



VMware Infrastructure

Automating High Availability (HA) Services with VMware HA

VMware® Infrastructure 3 is the first full infrastructure virtualization suite to empower enterprises and small businesses alike to transform, manage, and optimize their IT infrastructure through virtualization. VMware Infrastructure 3 delivers comprehensive virtualization, management, resource optimization, application availability and operational automation capabilities in an integrated offering. VMware HA, a new capability in VMware Infrastructure 3, helps customers improve service levels for any application by implementing cost-effective virtualization-based high-availability solutions that are both easy to use and easy to configure.

This white paper provides an architectural and conceptual overview of VMware HA and describes how you can use HA to provide high availability for any applications running in virtual machines at lower cost than would be possible with static, physical infrastructure. Using VMware HA, virtual machines are automatically restarted in the event of hardware failure without investing in costly one-to-one mapping of production and backup hardware.

This white paper covers the following topics:

- [Introduction to VMware Infrastructure and VMware HA](#)
- [VMware HA Architecture and Conceptual Overview](#)
- [Using VMware HA](#)
- [VMware HA Requirements and Best Practices](#)

This white paper is intended for VMware partners, resellers, and VMware customers who want to implement virtual infrastructure solutions and want to know how to use distributed infrastructure services such as VMware HA.

Introduction to VMware Infrastructure and VMware HA

With the introduction of VMware Infrastructure 3, VMware extends the evolution of virtual infrastructure and virtual machines that began with VMware ESX Server v1.0. VMware Infrastructure 3 also introduces a revolutionary new set of infrastructure-wide services for resource optimization, high availability, and data protection that deliver capabilities which previously required complex or expensive solutions to implement using only physical machines.



Use of these services provides significantly higher hardware utilization and better alignment of IT resources with business goals and priorities.

VMware Infrastructure introduces two new concepts:

- **Clusters** that aggregate and manage the combined resources of multiple hosts as a single collection.
- **Resource pools** that simplify control over the resources of a host or a cluster.

VMware Infrastructure virtualizes and aggregates industry-standard servers (processors, memory, their attached network and storage capacity) into logical resource pools (from a single ESX Server host or from a VMware cluster) that can be allocated to virtual machines on demand. Resource pools can also be nested and organized hierarchically so that the IT environment matches company organization. Individual business units can receive dedicated infrastructure while still profiting from the efficiency of resource pooling.

A set of virtualization-based distributed infrastructure services provide virtual machine monitoring and management to automate and simplify provisioning, optimize resource allocation, and provide operating system and application-independent high availability to applications at lower cost and without the complexity of solutions used with static, physical infrastructure. One of these distributed services, VMware HA, provides easy-to-use, cost-effective high availability for all applications running in virtual machines. In the event of server hardware failure, affected virtual machines are automatically restarted on other physical servers that have spare capacity. HA minimizes downtime and IT service disruption while eliminating the need for dedicated stand-by hardware and installation of additional software.

VMware HA provides uniform high availability across the entire virtualized IT environment without the cost and complexity of failover solutions tied to either operating systems or specific applications.

VMware HA Architecture and Conceptual Overview

Before discussing the specific details of how VMware HA works and how to use it to provide high availability, it's helpful to review a few basics about VMware Infrastructure and describe some of the key elements with which VMware distributed services such as VMware HA interact. The following sections provide basic information on VMware Infrastructure 3 architecture and components.

VMware Infrastructure

At the core of VMware Infrastructure, VMware ESX Server is the foundation for delivering virtualization-based distributed services to IT environments. ESX Server provides a robust virtualization layer that abstracts processor, memory, storage and networking resources into multiple virtual machines that run side-by-side on the same physical server.

ESX Server installs directly on the server hardware, or "bare metal," and inserts a robust virtualization layer between the hardware and the operating system. ESX Server partitions a physical server into multiple secure and portable virtual machines that run on the same physical server. Each virtual machine represents a complete system—with processors, memory, networking, storage and BIOS—so that Windows, Linux, Solaris and NetWare operating systems and software applications run in virtualized machines without any modification.



