

vSphere Availability

Update 1

ESXi 5.0

vCenter Server 5.0

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-000848-00

vmware®

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

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

The VMware Web site also provides the latest product updates.

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

docfeedback@vmware.com

Copyright © 2009–2012 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

About vSphere Availability	5
1 Business Continuity and Minimizing Downtime	7
Reducing Planned Downtime	7
Preventing Unplanned Downtime	8
vSphere HA Provides Rapid Recovery from Outages	8
vSphere Fault Tolerance Provides Continuous Availability	9
2 Creating and Using vSphere HA Clusters	11
How vSphere HA Works	11
vSphere HA Admission Control	16
vSphere HA Checklist	22
Creating a vSphere HA Cluster	23
Customizing vSphere HA Behavior	27
Best Practices for vSphere HA Clusters	30
3 Providing Fault Tolerance for Virtual Machines	35
How Fault Tolerance Works	36
Using Fault Tolerance with DRS	37
Fault Tolerance Use Cases	37
Fault Tolerance Checklist	38
Fault Tolerance Interoperability	39
Preparing Your Cluster and Hosts for Fault Tolerance	40
Providing Fault Tolerance for Virtual Machines	44
Viewing Information About Fault Tolerant Virtual Machines	46
Best Practices for Fault Tolerance	47
vSphere Fault Tolerance Configuration Recommendations	49
Index	51

About vSphere Availability

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

Intended Audience

This information is for anyone who wants to provide business continuity through the vSphere 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.

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.
- Reduce planned downtime for common maintenance operations.
- Provide automatic recovery 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 7
- [“Preventing Unplanned Downtime,”](#) on page 8
- [“vSphere HA Provides Rapid Recovery from Outages,”](#) on page 8
- [“vSphere Fault Tolerance Provides Continuous Availability,”](#) on page 9

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 vSphere vMotion[®] and Storage vMotion functionality in vSphere makes it possible for organizations to 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 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 availability 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 vSphere HA and Fault Tolerance features can minimize or eliminate unplanned downtime by providing rapid recovery from outages and continuous availability, respectively.

vSphere HA Provides Rapid Recovery from Outages

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

vSphere HA protects application availability in the following ways:

- It protects against a server failure by 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, vSphere HA provides the infrastructure to protect all workloads with the infrastructure:

- You do not need to install special software within the application or virtual machine. All workloads are protected by vSphere HA. After vSphere HA is configured, no actions are required to protect new virtual machines. They are automatically protected.
- You can combine vSphere HA with vSphere Distributed Resource Scheduler (DRS) to protect against failures and to provide load balancing across the hosts within a cluster.

vSphere HA has several advantages over traditional failover solutions:

Minimal setup

After a vSphere 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 vSphere HA, you must have sufficient resources to fail over the number of hosts you want to protect with vSphere 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 restarting nonresponsive virtual machines, it 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, vSphere HA can help recover from that failure.

vSphere Fault Tolerance Provides Continuous Availability

vSphere HA provides a base level of protection for your virtual machines by restarting virtual machines in the event of a host failure. vSphere 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 provides continuous availability by ensuring that the states of the Primary and Secondary VMs are identical at any point in the instruction execution of the virtual machine. This is done using the VMware vLockstep technology on the ESXi host platform. 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) executes the workload.

If either the host running the Primary VM or the host running the Secondary VM fails, an immediate and transparent failover occurs. The functioning ESXi host seamlessly becomes the Primary VM host without losing network connections or in-progress transactions. With transparent failover, there is no data loss and network connections are maintained. After a transparent failover occurs, a new Secondary VM is respawned and redundancy is re-established. The entire process is transparent and fully automated and occurs even if vCenter Server is unavailable.

Creating and Using vSphere HA Clusters

2

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

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

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

This chapter includes the following topics:

- [“How vSphere HA Works,”](#) on page 11
- [“vSphere HA Admission Control,”](#) on page 16
- [“vSphere HA Checklist,”](#) on page 22
- [“Creating a vSphere HA Cluster,”](#) on page 23
- [“Customizing vSphere HA Behavior,”](#) on page 27
- [“Best Practices for vSphere HA Clusters,”](#) on page 30

How vSphere HA Works

vSphere HA provides high availability for virtual machines by pooling the virtual machines 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.

When you create a vSphere HA cluster, a single host is automatically elected as the master host. The master host communicates with vCenter Server and monitors the state of all protected virtual machines and of the slave hosts. Different types of host failures are possible, and the master host must detect and appropriately deal with the failure. The master host must distinguish between a failed host and one that is in a network partition or that has become network isolated. The master host uses datastore heartbeating to determine the type of failure.

Master and Slave Hosts

When you add a host to a vSphere HA cluster, an agent is uploaded to the host and configured to communicate with other agents in the cluster. Each host in the cluster functions as a master host or a slave host.

When vSphere HA is enabled for a cluster, all active hosts (those not in standby or maintenance mode, or not disconnected) participate in an election to choose the cluster's master host. The host that mounts the greatest number of datastores has an advantage in the election. Only one master host exists per cluster and all other hosts are slave hosts. If the master host fails, is shut down, or is removed from the cluster a new election is held.

The master host in a cluster has a number of responsibilities:

- Monitoring the state of slave hosts. If a slave host fails or becomes unreachable, the master host identifies which virtual machines need to be restarted.
- Monitoring the power state of all protected virtual machines. If one virtual machine fails, the master host ensures that it is restarted. Using a local placement engine, the master host also determines where the restart should be done.
- Managing the lists of cluster hosts and protected virtual machines.
- Acting as vCenter Server management interface to the cluster and reporting the cluster health state.

The slave hosts primarily contribute to the cluster by running virtual machines locally, monitoring their runtime states, and reporting state updates to the master host. A master host can also run and monitor virtual machines. Both slave hosts and master hosts implement the VM and Application Monitoring features.

One of the functions performed by the master host is virtual machine protection. When a virtual machine is protected, vSphere HA guarantees that it attempts to power it back on after a failure. A master host commits to protecting a virtual machine when it observes that the power state of the virtual machine changes from powered off to powered on in response to a user action. If a failover occurs, the master host must restart the virtual machines that are protected and for which it is responsible. This responsibility is assigned to the master host that has exclusively locked a system-defined file on the datastore that contains a virtual machine's configuration file.

vCenter Server reports whether a host is a master host or a slave host using a vSphere HA host state. This state is reported on the host's **Summary** tab in the vSphere Client and in the Host List view for a cluster or datacenter, if the HA State column has been enabled. An HA state of "Running (Master)" indicates the host is serving as a vSphere HA master host. A state of "Connected (Slave)" indicates the host is serving as a vSphere HA slave host. Several other states are provided to indicate when an election is underway or an error condition has occurred. The host's **Summary** tab provides a link next to the vSphere HA state of the host that explains the current state. For more information about the error states see the *vSphere Troubleshooting* documentation.

NOTE If you disconnect a host from a cluster, all of the virtual machines registered to that host are unprotected by vSphere HA.

Host Failure Types and Detection

The master host of a vSphere HA cluster is responsible for detecting the failure of slave hosts. Depending on the type of failure detected, the virtual machines running on the hosts might need to be failed over.

In a vSphere HA cluster, three types of host failure are detected:

- A host stops functioning (that is, fails).
- A host becomes network isolated.
- A host loses network connectivity with the master host.

The master host monitors the liveness of the slave hosts in the cluster. This communication is done through the exchange of network heartbeats every second. When the master host stops receiving these heartbeats from a slave host, it checks for host liveness before declaring the host to have failed. The liveness check that the master host performs is to determine whether the slave host is exchanging heartbeats with one of the datastores. See [“Datastore Heartbeating,”](#) on page 14. Also, the master host checks whether the host responds to ICMP pings sent to its management IP addresses.

If a master host is unable to communicate directly with the agent on a slave host, the slave host does not respond to ICMP pings, and the agent is not issuing heartbeats it is considered to have failed. The host's virtual machines are restarted on alternate hosts. If such a slave host is exchanging heartbeats with a datastore, the master host assumes that it is in a network partition or network isolated and so continues to monitor the host and its virtual machines. See [“Network Partitions,”](#) on page 13.

