template virtual machine	FT cannot be enabled on virtual machines which are templates. Use a non-template virtual machine for FT.
The Fault Tolerance configuration of the entity {entityName} has an issue: Virtual machine with multiple virtual CPUs	FT is only supported on virtual machines with a single vCPU configured. Use a single vCPU virtual machine for FT.
The Fault Tolerance configuration of the entity {entityName} has an issue: Host is inactive	You must enable FT on an active host. An inactive host is one that is disconnected, in maintenance mode, or in standby mode.
The Fault Tolerance configuration of the entity {entityName} has an issue: Fault Tolerance not supported by host hardware	FT is only supported on specific processors and BIOS settings with Hardware Virtualization (HV) enabled. To resolve this issue, use hosts with supported CPU models and BIOS settings.
The Fault Tolerance configuration of the entity {entityName} has an issue: Fault Tolerance not supported by VMware Server 2.0	Upgrade to VMware ESX or ESXi 4.0 or later.
The Fault Tolerance configuration of the entity {entityName} has an issue: No VMotion license or no virtual NIC configured for VMotion	Verify that you have correctly configured networking on the host. See "Configure Networking for Host Machines," on page 36. If it is, then you might need to acquire a VMotion license.

Table A-1. Fault Tolerance Error Messages (Continued)

Error Message	Description and Solution
The Fault Tolerance configuration of the entity {entityName} has an issue: No virtual NIC configured for Fault Tolerance logging	An FT logging NIC has not been configured. See “Configure Networking for Host Machines,” on page 36 for instructions.
The Fault Tolerance configuration of the entity {entityName} has an issue: Check host certificates flag not set for vCenter Server	The "check host certificates" box is not checked in the SSL settings for vCenter Server. You must check that box. See “Enable Host Certificate Checking,” on page 36.
The Fault Tolerance configuration of the entity {entityName} has an issue: The virtual machine has one or more snapshots	FT does not support virtual machines with snapshots. Enable FT on a virtual machine without snapshots or use the snapshot manager to delete all snapshots associated with this virtual machine.
The Fault Tolerance configuration of the entity {entityName} has an issue: No configuration information for the virtual machine	vCenter Server has no information about the configuration of the virtual machine. Determine if it is misconfigured. You can try removing the virtual machine from the inventory and re-registering it.
The Fault Tolerance configuration of the entity {entityName} has an issue: Record and replay functionality not supported by the virtual machine	Upgrade the hardware the virtual machine is running on and then turn on FT. Potential configuration issues include: <ul style="list-style-type: none"> ■ Software virtualization with FT is unsupported. ■ FT is not supported for SMP virtual machines. ■ Paravirtualization (VMI) with FT is not supported. ■ VM has device that is not supported with FT. ■ Combination of guest operating system, CPU type and configuration options is incompatible with FT. See “Fault Tolerance Interoperability,” on page 34 for more details about these requirements.
The Fault Tolerance configuration of the entity {entityName} has an issue: The virtual machine's current configuration does not support Fault Tolerance	This error occurs when you attempt to turn on FT for a powered-on virtual machine that does not meet all of the configuration requirements for FT. Power off the virtual machine, address the configuration issue, then Turn On Fault Tolerance. Potential configuration issues include: <ul style="list-style-type: none"> ■ Software virtualization with FT is unsupported. ■ FT is not supported for SMP virtual machines. ■ Paravirtualization (VMI) with FT is not supported. ■ VM has device that is not supported with FT. ■ Combination of guest operating system, CPU type and configuration options is incompatible with FT. See “Fault Tolerance Interoperability,” on page 34 for more details about these requirements.
The virtual machine has {numCpu} virtual CPUs and is not supported for reason: Fault Tolerance	This error occurs when you attempt to reconfigure a Primary VM with more than one vCPU. You must modify the number of vCPUs to one.
The file backing ({backingFilename}) for device Virtual Floppy is not supported for Fault Tolerance	FT is not supported on a virtual machine with a virtual floppy device that has file backing not accessible to the host upon which the Secondary VM resides. To turn on FT for this virtual machine, first remove the unsupported device.
The file backing ({backingFilename}) for device Virtual CDROM is not supported for Fault Tolerance	FT is not supported on a virtual machine with a virtual CDROM device that has file backing not accessible to the host upon which the Secondary VM resides. To turn on FT for this virtual machine, first remove the unsupported device.
The file backing ({backingFilename}) for device Virtual serial port is not supported for Fault Tolerance	FT is not supported on a virtual machine with a virtual serial port device that has file backing not accessible to the host upon which the Secondary VM resides. To turn on FT for this virtual machine, first remove the unsupported device.
The file backing ({backingFilename}) for device Virtual parallel port is not supported for Fault Tolerance	FT is not supported on a virtual machine with a virtual parallel port device that has file backing not accessible to the host upon which the Secondary VM resides. To turn on FT for this virtual machine, first remove the unsupported device.

Table A-1. Fault Tolerance Error Messages (Continued)

Error Message	Description and Solution
The file backing ({backingFilename}) for device Virtual disk is not supported for Fault Tolerance	FT is not supported on a virtual machine with a physical disk that has file backing not accessible to the host upon which the Secondary VM resides. To turn on FT for this virtual machine, first remove the unsupported device.
vCenter disabled Fault Tolerance on VM {vmName} because the Secondary VM could not be powered on	To diagnose why the Secondary VM could not be powered on, see “Troubleshooting Fault Tolerance,” on page 42.
Starting the Secondary VM {vmName} timed out within {timeout} ms	You might be experiencing network latency that is causing the timeout. See “Troubleshooting Fault Tolerance,” on page 42.
Resynchronizing Primary and Secondary VMs	Fault Tolerance has detected a difference between the Primary and Secondary VMs. This can be caused by transient events which occur due to hardware or software differences between the two hosts. FT has automatically started a new Secondary VM, and no action is required. If you see this message frequently, you should alert support to determine if there is an issue.