Another key building block of VMware Infrastructure, VirtualCenter, is used to manage all ESX Server hosts and virtual machines. VirtualCenter Management Server also provides critical services such as:

- Centralized server and virtual machine management
- Virtual machine provisioning
- Performance monitoring
- Operational automation
- Secure access control
- Migration of live virtual machines

Figure 1 shows the architecture and typical configuration of VMware Infrastructure:




Figure 1. VMware Infrastructure Configuration

VMware Infrastructure simplifies management with a single client called the Virtual Infrastructure (VI) Client that you can use to perform all tasks. Every ESX Server configuration task, from configuring storage and network connections, to managing the service console, can be accomplished centrally through the VI Client.

The VI Client connects to ESX Server hosts, even those not under VirtualCenter management, and lets you remotely connect to any virtual machine for console access. There is a Windows version of the VI Client, and for access from any networked device, a web browser application provides virtual machine management and VMware Console access. The browser version of the client, Virtual Infrastructure Web Access, makes it as easy to give a user access to a virtual machine as sending a bookmark URL.

VirtualCenter user access controls provide customizable roles and permissions, so you create your own user roles by selecting from an extensive list of permissions to grant to each role.



Responsibilities for specific VMware Infrastructure components such as resource pools can be delegated based on business organization, or ownership. VirtualCenter also provides full audit tracking to provide a detailed record of every action or operation performed on the virtual infrastructure and who did it

Users can also access virtualization-based distributed services provided by VMotion™, DRS, and HA directly through VirtualCenter and the VI Client. In addition, VirtualCenter exposes a rich programmatic Web Service interface for integration with third party system management products and extension of the core functionality.

- **VMware VMotion** enables the live migration of running virtual machines from one physical server to another. Live migration of virtual machines enables companies to perform hardware maintenance without scheduling downtime and disrupting business operations. VMotion also allows the mapping of virtual machines to hosts to be continuously and automatically optimized within clusters for maximum hardware utilization, flexibility, and availability.
- **VMware DRS** works with VMotion to provide automated resource optimization and virtual machine placement and migration to help align available resources with pre-defined business priorities while maximizing hardware utilization.
- **VMware HA** enables broad-based, cost-effective application availability, independent of specific hardware and operating systems.
- **VMware Consolidated Backup** provides an easy-to-use, centralized facility for LAN-free backup of virtual machines. Full and incremental file-based backup is supported for virtual machines running Microsoft Windows operating systems. Full image backup for disaster recovery scenarios is available for all virtual machines regardless of guest operating system.

VMware Clusters

Clusters, a new concept in virtual infrastructure management, give you the power of multiple hosts with the simplicity of managing a single entity. New cluster support in VMware Infrastructure 3 reduces management complexity by combining standalone hosts into a single cluster with pooled resources and inherently higher availability.

:




Figure 2. Resource Aggregation in VMware Clusters



VMware clusters let you aggregate the hardware resources of individual ESX Server hosts but manage the resources as if they resided on a single host. Now, when you power on a virtual machine, it can be given resources from anywhere in the cluster, rather than be tied to a specific ESX Server host.

VMware Infrastructure 3 provides two services to help with the management of VMware clusters, VMware HA and VMware DRS. VMware HA allows virtual machines running on specific hosts to be automatically restarted using other host resources in the cluster in the case of host machine failures. VMware DRS provides automatic initial virtual machine placement and makes automatic resource relocation and optimization decisions as hosts are added or removed from the cluster or the load on individual virtual machines goes up or down. DRS also makes cluster-wide resource pools possible.

Note: For more information about resource pools and using VMware DRS to manage operations such as virtual machine placement and providing dynamic resource allocation for virtual machines running on VMware cluster hosts, see the VMware Infrastructure 3 white paper titled "Resource Management with VMware DRS."