Host network isolation occurs when a host is still running, but it can no longer observe traffic from vSphere HA agents on the management network. If a host stops observing this traffic, it attempts to ping the cluster isolation addresses. If this also fails, the host declares itself as isolated from the network.

The master host monitors the virtual machines that are running on an isolated host and if it observes that they power off, and the master host is responsible for the virtual machines, it restarts them.

NOTE If you ensure that the 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.

Network Partitions

When a management network failure occurs for a vSphere HA cluster, a subset of the cluster's hosts might be unable to communicate over the management network with the other hosts. Multiple partitions can occur in a cluster.

A partitioned cluster leads to degraded virtual machine protection and cluster management functionality. Correct the partitioned cluster as soon as possible.

- **Virtual machine protection.** vCenter Server allows a virtual machine to be powered on, but it is protected only if it is running in the same partition as the master host that is responsible for it. The master host must be communicating with vCenter Server. A master host is responsible for a virtual machine if it has exclusively locked a system-defined file on the datastore that contains the virtual machine's configuration file.
- **Cluster management.** vCenter Server can communicate with only some of the hosts in the cluster, and it can connect to only one master host. As a result, changes in configuration that affect vSphere HA might not take effect until after the partition is resolved. This failure could result in one of the partitions operating under the old configuration, while another uses the new settings.

If a vSphere HA cluster contains pre-ESXi 5.0 hosts and a partition occurs, vSphere HA might incorrectly power on a virtual machine that was powered off by the user or it might fail to restart a virtual machine that failed.

NOTE When a vSphere HA cluster is partitioned, you can add a host only to the partition that vCenter Server is communicating with.

Datastore Heartbeating

When the master host in a vSphere HA cluster can not communicate with a slave host over the management network, the master host uses datastore heartbeating to determine whether the slave host has failed, is in a network partition, or is network isolated. If the slave host has stopped datastore heartbeating, it is considered to have failed and its virtual machines are restarted elsewhere.

vCenter Server selects a preferred set of datastores for heartbeating. This selection is made to maximize the number of hosts that have access to a heartbeating datastore and minimize the likelihood that the datastores are backed by the same storage array or NFS server. To replace a selected datastore, use the Cluster Settings dialog box of the vSphere Client to specify the heartbeating datastores. The **Datastore Heartbeating** tab lets you specify alternative datastores. Only datastores mounted by at least two hosts are available. You can also see which datastores vSphere HA has selected for use by viewing the **Heartbeat Datastores** tab of the HA Cluster Status dialog box.

You can use the advanced attribute `das.heartbeatdsperhost` to change the number of heartbeat datastores selected by vCenter Server for each host. The default is two and the maximum valid value is five.

vSphere HA creates a directory at the root of each datastore that is used for both datastore heartbeating and for persisting the set of protected virtual machines. The name of the directory is `.vSphere-HA`. Do not delete or modify the files stored in this directory, because this can have an impact on operations. Because more than one cluster might use a datastore, subdirectories for this directory are created for each cluster. Root owns these directories and files and only root can read and write to them. The disk space used by vSphere HA depends on several factors including which VMFS version is in use and the number of hosts that use the datastore for heartbeating. With vmfs3, the maximum usage is approximately 2GB and the typical usage is approximately 3MB. With vmfs5 the maximum and typical usage is approximately 3MB. vSphere HA use of the datastores adds negligible overhead and has no performance impact on other datastore operations.

vSphere HA Security

vSphere HA is enhanced by several security features.

Select firewall ports opened

vSphere HA uses TCP and UDP port 8182 for agent-to-agent communication. The firewall ports open and close automatically to ensure they are open only when needed.

Configuration files protected using file system permissions

vSphere HA stores configuration information on the local storage or on ramdisk if there is no local datastore. These files are protected using file system permissions and they are accessible only to the root user.

Detailed logging

The location where vSphere HA places log files depends on the version of host.

- For ESXi 5.x hosts, vSphere HA writes to syslog only by default, so logs are placed where syslog is configured to put them. The log file names for vSphere HA are prepended with `fdm`, fault domain manager, which is a service of vSphere HA.
- For legacy ESXi 4.x and 3.x hosts, vSphere HA writes to `/var/log/vmware/fdm` on local disk, as well as syslog if it is configured.
- For legacy ESX 4.x and 3.x hosts, vSphere HA writes to `/var/log/vmware/fdm`.

Secure vSphere HA logins

vSphere HA logs onto the vSphere HA agents using a user account, `vpuser`, created by vCenter Server. This account is the same account used by vCenter Server to manage the host. vCenter Server creates a random password for this account and changes the password periodically. The time period is set by the

vCenter Server `VirtualCenter.VimPasswordExpirationInDays` setting. You can change the setting using the Advanced Settings control in the vSphere Client. Users with administrative privileges on the root folder of the host can also log in to the agent.

Secure communication

All communication between vCenter Server and the vSphere HA agent is done over SSL. Agent-to-agent communication also uses SSL except for election messages, which occur over UDP. Election messages are verified over SSL so that a rogue agent can prevent only the host on which the agent is running from being elected as a master host. In this case, a configuration issue for the cluster is issued so the user is aware of the problem.

Host SSL certificate verification required

vSphere HA requires that each host have a verified SSL certificate. Each host generates a self-signed certificate when it is booted for the first time. This certificate can then be regenerated or replaced with one issued by an authority. If the certificate is replaced, vSphere HA needs to be reconfigured on the host. If a host becomes disconnected from vCenter Server after its certificate is updated and the ESXi or ESX Host agent is restarted, then vSphere HA is automatically reconfigured when the host is reconnected to vCenter Server. If the disconnection does not occur because vCenter Server host SSL certificate verification is disabled at the time, verify the new certificate and reconfigure vSphere HA on the host.

Using vSphere HA and DRS Together

Using vSphere HA with Distributed Resource Scheduler (DRS) combines automatic failover with load balancing. This combination can result in a more balanced cluster after vSphere HA has moved virtual machines to different hosts.

When vSphere 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. vSphere HA uses the virtual machine's CPU and memory reservation to determine if a host has enough spare capacity to accommodate the virtual machine.

In a cluster using DRS and vSphere HA with admission control turned on, virtual machines might not be evacuated from hosts entering maintenance mode. This behavior occurs because of the resources reserved for restarting virtual machines in the event of a failure. You must manually migrate the virtual machines off of the hosts using vMotion.

In some scenarios, vSphere HA might not be able to fail over virtual machines because of resource constraints. This can occur for several reasons.

- HA admission control is disabled and Distributed Power Management (DPM) is enabled. This can result in DPM consolidating virtual machines onto fewer hosts and placing the empty hosts in standby mode leaving insufficient powered-on capacity to perform a failover.
- VM-Host affinity (required) rules might limit the hosts on which certain virtual machines can be placed.
- There might be sufficient aggregate resources but these can be fragmented across multiple hosts so that they can not be used by virtual machines for failover.

In such cases, vSphere HA can use DRS to try to adjust the cluster (for example, by bringing hosts out of standby mode or migrating virtual machines to defragment the cluster resources) so that HA can perform the failovers.

If DPM is in manual mode, you might need to confirm host power-on recommendations. Similarly, if DRS is in manual mode, you might need to confirm migration recommendations.

If you are using VM-Host affinity rules that are required, be aware that these rules cannot be violated. vSphere HA does not perform a failover if doing so would violate such a rule.

For more information about DRS, see the *vSphere Resource Management* documentation.

vSphere 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.

Three types of admission control are available.

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.
vSphere 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 the following:

- 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 vSphere HA admission control can be disabled. However, without it there is no assurance that the expected number of virtual machines can be restarted after a 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 Distributed Power Management (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 vSphere Update Manager).
- If you need to perform testing or maintenance operations.

NOTE When vSphere HA admission control is disabled, vSphere HA ensures that there are at least two powered-on hosts in the cluster even if DPM is enabled and can consolidate all virtual machines onto a single host. This is to ensure that failover is possible.

