

VMware vSphere HA for Cisco Prime Collaboration

Introduction to High Availability

High availability (HA) is a system design approach that ensures specific applications or services are accessible at an agreed upon level of performance for a specified period of time. Essentially, different design techniques are put in place to ensure the application or service is up when it needs to be, minimizing downtime (the time the application is unavailable).

HA Layers

Generically speaking, there are multiple layers at which HA can be implemented into a design. These layers may include:

- The application layer: The application under consideration
- The OS layer: Linux server, Windows server, and so on
- The virtualization layer: Virtualization software
- The physical layer: Servers, networking infrastructure, storage infrastructure, and so on

This white paper covers recommendations for enabling HA at the virtualization layer for Cisco Prime™ Collaboration. As such, while we will discuss basic hardware requirements for establishing HA utilizing VMware HA in conjunction with Cisco Prime Collaboration 10.0, discussing HA at other layers is outside the scope of this paper.

Introduction to VMware vSphere HA

vSphere HA provides high availability and reduces downtime for virtualized applications by making use of multiple hosts set up in a cluster configuration. This allows for relatively quick recovery from various types of outages that may occur.

vSphere HA provides for application availability at the virtualization layer by:

- Reacting to hardware failures and network disruptions. When this occurs, it will restart the application on an active host within the cluster.
- Detecting OS failures and monitoring virtual machines (VMs). When this occurs, it will restart the virtual machine itself as required.

Details on vSphere HA can be found here:

<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-availability-guide.pdf>

vSphere HA Requirements

- All hosts must be licensed for vSphere HA
- There must be a minimum of two hosts in the cluster
- All hosts should either be configured with static IP addresses or the addresses must be persistent across reboots if using Dynamic Host Configuration Protocol (DHCP)

-
- There should minimally exist one management network common among all hosts, though it is recommended to have two
 - All hosts should have access to the same virtual machine networks and data stores
 - Virtual machines must be located on shared storage
 - Host certificate checking should be enabled
 - Hosts within a cluster must be either on IPv4 or IPv6, not a mix, or this can result in a network partition

Details on vSphere HA requirements can be found here:

<http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-availability-guide.pdf>

Cisco Prime Collaboration Requirements

There are no special requirements for Cisco Prime Collaboration to enable vSphere HA, nor is any configuration required on Cisco Prime Collaboration.

General VMware requirements for Cisco Prime Collaboration can be found in the Quick Start Guides here:

<http://www.cisco.com/c/en/us/support/cloud-systems-management/prime-collaboration/products-installation-guides-list.html>

The Cisco Prime Collaboration Quick Start Guides include detailed virtual machine requirements for Cisco Prime Collaboration Assurance (including Analytics) and Cisco Prime Collaboration Provisioning.

vSphere HA Configuration

vSphere HA is configured clusterwide. It also comes with a number of options for further fine-tuning and control. Recommendations for Cisco Prime Collaboration are offered below.

Admission Control and Admission Control Policy

These options specify policies based on available cluster capacity when vSphere HA is enabled. For the Admission Control option, Cisco recommends **Enabled**. This will prevent new VMs from turning on in situations in which doing so will exceed the configured maximum host failover capacity, and it will help prevent overcommitment in the event of host failures. Once enabled, three Admission Control Policy options will help determine how much spare capacity is available in the event of host failures.

The first option looks at spare capacity from the perspective of the number of host failures you want to withstand, given your hardware slot size (a metric that sizes your host based on the amount of RAM and CPU capacity). The second option allows you to choose manually the amount of spare RAM and CPU percentages. The third option allows you to specify failover hosts - essentially placing these hosts in maintenance mode until there's a failure.

The most conservative and safest approach is to size failover capacity to enable restarting of all machines (high enough to allow all machines to restart). If this cannot be done, then Cisco recommends an approach that allows for the maximum amount of spare capacity in your setup to tolerate failover.

Virtual Machine Options

Within this area are two configurable items.

The first is VM Restart Priority and the second is Host Isolation Response. You can configure these both at the cluster level by selecting from the drop-down boxes, or you can configure them individually by clicking the specific machine and changing the settings manually.

For VM Restart Priority, your choice here is largely dependent on which application is critical to restart during reduced resource availability. If you desire to restart Cisco Prime Collaboration immediately, particularly if you've enabled the Analytics option, choose "High" as the restart priority.

Please note that if you opt to disable Admission Control, it is possible that failover capacity will be exceeded, and, during a failure, machines with a lower priority that were not on a failing host will continue to run while machines of a higher priority on a failing host will not restart due to a lack of resources.

VM Monitoring

Choose the "VM Monitoring Only" option. This will monitor for guest OS failures, and upon detection will restart the VM. Cisco recommends setting this to "High," which will cause a restart if 30 seconds elapse without a heartbeat response from the guest OS. This can be set either at the cluster level or for individual machines.

Data Store Heartbeating

Data store heartbeating is a new feature in vSphere HA 5.x, providing an additional means of backup from network isolation. Cisco has no specific recommendations as they relate to Cisco Prime Collaboration, and these settings will depend largely on administrator preference.

Testing and Verifying vSphere HA with Cisco Prime Collaboration

Option 1: vMotion Test

One of the easiest ways to test that your setup works is to perform a vMotion migration of the virtual machine from one host to another, as this is the same mechanism used when hosts fail and machines have to be restarted - those machines will be migrated (using vMotion) to the hosts that are up.

Option 2: Power Outage Test - Host Shutdown

A second option is to turn off one host and verify that all of the virtual machines on the host that went off come back online on the remaining or configured failover hosts.

Option 3: Network Isolation

There are various ways of simulating a network isolation event. In vSphere 5.x, you should turn off data store heartbeating in your cluster settings. After this, either disconnecting the patch cables where the physical network interface cards (NICs) are uplinked or shutting down your ports using your switch command-line interface (CLI) will simulate a network isolation event.

The following reference provides more detail on Options 2 and 3 for testing purposes:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2056634

vSphere HA Testing Observations

During HA testing, Cisco observed that VMs were automatically moved from hosts made unavailable to hosts that were available as backup. The time to move the VM was quick - usually less than one minute. Once the VM was moved, it was automatically powered on and resumed activity. As mentioned earlier, in a production environment the time taken for this to occur can vary widely depending on your priority settings and availability capacity for outage and network isolation events.

Troubleshooting

Cisco Prime Collaboration 10.0 Troubleshooting

During the HA process the VM is essentially rebooted. Before using Cisco Prime Collaboration after HA, ensure that enough time is allowed for the application to fully come online and for all services to start. Refer to the Quick Start Guide and troubleshooting documents listed below on how to check Cisco Prime Collaboration services status.

Cisco Prime Collaboration Quick Start Guides can be found below:

<http://www.cisco.com/c/en/us/support/cloud-systems-management/prime-collaboration/products-installation-guides-list.html>

Troubleshooting Cisco Prime Collaboration:

http://docwiki.cisco.com/wiki/Troubleshooting_Cisco_Prime_Collaboration

vSphere HA 5.x Troubleshooting

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1001596.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)