VMware HA Overview

As described earlier, VMware HA provides easy-to-use, cost effective high availability for all applications running in virtual machines. In the event of server failure, affected virtual machines are automatically restarted on other host machines in the cluster that have spare capacity. HA minimizes downtime and IT service disruption while eliminating the need for dedicated stand-by hardware and installation of additional software. VMware HA provides uniform high availability across the entire virtualized IT environment without the cost and complexity of failover solutions tied to either operating systems or specific applications.

Traditional High Availability and Failover Solutions

Both VMware HA and traditional clustering and high availability solutions support automatic recovery from host failures. They are complementary, but differ somewhat in hardware and software requirements, time to recovery, and the degree to which they incorporate application and operating system awareness.

A traditional clustering solution such as Microsoft Cluster Service (MSCS) or Veritas Cluster Server aims to provide immediate recovery with minimal downtime for applications in case of host or virtual machine failure. To achieve this, the IT infrastructure must be set up as follows:

- Each machine (or virtual machine) must have a mirror virtual machine (potentially on a different host).
- The machine (or the virtual machine and its host) are set up to mirror each other using the clustering software. Generally, the primary virtual machine sends heartbeats to the mirror. In case of failure, the mirror takes over seamlessly.



The following illustration shows the typical host setup for virtual machines using a traditional clustering approach:




Figure 3. Traditional Clustering Configuration

Setup and maintenance of such a clustering solution is expensive and resource intensive. Each time you add a new virtual machine, additional virtual machines and possibly additional hosts are needed for failover. You have to set up, connect, and configure all new machines and update the clustering application's configuration.

To summarize, the traditional solution guarantees fast recovery, but is resource- and labor-intensive in addition to typically also being application and operating system dependent.

Because of the cost and complexity of clustering solutions, they are typically used for a small percentage of enterprise applications, leaving the vast majority of applications without any failover protection whatsoever.

VMware HA "democratizes" high availability by making it available and cost-justifiable for any application.

The VMware HA Solution

With VMware HA, a set of ESX Server hosts is combined into a cluster with a shared pool of resources. VMware HA monitors all hosts in the cluster. If one of the hosts fails, VMware HA immediately responds by restarting each affected virtual machine on a different host.




Figure 4. Host Failover using VMware HA

Using VMware HA has a number of advantages:

- Minimal setup and startup. The New Cluster wizard is used for initial setup. Hosts and new virtual machines can be added using the Virtual Infrastructure Client.



- Reduced hardware cost and setup. In a traditional clustering solution, duplicate hardware and software must be available, and the components must be connected and configured properly. When using VMware HA clusters, you must have sufficient resources to accommodate the number of hosts for which you want to guarantee failover. However, the VirtualCenter Server takes care of all other aspects of the resource management.
- VMware HA "democratizes" high availability by making it available and cost-justifiable for any application, regardless of hardware and operating system platform.

VMware HA is focused on hardware failure, not on operating system or software failure. If you need greater levels and guarantees of availability to handle those situations, you can consider using both VMware HA and traditional high availability approaches together.

VMware HA Features

Using a cluster enabled for VMware HA provides the following features:

- Automatic failover is provided on ESX Server host hardware failure for all running virtual machines within the bounds of failover capacity (see [Designating Failover Capacity](#) below). VMware HA provides automatic detection of server failures and initiates the virtual machine restart without any human intervention.
- VMware HA can take advantage of DRS to provide for dynamic and intelligent resource allocation and optimization of virtual machines after failover. After a host has failed and virtual machines have been restarted on other hosts, DRS can provide further migration recommendations or migrate virtual machines for more optimum host placement and balanced resource allocation.
- VMware HA supports easy-to-use configuration and monitoring using VirtualCenter. HA ensures that capacity is always available (within the limits of specified failover capacity) in order to restart all virtual machines affected by server failure (based on resource reservations configured for the virtual machines.)
- HA continuously monitors capacity utilization and "reserves" spare capacity to be able to restart virtual machines. Virtual Machines can fully utilize spare failover capacity when there hasn't been a failure.

Finally, VMware HA is compatible with traditional application-level failover approaches, so if requirements dictate, you can implement enhanced high availability and failover solutions using both methods.