Host Failures Cluster Tolerates Admission Control Policy

You can configure vSphere HA to tolerate a specified number of host failures. With the Host Failures Cluster Tolerates admission control policy, vSphere 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, vSphere HA performs admission control in the following way:

- 1 Calculates the slot size.

A slot is a logical representation of memory and CPU resources. By default, it is sized to 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.

Slot Size Calculation

Slot size is comprised of two components, CPU and memory.

- vSphere HA calculates 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 32MHz. You can change this value by using the `das.vmcputminmhz` advanced attribute.)
- vSphere HA calculates the memory component by obtaining the memory reservation, plus memory overhead, of each powered-on virtual machine and selecting the largest value. There is no default value for the memory reservation.

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. See [“vSphere HA Advanced Attributes,”](#) on page 28.

Using Slots to Compute the Current Failover Capacity

After the slot size is calculated, vSphere 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. The resource data for a host that is used by vSphere HA can be found by using the vSphere Client to connect to the host directly, and then navigating to the **Resource** tab for the host. If all hosts in your cluster are the same, this data can be obtained by dividing the cluster-level figures by the number of hosts. Resources being used for virtualization purposes are not included. Only hosts that are connected, not in maintenance mode, and that have no vSphere 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 smaller number 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 vSphere 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. This is because some virtual machines can take up multiple slots.


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

Example: Admission Control Using Host Failures Cluster Tolerates Policy

The way that slot size is calculated and used with this admission control policy is shown in 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 slots.
- 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).

Percentage of Cluster Resources Reserved Admission Control Policy

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

With the Percentage of Cluster Resources Reserved admission control policy, vSphere HA ensures that a specified percentage of aggregate CPU and memory resources are reserved for failover.

With the Cluster Resources Reserved policy, vSphere HA enforces admission control as follows:

- 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 corresponding Configured Failover Capacity (provided by the user).

If so, admission control disallows the operation.

vSphere HA 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 32MHz CPU is applied.

NOTE The Percentage of Cluster Resources Reserved admission control policy also checks that there are at least two vSphere HA-enabled hosts in the cluster (excluding hosts that are entering maintenance mode). If there is only one vSphere HA-enabled host, an operation is not allowed, even if there is a sufficient percentage of resources available. The reason for this extra check is that vSphere HA cannot perform failover if there is only a single host in the cluster.

Computing the Current Failover Capacity

The total resource requirements for the powered-on virtual machines is comprised of two components, CPU and memory. vSphere 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 32MHz (this value can be changed using the `das.vmcputminmhz` 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 adding 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 vSphere 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 used with this admission control policy is 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 Failover Hosts Admission Control Policy

You can configure vSphere HA to designate specific hosts as the failover hosts.

With the Specify Failover Hosts admission control policy, when a host fails, vSphere HA attempts to restart its virtual machines on one of the specified failover hosts. If this is not possible, for example the failover hosts have failed or have insufficient resources, then vSphere HA attempts to restart those virtual machines on other hosts in the cluster.

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

NOTE If you use the Specify Failover Hosts admission control policy and designate multiple failover hosts, DRS does not load balance failover hosts and VM-VM affinity rules are not supported.

The Current Failover Hosts appear in the vSphere HA section of the cluster's **Summary** tab in the vSphere Client. The status icon next to each host can be green, yellow, or red.

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

Choosing an Admission Control Policy

You should choose a vSphere 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 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 Failover Hosts policy, resources are not fragmented because hosts are 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 as a number of hosts. The Percentage of Cluster Resources policy allows you to designate up to 100% of cluster CPU or memory resources for failover. The Specify Failover Hosts policy allows you to specify a set of failover hosts.

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.

NOTE vSphere 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.

vSphere HA Checklist

The vSphere HA checklist contains requirements that you need to be aware of before creating and using a vSphere HA cluster.

Requirements for a vSphere HA Cluster

Review this list before setting up a vSphere HA cluster. For more information, follow the appropriate cross reference or see [“Creating a vSphere HA Cluster,”](#) on page 23.

- All hosts must be licensed for vSphere HA.

NOTE ESX/ESXi 3.5 hosts are supported by vSphere HA but must include a patch to address an issue involving file locks. For ESX 3.5 hosts, you must apply the patch ESX350-201012401-SG, while for ESXi 3.5 you must apply the patch ESXi350-201012401-I-BG. Prerequisite patches need to be applied before applying these patches.

- You need at least two hosts in the cluster.
- All hosts need to be configured with static IP addresses. If you are using DHCP, you must ensure that the address for each host persists across reboots.
- There should be at least one management network in common among all hosts and best practice is to have at least two. Management networks differ depending on the version of host you are using.
 - ESX hosts - service console network.
 - ESXi hosts earlier than version 4.0 - VMkernel network.
 - ESXi hosts version 4.0 and later ESXi hosts - VMkernel network with the **Management Network** checkbox enabled.

See [“Best Practices for Networking,”](#) on page 32.

- To ensure that any virtual machine can run on any host in the cluster, all hosts should have access to the same virtual machine networks and datastores. Similarly, virtual machines must be located on shared, not local, storage otherwise they cannot be failed over in the case of a host failure.

NOTE vSphere HA uses datastore heartbeating to distinguish between partitioned, isolated, and failed hosts. Accordingly, you must ensure that datastores reserved for vSphere HA are readily available at all times.

- For VM Monitoring to work, VMware tools must be installed. See [“VM and Application Monitoring,”](#) on page 26.
- Host certificate checking should be enabled. See [“Enable Host Certificate Checking,”](#) on page 41.
- vSphere HA supports both IPv4 and IPv6. A cluster that mixes the use of both of these protocol versions, however, is more likely to result in a network partition.

Creating a vSphere HA Cluster

vSphere HA operates in the context of a cluster of ESXi (or legacy 3.5, 4.0, and 4.1 ESX) hosts. You must create a cluster, populate it with hosts, and configure vSphere HA settings before failover protection can be established.

When you create a vSphere HA cluster, you must configure a number of settings that determine how the feature works. Before you do this, identify your cluster's nodes. These nodes are the ESXi hosts that will provide the resources to support virtual machines and that vSphere HA will use for failover protection. You should then 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 vSphere HA.

You can enable and configure vSphere 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 vSphere HA cluster. Automatic startup is not supported when used with vSphere HA.

Create a vSphere HA Cluster

Your cluster can be enabled for vSphere HA. A vSphere 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 vSphere HA settings.

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

Prerequisites

Verify that all virtual machines and their configuration files reside on shared storage. Verify that the hosts are configured to access that shared storage so that you can power on the virtual machines using different hosts in the cluster,

Verify that hosts are configured to have access to the virtual machine network.

NOTE VMware recommends redundant management network connections for vSphere HA. For information about setting up network redundancy, see [“Network Path Redundancy,”](#) on page 33.

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 vSphere 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 vSphere HA (and other) settings for the cluster.

- 7 On the Cluster Features page, select **Turn On vSphere HA**.
- 8 Configure the vSphere HA settings as appropriate for your cluster.

- Host Monitoring Status
- Admission Control
- Virtual Machine Options
- VM Monitoring
- Datastore Heartbeating

- 9 Click **OK** to close the cluster's Settings dialog box.

You have a configured vSphere HA cluster, populated with hosts, available.

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 vSphere 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 vSphere HA to enable vSphere Fault Tolerance on any virtual machine in the cluster.
Turn On vSphere 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

After you create a cluster, Host Monitoring enables the vSphere HA master host to respond to host or virtual machine failures and management network isolation.

If **Enable Host Monitoring** is selected, each 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 vSphere Fault Tolerance recovery process to work properly.

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

Enabling or Disabling Admission Control

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

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

Enable: Disallow VM power on operations that 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.

Disable: Allow VM power on operations that violate availability constraints

Disables admission control. Virtual machines can, for example, be powered on even if that causes insufficient failover capacity. When you do this, no warnings are presented, and the cluster does not turn red. If a cluster has insufficient failover capacity, vSphere HA can still perform failovers and it uses the VM Restart Priority setting to determine which virtual machines to power on first.

