



What's New in VMware vSphere™ 5.0 - Availability

TECHNICAL WHITE PAPER
V 1.0/UPDATED JUNE 2011

Table of Contents

Introduction 3

Management Availability 3

 Manageability 3

 Usability 4

 Application Support 4

Application Availability 5

Infrastructure Availability 5

 vSphere Fault Tolerance 5

 vSphere High Availability 5

 Scalability 5

 Reliability 6

 Usability 7

Summary 8

Introduction

The availability of an environment is not determined by the availability of a single component. As an example, availability of an application is of little use if the infrastructure is not highly available as well. Similarly, if the ability to manage the environment is lost, one cannot truly state that they have a highly available environment.

With this in mind, VMware provides several features that can be leveraged to increase the availability of a virtualized environment. This paper presents these features as they apply to availability of the applications, the infrastructure, and the management platform.

Management Availability

Being the primary management point of a VMware environment, the availability of VMware vCenter™ Server is a critical consideration when constructing a highly available environment. Providing this availability for VMware vCenter Server is the primary function of VMware vCenter Server Heartbeat.

With the latest release, several enhancements have been incorporated within VMware vCenter Server Heartbeat. These feature enhancements are the result of the focus of VMware in three key areas — manageability, usability, and application support.

Manageability

Providing protection for a component as critical as VMware vCenter Server requires a solution that is easy to manage. The easier it is to manage, the less likelihood of errors being introduced into the environment that would decrease the effectiveness of the solution.

VMware vCenter Server Heartbeat 6.4 provides an enhanced architecture that allows both the active and passive servers to be represented as unique entities within Microsoft Active Directory. This allows for the assignment of a unique IP address to both servers, making them accessible through the network at all times. With both servers readily accessible, administrators can easily perform maintenance actions, such as patching or upgrades on either server.

Now the VMware vCenter Server instance can be associated with a virtual IP address. When a failover is required, this virtual IP address follows the active VMware vCenter Server instance. By doing so, users must be aware only of a solitary address to connect to. This greatly reduces confusion.