Clusters and VirtualCenter Failure

You create and manage clusters using VirtualCenter. The VirtualCenter Management Server places an agent on each host in the cluster so each host can communicate with other hosts to maintain state information and know what to do in case of another host's failure. (The VirtualCenter Management Server does not provide a single point of failure.) If the VirtualCenter Management Server host goes down, HA functionality changes as follows. HA clusters can still restart virtual machines on other hosts in case of failure; however, the information about what extra resources are available will be based on the state of the cluster before the VirtualCenter Management Server went down.

Note: If you're also using DRS, the virtual machines running on VMware cluster hosts continue running using available resources. However, there are no further recommendations for resource optimization.



How does VMware HA work?

VMware HA continuously monitors all ESX Server hosts in a cluster and detects failures. An agent placed on each host maintains a "heartbeat" with the other hosts in the cluster and loss of a heartbeat initiates the process of restarting all affected virtual machines on other hosts.

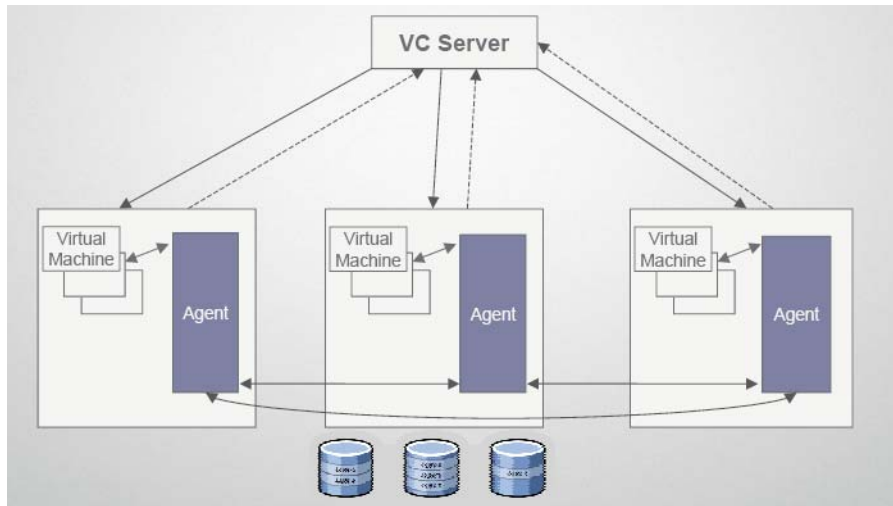


Figure 5. Host Failover using VMware HA

HA monitors whether sufficient resources are available in the cluster at all times in order to be able to restart virtual machines on different physical host machines in the event of host failure. Safe restart of virtual machines is made possible by the locking technology in the ESX Server storage stack, which allows multiple ESX Servers to have access to the same virtual machines file simultaneously.

Designating Failover Capacity

When you enable a cluster for HA, the New Cluster wizard prompts you to specify the maximum number of host failures you want to protect against. This number will be shown as the Configured Failover Capacity in the Virtual Infrastructure Client. VMware HA uses this number to continuously monitor whether there are enough resources to power on virtual machines in the cluster. You need to specify only the number of hosts for which you want failover capacity. VMware HA computes the resources that it requires to fail over virtual machines with the specified failover capacity.

This resource determination is based on the virtual machine's configured CPU and memory resource reservations and capability to handle the failure of the largest host(s) in the cluster. It helps to have more uniform hosts in the cluster, for example, to avoid situations in which virtual machines don't have enough resources to be restarted on new hosts. When the number of host failures exceeds configured spare capacity, virtual machines with the highest priorities are failed over first.

Note: You can choose to allow the cluster to power on virtual machines even when they violate availability constraints; however, this means that failover guarantees may no longer be valid.

Planning HA Clusters

When planning the size of HA clusters to provide the desired levels of failover capacity, keep in mind that each host requires some overhead memory and CPU and that each virtual machine



must be guaranteed its CPU and memory reservation. VMware HA factors in the worst-case failure scenarios when deciding to allow new virtual machines to be powered up. When computing required failover capacity, HA first considers the host with the largest capacity to run virtual machines with the highest resource requirements. HA might therefore be quite conservative in its estimates if the hosts in your cluster have a wide variance in the individual resources they provide.