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

- Host failures the cluster tolerates
- Percentage of cluster resources reserved as failover spare capacity
- Specify failover hosts

NOTE See [“Choosing an Admission Control Policy,”](#) on page 21 for more information about how vSphere HA admission control works.

Virtual Machine Options

Default virtual machine settings control the order in which virtual machines are restarted (VM restart priority) and how vSphere HA responds if hosts lose network connectivity with other hosts (host isolation response.)

These settings apply to all virtual machines in the cluster in the case of a host failure or isolation. You can also configure exceptions for specific virtual machines. See [“Customize vSphere HA Behavior for an Individual Virtual Machine,”](#) on page 30.

VM Restart Priority Setting

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 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.

The values for this setting are: Disabled, Low, Medium (the default), and High. If you select Disabled, vSphere HA is disabled for the virtual machine, which means that it is not restarted on other ESXi hosts if its host fails. The Disabled setting 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 setting 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 Setting

Host isolation response determines what happens when a host in a vSphere HA cluster loses its management network connections but continues to run. Host isolation responses require that Host Monitoring Status is enabled. If Host Monitoring Status is disabled, host isolation responses are also suspended. A host determines that it is isolated when it is unable to communicate with the agents running on the 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 (the default), Power off, and Shut down. 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. Shutting down is better than powering off the virtual machine, which does not flush most recent changes to disk or commit transactions. Virtual machines that are in the process of shutting 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 vSphere 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.

If a host has its isolation response disabled (that is, it leaves virtual machines powered on when isolated) and the host loses access to both the management and storage networks, a "split brain" situation can arise. In this case, the isolated host loses the disk locks and the virtual machines are failed over to another host even though the original instances of the virtual machines remain running on the isolated host. When the host comes out of isolation, there will be two copies of the virtual machines, although the copy on the originally isolated host does not have access to the vmdk files and data corruption is prevented. In the vSphere Client, the virtual machines appear to be flipping back and forth between the two hosts.

To recover from this situation, ESXi generates a question on the virtual machine that has lost the disk locks for when the host comes out of isolation and realizes that it cannot reacquire the disk locks. vSphere HA automatically answers this question and this allows the virtual machine instance that has lost the disk locks to power off, leaving just the instance that has the disk locks.

VM and Application Monitoring

VM Monitoring restarts individual virtual machines if their VMware Tools heartbeats are not received within a set time. Similarly, Application Monitoring can restart a virtual machine if the heartbeats for an application it is running are not received. You can enable these features and configure the sensitivity with which vSphere HA monitors non-responsiveness.

When you 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 and I/O activity from the VMware Tools process running inside the guest. If no heartbeats or I/O activity 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.

Occasionally, virtual machines or applications that are still functioning properly stop sending heartbeats. To avoid unnecessary resets, 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`.

To enable Application Monitoring, you must first obtain the appropriate SDK (or be using an application that supports VMware Application Monitoring) and use it to set up customized heartbeats for the applications you want to monitor. After you have done this, Application Monitoring works much the same way that VM Monitoring does. If the heartbeats for an application are not received for a specified time, its virtual machine is restarted.

You can 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 or application 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.

The default settings for monitoring sensitivity are described in [Table 2-1](#). You can also specify custom values for both monitoring sensitivity and the I/O stats interval by selecting the **Custom** checkbox.

Table 2-1. VM Monitoring Settings

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

After failures are detected, vSphere HA resets virtual machines. The reset ensures 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, vSphere HA makes no further attempts to reset the virtual machines after subsequent failures until after the specified time has elapsed. You can configure the number of resets using the **Maximum per-VM resets** custom setting.

Customizing vSphere HA Behavior

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

Review the advanced settings you can use to optimize the vSphere HA clusters in your environment. Because these attributes affect the functioning of vSphere HA, change them with caution.

Set Advanced vSphere HA Options

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

Prerequisites

A vSphere HA cluster for which to modify settings.

Cluster administrator privileges.

Procedure

- 1 In the cluster's Settings dialog box, select **vSphere 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.

vSphere HA Advanced Attributes

You can set advanced attributes that affect the behavior of your vSphere HA cluster.

Table 2-2. vSphere HA Advanced Attributes

Attribute	Description
das.isolationaddress[...]	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 management 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: das.isolationaddressX, where X = 1-10. Typically you should specify one per management network. Specifying too many addresses makes isolation detection take too long.
das.usedefaultisolationaddress	By default, vSphere HA uses the default gateway of the console network as an isolation address. This attribute specifies whether or not this default is used (true false).
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.vmcputminmhz	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 32MHz.
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.

Table 2-2. vSphere HA Advanced Attributes (Continued)

Attribute	Description
das.ignoreinsufficienthbdatastore	Disables configuration issues created if the host does not have sufficient heartbeat datastores for vSphere HA. Default value is false.
das.heartbeatdsperhost	Changes the number of heartbeat datastores required. Valid values can range from 2-5 and the default is 2.

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

- das.isolationaddress[...]
- das.usedefaultisolationaddress
- das.isolationshutdowntimeout

Options No Longer Supported

In vCenter Server 5.0 a number of advanced configuration options for vSphere HA are no longer supported. The following options are no longer supported.

- das.consoleUser
- das.consoleNode
- das.consolePerm
- das.primaryCount
- das.checkVmStateDelay
- das.trace
- das.traceLevel
- das.traceOutput
- das.preferredPrimaries
- das.disableUWSwapRequirement
- das.sensorPollingFreq
- das.bypassNetCompatCheck
- das.defaultfailoverhost
- das.failureDetectionTime
- das.failureDetectionInterval

If you attempt to set one of the unsupported options, vCenter Server 5.0 reports that the option is not valid. Also, if you upgrade to vCenter Server 5.0 from a previous version with any of these options defined, they are deleted and are no longer in effect.

Customize vSphere HA Behavior for an Individual Virtual Machine

Each virtual machine in a vSphere 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.

Procedure

- 1 Select the cluster and select **Edit Settings** from the right-click menu.
- 2 Select **Virtual Machine Options** under vSphere 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 vSphere 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 vSphere HA Clusters

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

Setting Alarms to Monitor Cluster Changes

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

Several default vSphere HA alarms are available.

- Insufficient failover resources (a cluster alarm)
- Cannot find master (a cluster alarm)
- Failover in progress (a cluster alarm)
- Host HA status (a host alarm)
- VM monitoring error (a virtual machine alarm)
- VM monitoring action (a virtual machine alarm)
- Failover failed (a virtual machine alarm)

NOTE The default alarms include the feature name, vSphere HA.

Monitoring Cluster Validity

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

A cluster enabled for vSphere 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** tab in the vSphere Client displays a list of configuration issues for clusters. The list explains what has caused the cluster to become invalid or overcommitted (yellow).

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

Checking the Operational Status of the Cluster

Configuration issues and other errors can occur for your cluster or its hosts that adversely affect the proper operation of vSphere HA. You can monitor these errors by looking at the Cluster Operational Status screen, which is accessible in the vSphere Client from the vSphere HA section of the cluster's **Summary** tab. Address issues listed here.

Most configuration issues have a matching event that is logged. All vSphere HA events include "vSphere HA" in the description. You can search for this term to find the corresponding events.

vSphere HA and Storage vMotion Interoperability in a Mixed Cluster

In clusters where ESXi 5.0 hosts and ESX/ESXi 4.1 or prior hosts are present and where Storage vMotion is used extensively or Storage DRS is enabled, VMware recommends that you do not deploy vSphere HA. vSphere HA might respond to a host failure by restarting a virtual machine on a host with an ESXi version different from the one on which the virtual machine was running before the failure. A problem can occur if, at the time of failure, the virtual machine was involved in a Storage vMotion action on an ESXi 5.0 host, and vSphere HA restarts the virtual machine on a host with a version prior to ESXi 5.0. While the virtual machine might power on, any subsequent attempts at snapshot operations could corrupt the vdisk state and leave the virtual machine unusable.

Admission Control Best Practices

The following recommendations are best practices for vSphere HA admission control.

