White Paper

High Availability for Cisco Unified Communications on the Cisco Unified Computing System (UC on UCS)

Updated January 2013

Contents

UC on UCS High Availability

High availability and redundancy are important considerations in any network design. When deploying Unified Communications (UC) applications, these principles are even more important as UC applications continue to play a greater role in advancing productivity in the workplace. UC applications provide a variety of important capabilities including voice and video call control, voicemail, rich media conferencing, IM and presence and customer collaboration. UC applications also extend capabilities to mobile devices such as tablets and smart phones, ensuring user productivity while on the go. Users expect that regardless of what medium of communication they use, they will have a reliable and rich experience.

In order to provide these key benefits, it is important that UC applications are deployed in ways that ensure system uptime and availability. Redundancy, high availability and disaster recovery are therefore important considerations that UC architects should account for in their designs.

The Cisco Unified Computing System (UCS) platform offers the best platform for deploying Cisco UC applications. With Cisco UCS, customers have the choice of using rack mounted C-Series servers or B-Series blade servers. UC workloads can be deployed using a variety of options. When considering their customers' UC solution redundancy and highly availability requirements, Cisco and Partner sales engineers should engage their data center Product Sales Specialist (PSS) and data center Consulting Systems Engineer (CSE) in order to fully discuss, explore and understand the UCS, storage, and VMware features and functionality pertaining to redundancy and high availability. Additionally, fully engaging their UC PSS and CSE counterparts will also ensure that UC considerations are fully articulated and understood, including considerations such as how UC applications work in a VMware environment and which VMware features are supported.

While this white paper provides general recommendations and guidelines for high availability, disaster recovery and redundancy, ultimately, the choice for your customer of which features to implement in order to achieve their requirements depends on a number of factors. Some of these factors may include:

- The customer's budgetary constraints (what can they afford?)
- Business requirements for uptime and business continuity (what will the impact to the business be? How long can the customer afford to be hampered by an outage of an application or server?)
- Perceived cost/benefit balance, specifically between the cost of implementing a redundancy and high availability scheme and the effectiveness of such redundancy scheme. (Is it affordable? What does the cost buy the customer that they would otherwise lose)

The choice and recommendation of the solution design should therefore be arrived at after carefully considering these factors and dependencies.

This white paper is intended to supplement and summarize (but not replace) the information available from these sources:

- Support policies for UC virtualization at www.cisco.com/go/uc-virtualized at the time of this writing.
- General design recommendations and decision criteria in CiscoLive! sessions at the time of this writing (e.g. BRKUCC-2225 "Planning and Designing Virtual Unified Communication Solution") at https://www.ciscolive365.com)
- The document "Cisco Unified Communications Manager on Virtual Servers" at http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/virtual/servers.html
- Cisco Unified Communications design guides at www.cisco.com/go/ucsrnd
- Cisco Unified Computing System technical documentation at www.cisco.com/go/ucs
- VMware technical documentation at www.vmware.com

Hardware Redundancy

In a UC on UCS deployment, the goal of hardware redundancy is the same as in non-virtualized ("bare-metal") deployments. It is to provide a means of maintaining survivability in the event of a failure of either a component of, or the entire hardware system, which could include a physical server or a chassis that contains blades. A failure could be caused by a loss of functionality of any one of the hardware components including network interface card (NIC), CPU, power supply, hard drive, or even a cable connecting a server to the LAN or SAN.

Recommendations for how to reduce the impact of hardware failures include the following:

- Redundant physical servers (redundant blades in redundant blade server chassis, or redundant rack-mount servers)
- Redundant power supplies, for example, dual power supplies on servers, or grid power supply configurations on UCS blade server chassis.
- Use of Uninterrupted Power Supply (UPS) systems.
- Redundant LAN links
- RAID with redundant disks where Direct Attached Storage (DAS) is used.
- RAID on array + redundant SAN switching and links where Storage Attached Networks (SAN) is used.

Redundant physical servers

While virtualization and co-residency help to reduce the number of servers needed in UC deployments when compared to deployments on physical servers, the need for multiple physical servers does not go away completely. Multiple physical servers may still be needed for redundancy, for geographic coverage or to scale as the size of the deployment grows and more applications or instances of an application are added to the solution.

Redundant physical servers are a common deployment requirement providing a means for a secondary server to take over functionality that may be lost upon the failure of a primary server. This means, for example, providing redundant B-Series blade servers or C-Series rack servers. This provides not only physical server high availability, but also an easy way of ensuring application survivability and redundancy.

A catastrophic hardware failure such as a CPU failure or other critical failure that decommissions a rack server can be mitigated by deploying redundant rack servers. A catastrophic failure that decommissions a blade in a chassis can be mitigated by deploying redundant blades. A catastrophic failure that decommissions a chassis and all the blades in it can be mitigated by deploying a redundant chassis.

As a general rule, provide sufficient levels of hardware redundancy to allow the applications hosted on each server to maintain full functionality on a secondary server in the event of a failure of the primary server. For example, in a Business Edition 6000 deployment, consider deploying a redundant UCS server configuration.

The diagrams below depict sample configurations for hardware redundancy.