Using VMware HA

This section describes some of the setup and operation tasks you can perform using HA and VirtualCenter—creating HA clusters, adding or removing hosts from clusters, planning failover capacity, setting properties, and so on.

Enabling HA

VMware HA is included as an integrated component in VMware Infrastructure 3 Enterprise. It is also available as add-on license options to VMware Infrastructure 3 Starter and VMware Infrastructure 3 Standard. To enable HA when you create a VMware cluster, you need to set the **Enable VMware HA** option.

For clusters enabled for HA, the resources of all included hosts are assigned to the cluster. If clusters are also DRS-enabled, you can use DRS to provide dynamic and intelligent resource allocation, optimization, and load-balancing of virtual machines, after failover.

Creating a VMware Cluster

A cluster is a collection of ESX Server hosts and associated virtual machines with shared resources and a shared management interface. When you add a host to a cluster, the host's resources become part of the cluster's resources. When you create a cluster, you can enable it for DRS, HA, or both. If DRS is enabled, the cluster supports shared resource pools and performs placement and dynamic load balancing for virtual machines in the cluster. If HA is enabled, the cluster supports failover. When a host fails, HA will automatically restart virtual machines on a different host. If clusters are enabled for both DRS and HA, DRS will optimize host placement and balanced resource allocation after failover and restart of virtual machines on new hosts.

Your system must also meet certain prerequisites to use VMware cluster features successfully. See [VMware HA Requirements and Best Practices](#), later in this white paper, for more specific requirements and recommendations.

VirtualCenter provides a New Cluster wizard to take you through the steps of creating a new cluster. When you first invoke the wizard, it prompts you to choose whether to create a cluster that supports VMware DRS, VMware HA, or both. Following that, you are prompted for the corresponding configuration information.

Note: When you create a cluster, it initially does not include any hosts or virtual machines.

Using HA and DRS Together

When HA performs failover and restarts virtual machines on different hosts, its first priority is immediate availability of all virtual machines. After the virtual machines have been restarted, those hosts in which they were powered on are usually heavily loaded, while other hosts are comparatively lightly loaded.

Using HA and DRS together combines automatic failover with load balancing. This combination can result in a fast rebalancing of virtual machines after HA has moved virtual machines to



different hosts. You can set up affinity and anti-affinity rules to start two or more virtual machines preferentially on the same host (affinity) or on different hosts.

Note: For more information about resource pools and using VMware DRS to manage operations such as virtual machine placement and providing dynamic resource allocation for virtual machines running on VMware cluster hosts, see the VMware Infrastructure 3 white paper titled "Resource Management with VMware DRS."

Selecting High Availability Options (HA)

If you have enabled HA, the New Cluster Wizard allows you to set the following options.

Option	Description
Host Failures	Specifies the number of host failures (or failure capacity) for which you want to guarantee failover of virtual machines.
Admission Control	Offers two choices about how decisions are made to allow new virtual machines to be powered up: <ul style="list-style-type: none"> • Do not power on virtual machines if they violate availability constraints and enforce the specified failover capacity limits. • Allow virtual machines to be powered on even if they violate availability constraints. This allows you to power on virtual machines even if failover of the number of specified hosts can no longer guaranteed. (A warning is issued.)

After initial creation of the cluster, you can add hosts and virtual machines to the cluster, or specify additional cluster customization such as setting the priority for individual virtual machines. HA uses virtual machine priority to decide order of restart in case of a red cluster (when configured failover capacity exceeds current failover capacity).

Note: If you are using a cluster enabled for HA, that cluster might be marked with a red warning icon until you have added enough hosts to satisfy the specified failover capacity. See [Cluster Status Information](#) later in this paper.

Adding Hosts to a HA Cluster

The VirtualCenter inventory panel displays all clusters and hosts managed by that VirtualCenter Management Server. Adding managed hosts to an HA cluster is as simple as selecting and dragging a host machine to the desired target cluster.