- Select the Percentage of Cluster Resources Reserved admission control policy. This policy offers the most flexibility in terms of host and virtual machine sizing. In most cases, a calculation of $1/N$, where N is the number of total nodes in the cluster, yields adequate sparing.
- Ensure that you size all cluster hosts equally. An unbalanced cluster results in excess capacity being reserved to handle failure of the largest possible node.
- Try to keep virtual machine sizing requirements similar across all configured virtual machines. The Host Failures Cluster Tolerates admission control policy uses slot sizes to calculate the amount of capacity needed to reserve for each virtual machine. The slot size is based on the largest reserved memory and CPU needed for any virtual machine. When you mix virtual machines of different CPU and memory requirements, the slot size calculation defaults to the largest possible, which limits consolidation.

Using Auto Deploy with vSphere HA

You can use vSphere HA and Auto Deploy together to improve the availability of your virtual machines. Auto Deploy provisions hosts when they power up and you can also configure it to install the vSphere HA agent on such hosts during the boot process. To have Auto Deploy install the vSphere HA agent, the image profile you assign to the host must include the vmware-fdm VIB. See the Auto Deploy documentation included in vSphere Installation and Setup for details.

Suspending vSphere HA

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

Best Practices for Networking

VMware recommends best practices for the configuration of host NICs and network topology for vSphere HA. Best Practices include recommendations for your ESXi hosts, and 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 because of dropped vSphere HA heartbeats.

- When making changes to the networks that your clustered 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 vSphere 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 ESXi hosts themselves, for example, adding port groups, or removing vSwitches, VMware recommends that in addition to suspending Host Monitoring, you place the hosts on which the changes are being made into maintenance mode. When the host comes out of maintenance mode, it is reconfigured, which causes the network information to be reinspected for the running host. If not put into maintenance mode, the vSphere HA agent runs using the old network configuration information.

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

Networks Used for vSphere HA Communications

To identify which network operations might disrupt the functioning of vSphere HA, you should know which management networks are being used for heart beating and other vSphere HA communications.

- On legacy ESX hosts in the cluster, vSphere HA communications travel over all networks that are designated as service console networks. VMkernel networks are not used by these hosts for vSphere HA communications.
- On ESXi hosts in the cluster, vSphere HA communications, by default, travel over VMkernel networks, except those marked for use with vMotion. If there is only one VMkernel network, vSphere HA shares it with vMotion, if necessary. With ESXi 4.x and ESXi, you must also explicitly enable the **Management traffic** checkbox for vSphere HA to use this network.

NOTE To keep vSphere HA agent management traffic separate from other network traffic, VMware recommends that you configure hosts so vmkNICs used by vSphere HA do not share subnets with vmkNICs used for other purposes. vSphere HA agents send packets using any pNIC that is associated with a given subnet if there is also at least one vmkNIC configured for vSphere HA management traffic. Consequently, to ensure network flow separation, the vmkNICs used by vSphere HA and by other features must be on different subnets.

Network Isolation Addresses

A network isolation address is an IP address that is pinged to determine whether 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. Only one default gateway is specified, regardless of how many management networks have been defined. You should use the `das.isolationaddress[...]` advanced attribute to add isolation addresses for additional networks. See [“vSphere HA Advanced Attributes,”](#) on page 28.

Other Networking Considerations

You should consider other things when configuring the networking that supports your vSphere HA cluster.

- **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.
- **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.
- **Configure the management networks** so that the vSphere HA agent on a host in the cluster can reach the agents on any of the other hosts using one of the management networks. If you do not set up such a configuration, a network partition condition can occur after a master host is elected.

Network Path Redundancy

Network path redundancy between cluster nodes is important for vSphere HA reliability. A single management 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 management network, any failure between the host and the cluster can cause an unnecessary (or false) failover activity. 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 management network level. In most implementations, NIC teaming provides sufficient redundancy, but you can use or add management network redundancy if required. Redundant management 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 management 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 management network, 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 vSphere HA cluster, you must reconfigure vSphere 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 management network connection, which is attached to a separate virtual switch. The primary management network connection is used for network and management purposes. When the secondary management network connection is created, vSphere HA sends heartbeats over both the primary and secondary management network connections. If one path fails, vSphere HA can still send and receive heartbeats over the other path.

Providing Fault Tolerance for Virtual Machines

3

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

Fault Tolerance is built on the ESXi host platform (using the VMware vLockstep technology), 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, and the best practices for its usage.

NOTE You might encounter error messages when trying to use Fault Tolerance. For information on error messages related to Fault Tolerance, see the VMware knowledge base article at <http://kb.vmware.com/kb/1033634>.

This chapter includes the following topics:

- [“How Fault Tolerance Works,”](#) on page 36
- [“Using Fault Tolerance with DRS,”](#) on page 37
- [“Fault Tolerance Use Cases,”](#) on page 37
- [“Fault Tolerance Checklist,”](#) on page 38
- [“Fault Tolerance Interoperability,”](#) on page 39
- [“Preparing Your Cluster and Hosts for Fault Tolerance,”](#) on page 40
- [“Providing Fault Tolerance for Virtual Machines,”](#) on page 44
- [“Viewing Information About Fault Tolerant Virtual Machines,”](#) on page 46
- [“Best Practices for Fault Tolerance,”](#) on page 47
- [“vSphere Fault Tolerance Configuration Recommendations,”](#) on page 49

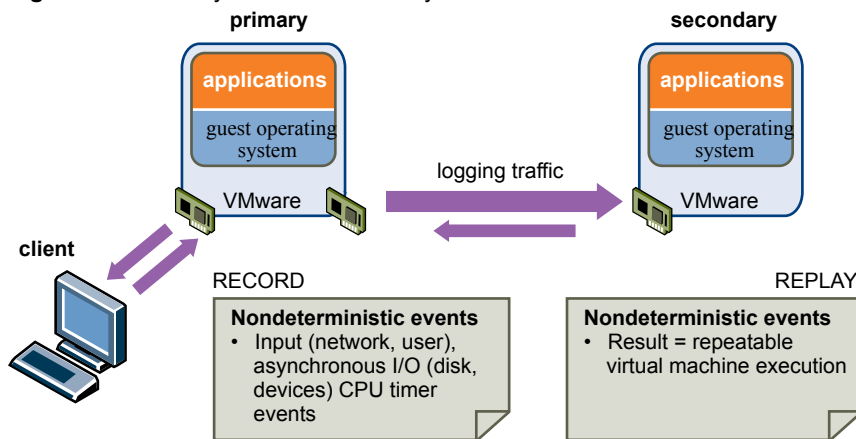
How Fault Tolerance Works

vSphere 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.

NOTE FT logging traffic between Primary and Secondary VMs is unencrypted and contains guest network and storage I/O data, as well as the memory contents of the guest operating system. This traffic can include sensitive data such as passwords in plaintext. To avoid such data being divulged, ensure that this network is secured, especially to avoid 'man-in-the-middle' attacks. For example, you could use a private network for FT logging traffic.

Figure 3-1. Primary VM and Secondary VM in Fault Tolerance Pair



The Primary and Secondary VMs continuously exchange heartbeats. This exchange 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. This restriction ensures that a host failure cannot result in the loss of both virtual machines. You can also use VM-Host affinity rules to dictate which hosts designated virtual machines can run on. If you use these rules, be aware that for any Primary VM that is affected by such a rule, its associated Secondary VM is also affected by that rule. For more information about affinity rules, see the *vSphere Resource Management* documentation.

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.

Using Fault Tolerance with DRS

You can use vSphere Fault Tolerance with vSphere Distributed Resource Scheduler (DRS) when the Enhanced vMotion Compatibility (EVC) feature is enabled. This process allows fault tolerant virtual machines to benefit from better initial placement and also to be included in the cluster's load balancing calculations.

When a cluster has EVC enabled, DRS makes the initial placement recommendations for fault tolerant virtual machines, moves them during cluster load rebalancing, and allows you to assign a DRS automation level to Primary VMs (the Secondary VM always assumes the same setting as its associated Primary VM.)