NOTE For errors related to CPU compatibility, see the VMware knowledge base article at <http://kb.vmware.com/kb/1008027> for information on supported processors.

Index

A

- admission control
 - enabling **23**
 - policy **23**
 - types **15**
 - VMware HA **15**
- admission control policy
 - choosing **20**
 - Host Failures Cluster Tolerates **15**
 - Percentage of Cluster Resources Reserved **18**
 - Specify a Failover Host **20**
- advanced attributes, VMware HA **25**
- Advanced Runtime Info **15**
- affinity rules **31**
- anti-affinity rules **31**

B

- best practices
 - Fault Tolerance **40**
 - VMware HA clusters **27**
 - VMware HA networking **28**
- business continuity **9**

C

- cluster settings **21**
- cluster validity **27**
- compliance check, Fault Tolerance **37**
- Configured Failover Capacity **15, 18**
- configuring VMware HA advanced options **26**
- creating a VMware HA cluster **21**
- Current Failover Capacity **15, 18**
- Current Failover Host **20**
- customizing VMware HA **25**

D

- das.defaultfailoverhost **25**
- das.failedetectioninterval **25**
- das.failedetectiontime **25, 28**
- das.iostatsInterval **24, 25**
- das.isolationaddress **25, 28**
- das.isolationShutdownTimeout **23, 25**
- das.slotCpuInMHz **15, 25**
- das.slotMemInMB **15, 25**
- das.usedefaultisolationaddress **25**

- das.vmCpuMinMHz **15, 18, 25**
- das.vmMemoryMinMB **25**
- default gateway **28**
- Distributed Power Management (DPM) **13, 15**
- Distributed Resource Scheduler (DRS) and Fault Tolerance **34**
 - Fault Tolerance errors **45**
 - turning on **22**
 - using with VMware HA **13**
- downtime
 - planned **9**
 - unplanned **10**

E

- educational support **7**
- error messages, Fault Tolerance **45**
- events and alarms, setting **27**
- Extended Page Tables (EPT) **34**

F

- failover host **20**
- Fault Tolerance
 - anti-affinity rules **31**
 - best practices **40**
 - compliance check **37**
 - configuration recommendations **41**
 - continuous availability **11**
 - enabling **35**
 - error messages **45**
 - interoperability **34**
 - Log Bandwidth **39**
 - logging **36, 42**
 - networking configuration **36**
 - overview **31**
 - prerequisites **33**
 - restrictions for turning on **37**
 - secondary location **39**
 - Total Secondary CPU **39**
 - Total Secondary Memory **39**
 - troubleshooting **42, 43**
 - turning on **38**
 - use cases **32**
 - validation checks **37**
 - vLockstep Interval **39**
 - vSphere configuration **33**

Fault Tolerance status

- Disabled **39**
- Need Secondary VM **39**
- Starting **39**
- VM not Running **39**

firewall ports **28**

ft.maxSwitchoverSeconds **43**

H

Hardware Virtualization (HV) **33, 37, 43**

host certificate checking **33, 36**

Host Failures Cluster Tolerates **15**

Host Isolation Response setting **23**

Host Monitoring **33**

Host Monitoring feature **22, 28**

hosts

- maintenance mode **13**
- network isolation **13**

I

I/O stats interval **24**

interoperability, Fault Tolerance **34**

iSCSI SAN **33**

ISO images **40**

M

Maximum per-VM resets **24**

minimizing downtime **9**

modifying cluster settings **21**

monitoring VMware HA **27**

N

N_Port ID Virtualization (NPIV) **34**

network isolation address **28**

network labels **28**

networking configuration, Fault Tolerance **36**

NIC teaming **29**

O

On-Demand Fault Tolerance **32**

P

paravirtualization **34**

Percentage of Cluster Resources Reserved **18**

planned downtime **9**

planning a VMware HA cluster **13**

port group names **28**

PortFast **28**

prerequisites, Fault Tolerance **33**

primary hosts in clusters **13**

R

Rapid Virtualization Indexing (RVI) **34**

RDM **33, 34**

resource fragmentation **20**

S

secondary hosts in clusters **13**

slot **15**

slot size calculation **15**

snapshots **34**

Specify a Failover Host **20**

storage

iSCSI **33**

NAS **33, 41**

NFS **33, 41**

Storage VMotion **9, 34**

suspending VMware HA **22**

Symmetric multiprocessor (SMP) **34**

T

technical support **7**

tolerating host failures **15**

transparent failover **11, 31**

troubleshooting Fault Tolerance **42**

turning on VMware HA **22**

U

unplanned downtime **10**

updated information **5**

upgrading hosts with FT virtual machines **41**

use cases, Fault Tolerance **32**

V

validation checks **37**

virtual machine overrides **23, 27**

Virtual Machine Startup and Shutdown
feature **21**

VM Monitoring **24**

VM Monitoring sensitivity **24**

VM Restart Priority setting **23**

VMDK **33**

VMFS **13, 28, 42**

VMware HA

advanced attributes **25**

advantages **10**

cluster settings **21**

customizing **25**

monitoring **27**

recovery from outages **10**

suspending **22**

turning on **22**

- VMware HA cluster
 - admission control **15**
 - best practices **27**
 - creating **21, 37**
 - heterogeneity **20**
 - planning **13**
 - primary hosts **13**
 - secondary hosts **13**
- VMware HA networking
 - best practices **28**
 - path redundancy **29**
- VMware Tools **24**
- VMware vLockstep **11, 31**