Note: You can also add unmanaged hosts by selecting the Add Host option and specifying the unmanaged host name, user name, and password.

Adding a host to the cluster spawns a system task "Configuring HA on the host." After this task has completed successfully, the host is included in the HA service and virtual machines deployed to the host become part of the cluster.

When a new host is added to a cluster

- The resources for that host immediately becomes available to the cluster for use in the cluster's root resource pool.
- Unless the cluster is also enabled for DRS, all resource pools are collapsed into the cluster's top-level (invisible) resource pool.
- Any capacity on the host beyond what is required or guaranteed for each running virtual machine becomes available as spare capacity in the cluster pool. This spare capacity can be used for starting virtual machines on other hosts in case of a host failure.



- If you add a host with several running virtual machines, and the cluster no longer fulfills its failover requirements because of that addition, a warning appears and the cluster's status is changed to invalid (Red).
- By default, all virtual machines on hosts that are added are given a restart priority of Medium. You can change the priority and specify other HA customization options to tailor individual virtual machine priorities and other settings. (See [Customizing Virtual Machine HA Options](#) for more information.)
- The system also monitors the status of the HA service on each host and displays information about configuration issues on the Summary page.

Viewing Cluster Information

When you select a cluster from the VirtualCenter inventory panel, the Summary page displays high-level information about the selected cluster.

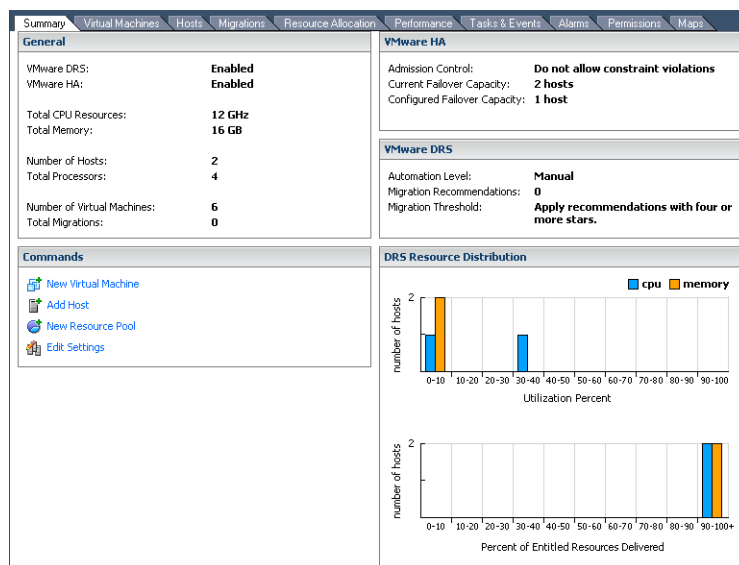


Figure 6. Viewing Cluster Information

The Summary page provides various information about the cluster and virtual machines assigned to the cluster. Regarding HA configuration, the page shows the cluster's admission control setting, current failover capacity, and configured failover capacity for clusters enabled for HA.

Note: For clusters enabled for DRS, the Summary page also displays automation level and migration threshold settings, outstanding migration recommendations and real-time histograms of "Utilization Percent" and "Percent of Entitled Resources Delivered," to show how balanced the cluster is.

The Summary page updates the current failover capacity whenever a host has been added to or removed from the cluster or when virtual machines have been powered on or powered off.

Cluster Status Information

As hosts and virtual machines are added or removed, clusters can become overcommitted or invalid because of HA or DRS violations. Messages displayed on the Summary page displays the status of the currently selected cluster.



The Virtual Infrastructure Client indicates whether a cluster is valid (green), overcommitted (yellow), or invalid (red).

Green (Valid) Clusters

A cluster is considered valid unless something happens that makes it overcommitted or it no longer satisfies failover capacity requirements. For example, an HA cluster becomes invalid if the current failover capacity is lower than the configured failover capacity.