DRS does not place more than a fixed number of Primary or Secondary VMs on a host during initial placement or load balancing. This limit is controlled by the advanced option `das.maxftvmsperhost`. The default value for this option is 4. However if you set this option to 0, DRS ignores this restriction.

When vSphere Fault Tolerance is used for virtual machines in a cluster that has EVC disabled, the fault tolerant virtual machines are given DRS automation levels of "disabled". In such a cluster, each Primary VM is powered on only on its registered host, its Secondary VM is automatically placed, and neither fault tolerant virtual machine is moved for load balancing purposes.

If you use affinity rules with a pair of fault tolerant virtual machines, a VM-VM affinity rule applies to the Primary VM only, while a VM-Host affinity rule applies to both the Primary VM and its Secondary VM. If a VM-VM affinity rule is set for a Primary VM, DRS attempts to correct any violations that occur after a failover (that is, after the Primary VM effectively moves to a new host).

Fault Tolerance Use Cases

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

Fault Tolerance provides a higher level of business continuity than vSphere 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 vSphere 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 vSphere 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 vSphere 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 Checklist

The following checklist contains cluster, host, and virtual machine requirements that you need to be aware of before using vSphere Fault Tolerance.

Review this list before setting up Fault Tolerance. You can also use the VMware SiteSurvey utility (download at http://www.vmware.com/download/shared_utilities.html) to better understand the configuration issues associated with the cluster, host, and virtual machines being used for vSphere FT.

NOTE The failover of fault tolerant virtual machines is independent of vCenter Server, but you must use vCenter Server to set up your Fault Tolerance clusters.

Cluster Requirements for Fault Tolerance

You must meet the following cluster requirements before you use Fault Tolerance.

- Host certificate checking enabled. See “[Enable Host Certificate Checking](#),” on page 41.
- At least two FT-certified hosts running the same Fault Tolerance version or host build number. The Fault Tolerance version number appears on a host's **Summary** tab in the vSphere Client.

NOTE For legacy hosts prior to ESX/ESXi 4.1, this tab lists the host build number instead. Patches can cause host build numbers to vary between ESX and ESXi installations. To ensure that your legacy hosts are FT compatible, do not mix legacy ESX and ESXi hosts in an FT pair.

- ESXi hosts have access to the same virtual machine datastores and networks. See “[Best Practices for Fault Tolerance](#),” on page 47.
- Fault Tolerance logging and VMotion networking configured. See “[Configure Networking for Host Machines](#),” on page 41.
- vSphere HA cluster created and enabled. See “[Creating a vSphere HA Cluster](#),” on page 23. vSphere HA must be enabled before you can power on fault tolerant virtual machines or add a host to a cluster that already supports fault tolerant virtual machines.

Host Requirements for Fault Tolerance

You must meet the following host requirements before you use Fault Tolerance.

- Hosts must have processors from the FT-compatible processor group. It is also highly recommended that the hosts' processors are compatible with one another. See the VMware knowledge base article at <http://kb.vmware.com/kb/1008027> for information on supported processors.
- Hosts must be licensed for Fault Tolerance.
- Hosts must be certified for Fault Tolerance. See <http://www.vmware.com/resources/compatibility/search.php> and select **Search by Fault Tolerant Compatible Sets** to determine if your hosts are certified.
- The configuration for each host must have Hardware Virtualization (HV) enabled in the BIOS.

To confirm the compatibility of the hosts in the cluster to support Fault Tolerance, you can also run profile compliance checks as described in “[Create vSphere HA Cluster and Check Compliance](#),” on page 43.

NOTE When a host is unable to support Fault Tolerance you can view the reasons for this on the host's **Summary** tab in the vSphere Client. Click the blue caption icon next to the **Host Configured for FT** field to see a list of Fault Tolerance requirements that the host does not meet.

Virtual Machine Requirements for Fault Tolerance

You must meet the following virtual machine requirements before you use Fault Tolerance.

- No unsupported devices attached to the virtual machine. See “[Fault Tolerance Interoperability](#),” on page 39.
- Virtual machines must be stored in virtual RDM or virtual machine disk (VMDK) files that are thick provisioned. If a virtual machine is stored in a VMDK file that is thin provisioned and an attempt is made to enable Fault Tolerance, a message appears indicating that the VMDK file must be converted. To perform the conversion, you must power off the virtual machine.
- Incompatible features must not be running with the fault tolerant virtual machines. See “[Fault Tolerance Interoperability](#),” on page 39.
- Virtual machine files must be stored on shared storage. Acceptable shared storage solutions include Fibre Channel, (hardware and software) iSCSI, NFS, and NAS.
- Only virtual machines with a single vCPU are compatible with Fault Tolerance.
- 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 vSphere Fault Tolerance, you should be aware of the features and products Fault Tolerance cannot interoperate with.

vSphere Features Not Supported with Fault Tolerance

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.
- Linked clones. You cannot enable Fault Tolerance on a virtual machine that is a linked clone, nor can you create a linked clone from an FT-enabled virtual machine.
- Virtual Machine Backups. You cannot back up an FT-enabled virtual machine using Storage API for Data Protection, VMware Data Recovery, or similar backup products that require the use of a virtual machine snapshot, as performed by ESXi. To back up a fault tolerant virtual machine in this manner, you must first disable FT, then re-enable FT after performing the backup. Storage array-based snapshots do not affect FT.

Features and Devices 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 with 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).	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.
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.
vlan networking drivers.	Fault Tolerance does not support virtual machines that are configured with vlane virtual NIC cards. However, vmxnet2, vmxnet3, and e1000 are fully supported.
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. You must power off the virtual machine to trigger this conversion.
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.
Serial or parallel ports	Remove these devices from the virtual machine.
IPv6	Use IPv4 addresses with the FT logging NIC.
Video devices that have 3D enabled.	Fault Tolerance does not support video devices that have 3D enabled.

Preparing Your Cluster and Hosts for Fault Tolerance

To enable vSphere 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 the following:

- Enable host certificate checking (if you are upgrading from a previous version of vCenter Server).
- Configure networking for each host.
- Create the vSphere 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 45.

Enable Host Certificate Checking

Using host certificate checking, you can configure ESXi hosts to verify each other's identities, helping to ensure a more secure environment. Host certificate checking is required for ESXi hosts on which fault tolerant virtual machines reside.

If you installed VMware vCenter Server version 4.1, host certificate checking is enabled 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 **vCenter requires verified host SSL certificates** box.
- 5 Click **OK**.

Configure Networking for Host Machines

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

To enable Fault Tolerance for a host, you must 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.

Prerequisites

Multiple gigabit Network Interface Cards (NICs) are required. For each host supporting Fault Tolerance, you need a minimum of two physical gigabit NICs. For example, you need one dedicated to Fault Tolerance logging and one dedicated to vMotion. VMware recommends three or more NICs to ensure availability.

NOTE The vMotion and FT logging NICs must be on different subnets and IPv6 is not supported on the FT logging NIC.

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.

- 7 Select either **Use this port group for vMotion** or **Use this port group for Fault Tolerance logging** and click **Next**.
- 8 Provide an IP address and subnet mask and click **Next**.
- 9 Click **Finish**.

After you create both a vMotion and Fault Tolerance logging virtual switch, you can create other virtual switches, as needed. You should then 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 **Host Configured for FT** should show yes.

NOTE If you configure networking to support FT but subsequently disable the Fault Tolerance logging port, pairs of fault tolerant virtual machines that are already powered on remain powered on. 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.

Fault Tolerance Host Networking Configuration Example

This example describes the host network configuration for Fault Tolerance in a typical deployment with four 1GB NICs. This is one possible deployment that ensures adequate service to each of the traffic types identified in the example and could be considered a best practice configuration.

Fault Tolerance provides full uptime during the course of a physical host failure due to power outage, system panic, or similar reasons. Network or storage path failures or any other physical server components that do not impact the host running state may not initiate a Fault Tolerance failover to the Secondary VM. Therefore, customers are strongly encouraged to use appropriate redundancy (for example, NIC teaming) to reduce that chance of losing virtual machine connectivity to infrastructure components like networks or storage arrays.