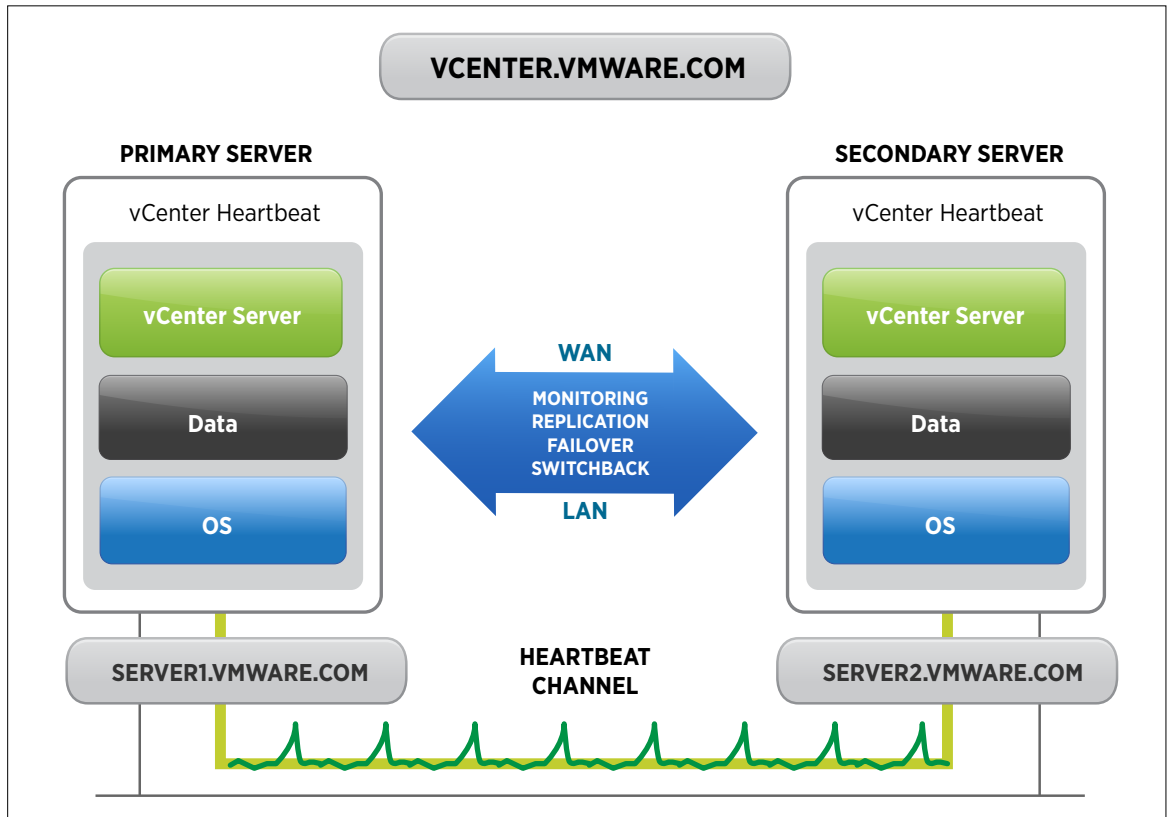


Figure 1. Architecture of VMware vCenter Server Heartbeat

Maintenance can now be performed with minimal downtime for the VMware vCenter Server instance. Patches can be applied to the passive host through its unique IP without affecting the active host. The VMware vCenter Server instance can then be failed over from the active server to the passive server in preparation for patching the active server. In the event that the maintenance activity causes an issue with a host, administrators can simply fail the VMware vCenter Server instance back to the other host until the problem is rectified.

Usability

Being able to effectively use VMware vCenter Server Heartbeat eliminates complexity and increases the efficiency of the user. VMware vSphere™ 5.0 ("vSphere") brings enhancements designed to leverage the VMware vCenter Server client as a single pane of glass for VMware vCenter Server Heartbeat operations. A new vSphere client plug-in provides administrators with the ability to monitor and perform management functions right from the client. VMware vCenter Server Heartbeat operational tasks and alarms are now registered with VMware vCenter Server. This allows users to view this information directly from the VMware vCenter Server client.

Application Support

In addition to providing availability for VMware vCenter Server 5.0, VMware vCenter Server Heartbeat now also provides availability of VMware View Composer and Microsoft SQL Server 2008 R2. This support increases the platform support to match common customer deployment preferences.

Application Availability

Applications drive the need for high availability within an environment. It is critical, then, to monitor the availability of an application. The ability to induce an action based on that monitoring is also critical.

Customers have attempted to provide this ability in a number of ways, such as implementing a clustering product within the guest operating system (OS) of a virtual machine. One of the problems with this method is that these solutions often have no visibility into the infrastructure layer, so they could not trigger the infrastructure to take action when required.

VMware now provides a mechanism for this through the use of the vSphere® High Availability (vSphere HA) application monitoring API, which is included with vSphere 5.0. This API allows a developer to create a monitoring program, running within the guest OS that communicates with vSphere HA using heartbeats. These heartbeats inform vSphere HA that an application is functioning properly. Upon loss of a heartbeat, vSphere HA will restart the virtual machine.

By leveraging this API, customers and vendors can more easily craft solutions that take advantage of vSphere HA and the underlying infrastructure protection it provides.

Infrastructure Availability

The infrastructure itself provides the base level of availability for an environment. The availability requirements for the infrastructure are often determined by the availability requirements of the applications for which the infrastructure provides services. The more critical an application is, the greater the need for availability of the infrastructure.

VMware understands that customers have a varying degree of availability requirements when it comes to their infrastructure. To this end, VMware provides multiple solutions that can be enabled by the customer. These include vSphere® Fault Tolerance (Fault Tolerance) for continuous availability and vSphere HA for high availability.

Fault Tolerance

Fault Tolerance is designed to provide continuous availability for a virtual machine. When enabled, Fault Tolerance creates an exact replica of a virtual machine. It keeps this replica on standby, ready to be brought into use upon a failure of the primary virtual machine.

In vSphere 5.0, the Fault Tolerance feature now encompasses new support for CPU architectures and guest operating systems. Refer to <http://kb.vmware.com/kb/1008027> for the most up-to-date information on the supported CPUs and guest operating systems for Fault Tolerance.

vSphere HA

The most widely used feature for availability in a virtualized environment is vSphere HA. vSphere HA provides the foundation for a highly available environment by monitoring the virtual machines and the hosts upon which they run.

In vSphere 5.0, the vSphere HA feature has been completely rewritten by VMware from the ground up. This was done to increase the scalability, reliability, and usability of vSphere HA. Each of these factors, and the enhancements incorporated to support them, are described below.

Scalability

With the increased usage of VMware products in today's modern datacenter comes a need to provide a scalable solution for high availability. The redesign of vSphere HA provides this foundation.

One of the biggest changes with vSphere HA is that the concept of primary and secondary nodes has been completely removed. The new model incorporates a master-slave relationship between the nodes in a cluster, where one node is elected to be a master and the rest are slaves. The master node coordinates all availability actions with the other nodes and is responsible for communicating that state to the VMware vCenter Server. This model eliminates a significant amount of planning in the architecture design of a highly available environment. No longer must administrators worry about what hosts are their primary nodes and where they are located. This is especially significant when implementing vSphere HA on blade chassis and in stretched cluster environments.