If a cluster is labeled green (valid), this indicates there are enough resources to meet all reservations and to support all running virtual machines. In addition, there is at least one host with enough resources to run each virtual machine assigned to the cluster. If you use a particularly large virtual machine (for example, a virtual machine with a 16GB reservation), you must have at least one host with that much memory. It's not enough if two hosts together fulfill the requirement.

Yellow (Overcommitted) Cluster

There is no yellow HA cluster indication resulting from failover requirements either being met or not met. A cluster becomes yellow when the tree of resource pools and virtual machines is internally consistent but the user has allowed violation of certain requirements. This happens when capacity is removed from the cluster, for example, because a host fails or is removed and there are no longer enough resources to support all requirements.

Red (Invalid) Cluster

A cluster enabled for HA becomes red when the number of virtual machines powered on exceeds the requirements of strict failover, that is, current failover capacity is smaller than the configured failover capacity. This can happen, for example, if you first check **Allow virtual machines to be started even if they violate availability constraints** for that cluster and later power on so many virtual machines that there are no longer sufficient resources to guarantee failover for the specified number of hosts. A cluster can also become red if you power on a virtual machine or perform other operations directly on the host.

A cluster can also become red, for example, if HA is set up for two-host failure in a four-host cluster and one host fails. The remaining three hosts might no longer be able to satisfy a two-host failure.

If a cluster enabled for HA becomes red, it can no longer guarantee failover for the specified number of hosts, but does continue performing failover. In case of host failure, HA first fails over the virtual machines of one host in order of priority, then the virtual machines of the second host in order of priority, and so on.

Customizing Virtual Machine HA Options

Reconfiguring HA can mean turning it off or reconfiguring its options. To turn off HA, you can select the cluster and deselect the **VMware HA** check box from the Edit Settings panel.



To reconfigure VMware HA, for example, to customize HA behavior for individual virtual machines, you can select the cluster and select **HA Services** from the **Edit Settings** -> **Cluster Settings** dialog box.




Figure 7. VMware HA Cluster Settings dialog box

This dialog box lets you make changes to the number of host failovers or the admission control behavior.

You can customize HA for restart priority and isolation response:

- **Restart priority** determines the order in which virtual machines are restarted upon host failure. (Higher priority virtual machines are started first.) Restart priority is always considered, but is especially important in the following cases:
 - If you've set host failure capacity to a certain number of hosts and more than that number of hosts actually fails.
 - If you've turned off strict admission control and have started more virtual machines than HA has been set up to support.

Note: This priority applies only on a per-host basis. If multiple hosts fail, VirtualCenter first migrates all virtual machines from the first host, in order of their priority, then all virtual machines from the second host in order of priority, and so on.

- **Isolation response** determines what a host in a HA cluster should do with running virtual machines when the host loses its network connectivity (not receiving HA heartbeats and unable to ping gateway). By default, virtual machines are powered off in case of a host isolation incident. This releases their shared storage locks, which allows the virtual machines to be started up on other hosts.

You can change this default behavior for individual virtual machines and choose **Leave running** to indicate the virtual machine on isolated hosts should continue running even if the host can no longer communicate with other hosts in the cluster. If you choose to do this and it turns out that the original host can't access shared storage, the virtual machine lock will time out and the virtual machine may be started on a second host (a condition commonly referred to as *split-brain*). This condition is more likely to occur with NAS or iSCSI storage, in the case of network failures, since both methods are TCP/IP based. For these types of storage, keeping the Isolation Response at **Power off** (the default) is highly recommended.

Note: If you add a host to a cluster, all virtual machines in the cluster default to a Restart priority of **Medium** and an Isolation Response of **Power off**.

Powering On Virtual Machines in a Cluster

When you power on a virtual machine on a host that is part of a cluster, the resulting VirtualCenter behavior depends on the type of cluster.



If you power on a virtual machine and HA is enabled, VirtualCenter first checks whether there are enough resources to continue supporting the specified number of host failovers if you power on the virtual machine.