NIC Teaming policies are configured on the vSwitch (vSS) Port Groups (or Distributed Virtual Port Groups for vDS) and govern how the vSwitch will handle and distribute traffic over the physical NICs (vmnics) from virtual machines and vmkernel ports. A unique Port Group is typically used for each traffic type with each traffic type typically assigned to a different VLAN.

Host Networking Configuration Guidelines

The following guidelines allow you to configure your host's networking to support Fault Tolerance with different combinations of traffic types (for example, NFS) and numbers of physical NICs.

- Distribute each NIC team over two physical switches ensuring L2 domain continuity for each VLAN between the two physical switches.
- Use deterministic teaming policies to ensure particular traffic types have an affinity to a particular NIC (active/standby) or set of NICs (for example, originating virtual port-id).
- Where active/standby policies are used, pair traffic types to minimize impact in a failover situation where both traffic types will share a vmnic.
- Where active/standby policies are used, configure all the active adapters for a particular traffic type (for example, FT Logging) to the same physical switch. This minimizes the number of network hops and lessens the possibility of oversubscribing the switch to switch links.

Configuration Example with Four 1Gb NICs

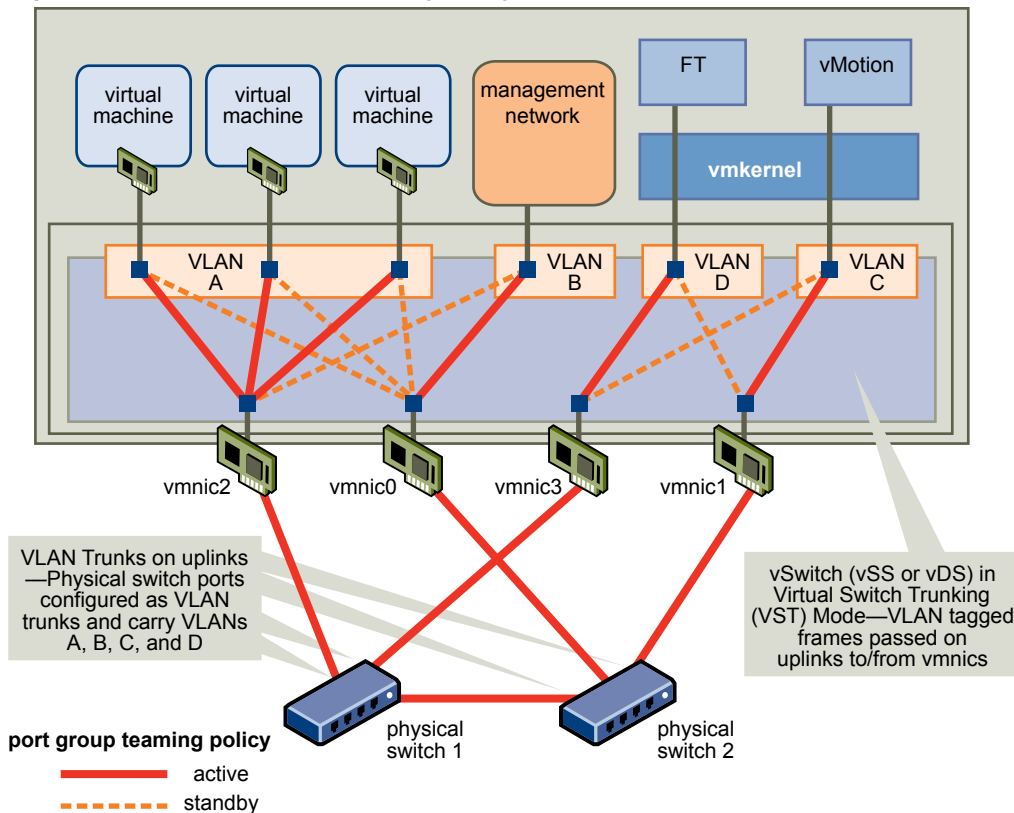
[Figure 3-2](#) depicts the network configuration for a single ESXi host with four 1GB NICs supporting Fault Tolerance. Other hosts in the FT cluster would be configured similarly.

This example uses four port groups configured as follows:

- VLAN A: Virtual Machine Network Port Group-active on vmnic2 (to physical switch #1); standby on vmnic0 (to physical switch #2.)
- VLAN B: Management Network Port Group-active on vmnic0 (to physical switch #2); standby on vmnic2 (to physical switch #1.)
- VLAN C: vMotion Port Group-active on vmnic1 (to physical switch #2); standby on vmnic3 (to physical switch #1.)
- VLAN D: FT Logging Port Group-active on vmnic3 (to physical switch #1); standby on vmnic1 (to physical switch #2.)

vMotion and FT Logging can share the same VLAN (configure the same VLAN number in both port groups), but require their own unique IP addresses residing in different IP subnets. However, separate VLANs might be preferred if Quality of Service (QoS) restrictions are in effect on the physical network with VLAN based QoS. QoS is of particular use where competing traffic comes into play, for example, where multiple physical switch hops are used or when a failover occurs and multiple traffic types compete for network resources.

Figure 3-2. Fault Tolerance Networking Configuration Example



Create vSphere HA Cluster and Check Compliance

vSphere Fault Tolerance is used in the context of a vSphere HA cluster. After you have configured networking on each host, create the vSphere 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 vSphere HA cluster, see [Chapter 2, “Creating and Using vSphere HA Clusters,”](#) on page 11.

Providing Fault Tolerance for Virtual Machines

After you have taken all of the required steps for enabling vSphere Fault Tolerance for your cluster, you can use the feature by turning it on for individual virtual machines.

The option to turn on Fault Tolerance is unavailable (dimmed) 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

Several 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 vSphere HA cluster or a mixed vSphere 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 vSphere HA disabled.
- The virtual machine must not have a video device with 3D enabled.

Several additional validation checks are performed for powered-on virtual machines (or those that are in the process of 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.
- Your hardware should be certified as compatible with Fault Tolerance. To confirm that it is, use the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php> and select **Search by Fault Tolerant Compatible Sets**.

- 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). See the VMware knowledge base article at <http://kb.vmware.com/kb/1008027> for information on supported combinations of processors and guest operating systems.
- 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. 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 entire state of the Primary VM is copied and 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 and placed on separate, compatible hosts. The Fault Tolerance Status that appears on the virtual machine's **Summary** tab in the vSphere Client is tagged **Protected**.

Turn On Fault Tolerance for Virtual Machines

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

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 single virtual machine and select **Fault Tolerance > Turn On Fault Tolerance**.

If you select more than one virtual machine, the **Fault Tolerance** menu is disabled. You must turn Fault Tolerance on for one virtual machine at a time.

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 vSphere Fault Tolerance section (pane) appears 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 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. No state is lost during a failover, regardless of the vLockstep Interval value.
Log Bandwidth	The amount of network capacity being used for sending vSphere Fault Tolerance log information from the host running the Primary VM to the host running the Secondary VM.

For each host configured to support Fault Tolerance, you can view information about its fault tolerant virtual machines by accessing the host's **Summary** tab in the vSphere Client. The **Fault Tolerance** section of this screen displays the total number of Primary and Secondary VMs residing on the host and the number of those virtual machines that are powered on. If the host is ESX/ESXi 4.1 or greater, this section also displays the Fault Tolerance version the host is running. Otherwise, it lists the host build number. For two hosts to be compatible they must have matching FT version numbers or matching host build numbers.

Best Practices for Fault Tolerance

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

In addition to the following information, see the white paper *VMware Fault Tolerance Recommendations and Considerations* at <http://www.vmware.com/resources/techresources/10040>.

Host Configuration

Consider 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 that 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 about how to configure instruction sets.

Homogeneous Clusters

vSphere Fault Tolerance can function in clusters with nonuniform hosts, but it works best in clusters with compatible nodes. When constructing your cluster, all hosts should have the following configuration:

- Processors from the same compatible processor group.
- Common access to datastores used by the virtual machines.
- The same virtual machine network configuration.
- The same ESXi version.
- The same Fault Tolerance version number (or host build number for hosts prior to ESX/ESXi 4.1).
- The same BIOS settings (power management and hyperthreading) for all hosts.