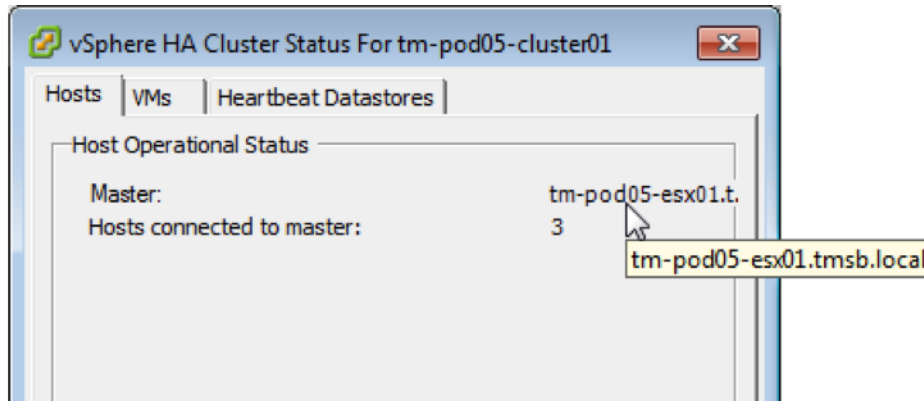


Figure 2. Role of the Nodes Is Easily Identifiable

The support for IPv6 networking in vSphere HA enables IT departments in need of a larger addresses space to fully leverage their network infrastructure.

vSphere HA now also includes an enhanced deployment mechanism. This enhancement allows administrators to complete tasks in a fraction of the time required previously, such as with deployment of the vSphere HA agent, vSphere HA configure, unconfigure, reconfigure, and so on.

Reliability

When a disaster event occurs, the last thing that an administrator wants to worry about is if the solution deployed will work properly. By analysis of the most common support calls made by customers about vSphere HA, VMware added features to ensure continued confidence in vSphere HA.

One enhancement was the elimination of dependencies on any external component by vSphere HA. Specifically, vSphere HA no longer has any type of dependency upon DNS resolution by each host in the cluster. Eliminating this reduces the odds that an outage of an external component will have an effect on the operation of vSphere HA.

Another enhancement is the ability to enable communication between the nodes within a cluster through the storage subsystem. Now vSphere HA will utilize multiple paths of communication, through the network and storage. Not only does this allow for a greater level of redundancy, but it also enables better identification of the health of a node and the virtual machines running on it.

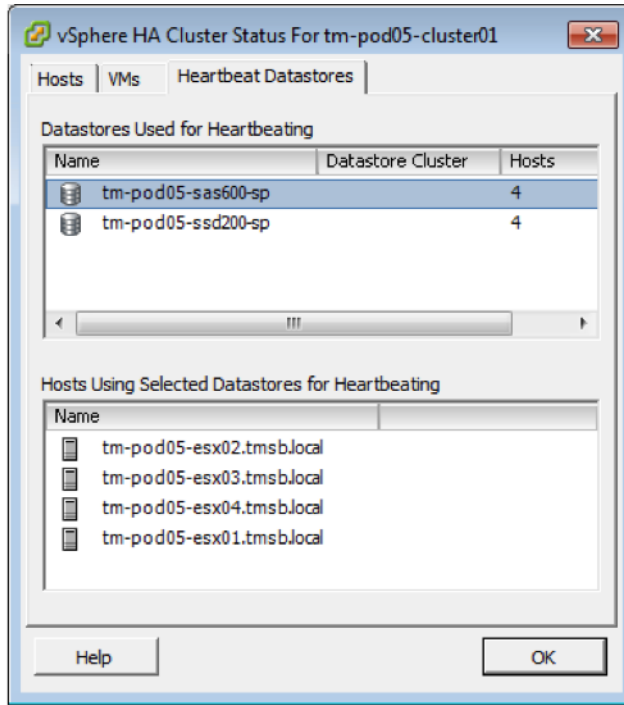


Figure 3. vSphere HA Heartbeat Datastores

Usability

Although most of the enhancements made to vSphere HA are not visible to the end user, there was a focus on improving the usability.

The user interface has been enhanced to allow for users to quickly identify the role a node in the cluster plays, as well as its state. Messages reporting error conditions have also been made easier to understand and act upon. When problems do occur, there is just one log file that must be reviewed, greatly decreasing the time to resolution.



Figure 4. vSphere HA Protection States of the Virtual Machines

Summary

The combination of VMware vCenter Server Heartbeat, the vSphere HA application monitoring API, and vSphere HA, provides holistic support for highly available environments. Used together, they provide a level of sophisticated protection from outages on all levels, including management, application, and infrastructure.