- If there are enough resources, the virtual machine is powered on.
- If there are not enough resources and strict admission control is being used (the default), a message informs you that the virtual machine cannot be powered on. If you are not using strict admission control, a message informs you that there are no longer enough resources to guarantee failover for all hosts. The virtual machine is powered on but the cluster turns red.

Host Removal and Virtual Machines

Both standalone hosts and hosts within a cluster support Maintenance Mode, which restricts virtual machine operations on the host to allow the user to shut down running virtual machines in preparation for host shut down. While in maintenance mode, the host does not allow you to deploy or power on new virtual machines. Virtual machines that are already running on the host continue to run normally. You can either migrate them to another host, or shut them down. When there are no more running virtual machines on the host or cluster, its icon changes and its Summary indicates the new state. In addition, menu and command options involving virtual machine deployment are disabled when this host or cluster is selected.

Because a host must be in maintenance mode before you can remove it from a cluster, all virtual machines must be powered off first (unless DRS is also enabled, in which case virtual machines are automatically removed from the host). When you then remove the host from the cluster, the virtual machines that are currently associated with the host are also removed from the cluster.

Removing Hosts with Virtual Machines from a Cluster

If you remove a host from a cluster, the available resources for the cluster decrease. When you remove a host with virtual machines from a cluster, all its virtual machines are removed as well. You can remove a host only if it is in maintenance mode or disconnected.

Note: If a cluster enabled for HA loses so many resources that it can no longer fulfill its failover requirements, a message appears and the cluster turns red. The cluster will fail over virtual machines in case of host failure, but is not guaranteed to have enough resources available to fail over all virtual machines.

Removing Virtual Machines from a Cluster

You can remove virtual machines by migrating them out of a cluster or removing a host with virtual machines from the cluster.

You can migrate a virtual machine from a cluster to a standalone host, or from a cluster to another cluster, using the standard drag-and-drop method or selecting Migrate from the virtual machine's right-button menu or the VirtualCenter menu bar. If the cluster is also DRS-enabled and the virtual machine is a member of a DRS cluster affinity rules group, VirtualCenter displays a warning before it allows the migration to proceed. The warning indicates that dependent virtual machines are not migrated automatically, so you have to acknowledge the warning before migration can proceed.



VMware HA Requirements and Best Practices

There are a few basic requirements that your virtual infrastructure system and hosts need to meet so that VMware cluster and HA features operate properly.

First, for clusters enabled for VMware HA, all virtual machines and their configuration files must reside on shared storage (Fibre Channel SAN, iSCSI SAN, or SAN iSCSI NAS), because you need to be able to power on the virtual machine on any host in the cluster. This also means that the hosts must be configured to have access to the same virtual machine networks, shared storage, and other resources.

Note: See the *VMware Infrastructure Server Configuration Guide*. Also see the *VMware SAN Configuration Guide* for additional information.

Second, VMware HA monitors heartbeat between hosts on the console network for failure detection. So, to have reliable failure detection for HA clusters, the console network should have redundant network paths. That way, if a host's first network connection fails, the second connection can broadcast heartbeats to other hosts. To set up redundancy, you need two physical network adapters on each host. You then connect them to the corresponding service console, either using two service console interfaces (each on a separate vSwitch or a separate physical network adapter) or using a single interface with NIC teaming.

Last, if you want to use DRS with HA for load balancing, the hosts in your cluster must be part of a vMotion network. If the hosts are not in the vMotion network, however, DRS can still make initial placement recommendations.

Summary

VMware HA, along with new VMware Infrastructure capabilities of clusters and resource pools, and tighter integration with other VMware tools such as VirtualCenter, vMotion, and DRS greatly simplifies virtual machine provisioning, resource allocation, load balancing, and migration, while also providing an easy-to-use, cost-effective high-availability and failover solution for applications running in virtual machines. Using VMware Infrastructure and HA helps to eliminate single points of failure in the deployment of business-critical virtual machine applications, while maintaining other inherent virtualization benefits such as higher system utilization, closer alignment of IT resources with business goals and priorities, and more streamlined, simplified, and automated administration of larger infrastructure installations and systems.