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

Performance

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

Store ISOs on Shared Storage for Continuous Access

Store ISOs that are accessed by virtual machines with Fault Tolerance enabled on shared storage that is accessible to both instances of the fault tolerant virtual machine. If you use this configuration, the CD-ROM in the virtual machine continues operating normally, even when a failover occurs.

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 can 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, noncritical operation such as an installation.

Avoid Network Partitions

A network partition occurs when a vSphere HA cluster has a management network failure that isolates some of the hosts from vCenter Server and from one another. See [“Network Partitions,”](#) on page 13. When a partition occurs, Fault Tolerance protection might be degraded.

In a partitioned vSphere HA cluster using Fault Tolerance, the Primary VM (or its Secondary VM) could end up in a partition managed by a master host that is not responsible for the virtual machine. When a failover is needed, a Secondary VM is restarted only if the Primary VM was in a partition managed by the master host responsible for it.

To ensure that your management network is less likely to have a failure that leads to a network partition, follow the recommendations in [“Best Practices for Networking,”](#) on page 32.

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 FT version number or host build number (for hosts prior to ESX/ESXi 4.1).

Prerequisites

Verify that you have cluster administrator privileges.

Verify that you have sets of four or more ESXi hosts that are hosting fault tolerant virtual machines that 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 ESXi build.
- 3 Turn off Fault Tolerance on the Primary VM.
- 4 Using vMotion, move the disabled Primary VM to one of the upgraded hosts.
- 5 Turn on 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 ESXi hosts in a cluster are upgraded.

vSphere 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 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. The memory reservation of a fault tolerant virtual machine is set to the virtual machine's memory size when Fault Tolerance is turned on. Without this excess in the resource pool, there might not be any memory available to use as overhead memory.
- Use a maximum of 16 virtual disks per fault tolerant virtual machine.
- To ensure redundancy and maximum Fault Tolerance protection, you should 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.

Index

A

- admission control
 - enabling **25**
 - policy **25**
 - types **16**
 - vSphere HA **16**
- admission control policy
 - choosing **21**
 - Host Failures Cluster Tolerates **16**
 - Percentage of Cluster Resources Reserved **19**
 - Specify Failover Hosts **20**
- advanced attributes, vSphere HA **27**
- Advanced Runtime Info **16**
- affinity rules **36, 37**
- anti-affinity rules **36**
- Application Monitoring **12, 26**
- Auto Deploy **30**

B

- best practices
 - Fault Tolerance **47**
 - vSphere HA clusters **30**
 - vSphere HA networking **32**
- business continuity **7**

C

- Cluster Operational Status **30**
- cluster settings **23**
- cluster validity **30**
- compliance check, Fault Tolerance **43**
- Configured Failover Capacity **16, 19**
- configuring vSphere HA advanced options **27**
- creating a vSphere HA cluster **23**
- Current Failover Capacity **16, 19**
- Current Failover Hosts **20**
- customizing vSphere HA **27**

D

- das.heartbeatdsperhost **14, 28**
- das.ignoreinsufficienthbdatastore **28**
- das.iostatsinterval **26, 28**
- das.isolationaddress **28, 32**
- das.isolationshutdowntimeout **25, 28**
- das.maxftvmsperhost **37**

- das.slotcpuinmhz **16, 28**
- das.slotmeminmb **16, 28**
- das.usedefaultisolationaddress **28**
- das.vmcipuminmhz **16, 19, 28**
- das.vmmemoryminmb **28**
- datastore heartbeating **12, 14**
- default gateway **32**
- Distributed Power Management (DPM) **15, 16**
- Distributed Resource Scheduler (DRS) and Fault Tolerance **39**
 - turning on **24**
 - using with vSphere Fault Tolerance **37**
 - using with vSphere HA **15**
- DNS lookup **22**
- downtime
 - planned **7**
 - unplanned **8**

E

- Enhanced vMotion Compatibility **37**
- error messages
 - Fault Tolerance **35**
 - vSphere HA **11**
- EVC **37**
- events and alarms, setting **30**
- Extended Page Tables (EPT) **39**

F

- failover hosts **20**
- Fault Tolerance
 - anti-affinity rules **36**
 - best practices **47**
 - checklist **38**
 - compliance check **43**
 - configuration recommendations **49**
 - continuous availability **9**
 - enabling **40**
 - error messages **35**
 - interoperability **39**
 - Log Bandwidth **46**
 - logging **41, 42**
 - networking configuration **41, 42**
 - overview **36**
 - prerequisites **38**
 - restrictions for turning on **44**

- secondary location **46**
- Total Secondary CPU **46**
- Total Secondary Memory **46**
- turning on **45**
- use cases **37**
- validation checks **44**
- version **38**
- vLockstep Interval **46**
- vSphere configuration **38**
- Fault Tolerance status
 - Disabled **46**
 - Need Secondary VM **46**
 - Starting **46**
 - VM not Running **46**
- firewall ports **14, 32**
- FT logging **36**

H

- Hardware Virtualization (HV) **38, 44**
- host certificate checking **38, 41**
- Host Failures Cluster Tolerates **16, 30**
- Host Isolation Response setting **25**
- Host Monitoring feature **24, 32**
- hosts
 - maintenance mode **12, 15**
 - network isolation **12**

I

- I/O stats interval **26**
- intended audience **5**
- interoperability, Fault Tolerance **39**
- IPv4 **22, 39**
- IPv6 **22, 39, 41**
- iSCSI SAN **38**
- ISO images **47**

L

- load balancing **37**
- log files **14**

M

- management network **22, 32**
- master host election **12**
- Maximum per-VM resets **26**
- minimizing downtime **7**
- modifying cluster settings **23**
- monitoring sensitivity **26**
- monitoring vSphere HA **30**

N

- N_Port ID Virtualization (NPIV) **39**
- network isolation address **32**

- network labels **32**
- network partition **12–14, 47**
- networking configuration, Fault Tolerance **41, 42**
- NIC teaming **33, 42**

O

- On-Demand Fault Tolerance **37**

P

- paravirtualization **39**
- Percentage of Cluster Resources Reserved **19, 30**
- planned downtime **7**
- planning a vSphere HA cluster **11**
- port group names **32**
- PortFast **32**
- prerequisites, Fault Tolerance **38**

R

- Rapid Virtualization Indexing (RVI) **39**
- RDM **38, 39**
- resource fragmentation **21**

S

- slot **16**
- slot size calculation **16**
- snapshots **39**
- Specify Failover Hosts **20**
- SSL Certificates **14**
- storage
 - iSCSI **38**
 - NAS **38, 49**
 - NFS **38, 49**
- Storage DRS **30**
- Storage vMotion **7, 30, 39**
- suspending vSphere HA **24**
- Symmetric multiprocessor (SMP) **39**

T

- TCP port **14**
- tolerating host failures **16**
- transparent failover **9, 36**
- turning on vSphere HA **24**

U

- UDP port **14**
- unplanned downtime **8**
- upgrading hosts with FT virtual machines **48**
- use cases, Fault Tolerance **37**

V

- validation checks **44**

- virtual machine overrides **25, 30**
- virtual machine protection **12, 13**
- Virtual Machine Startup and Shutdown
 - feature **23**
- VLAN **42**
- VM Monitoring **12, 26**
- VM Restart Priority setting **25**
- VM-VM affinity rules **20**
- VMDK **38**
- VMFS **14, 32**
- VMware Tools **26**
- VMware vLockstep **9, 35, 36**
- vpxuser user account **14**
- vSphere HA
 - advanced attributes **27**
 - advantages **8**
 - checklist **22**
 - cluster settings **23**
 - customizing **27**
 - error messages **11**
 - monitoring **30**
 - recovery from outages **8**
 - suspending **24, 30**
 - turning on **24**
- vSphere HA architecture **11**
- vSphere HA cluster
 - admission control **16**
 - best practices **30**
 - creating **23, 43**
 - heterogeneity **21**
 - master host **12, 13**
 - planning **11**
 - slave host **12**
- vSphere HA networking
 - best practices **32**
 - path redundancy **33**