Figure 1. Redundant Rack Mount Servers

Figure 2. Redundant Blades in the Same Chassis

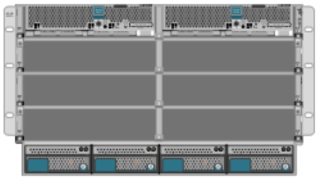

Figure 3. Redundant Blades in Redundant Chassis

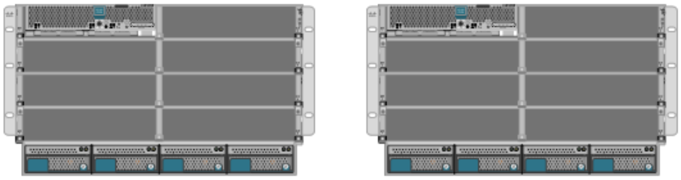

Redundant power supply

Redundant power supply provides a means for maintaining power to a server even when one or more of the power supplies has an outage. UCS servers provide options for power supply redundancy.
With C-Series UCS servers, customers have the option of redundant power supplies. Cisco recommends using dual power supplies whenever available. A redundant power supply provides a cost effective means of reducing the chances of a server outage due to a power supply unit failure.

With B-Series UCS solutions, power is supplied to the 5108 Blade Server chassis which has 4 power supply bays, and the fabric interconnects which have two power supply bays. You can configure power modes to either use the combined power provided by the installed power supply units or to provide power redundancy when there is a power outage.

Cisco Unified Computing System (UCS) enables end users to select the power mode that fits their redundancy needs through the Cisco Unified Computing System Manager (UCSM). Within the Cisco UCS Manager, the power redundancy mode is selected via the global policy tab, as shown below.

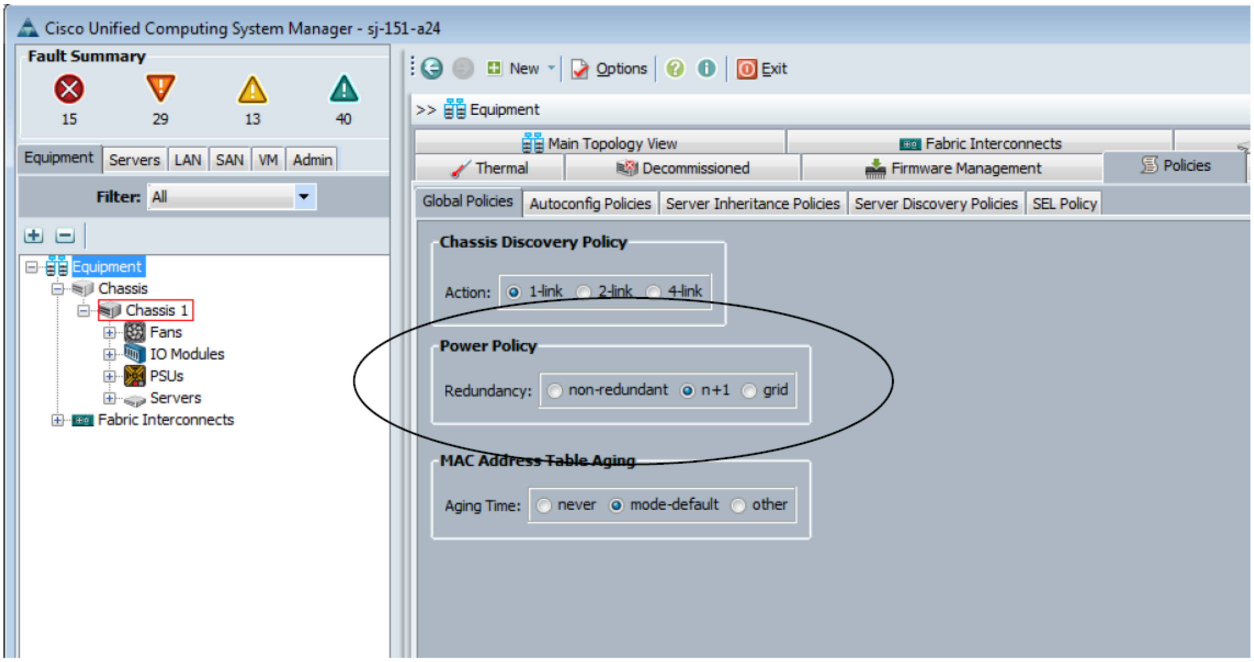

Chassis Discovery Policy

Action: ● 1-link ○ 2-link ○ 4-link

Power Policy

Redundancy: ○ non-redundant ● n+1 ○ grid

MAC Address Table Aging

Aging Time: ○ never ● mode-default ○ other

The power policy settings are described as follows:

- Non-redundant means that uptime cannot be guaranteed in the event of a failure.
- N+1 means that the system can tolerate the failure of one supply.
- Finally, grid redundancy (N+N) means that the system, if wired correctly into dual independent AC feeds, can tolerate the loss of one of those grids or half of the power supplies.

The power supplies are all operated in parallel output. You should connect two separate input sources (grids) to have the highest level of availability (grid redundancy). The system will operate on two power supplies (2+2 redundancy) for the Cisco UCS 5108 blade server chassis and one power supply (1+1 redundancy) for the Cisco UCS 6100 Series Fabric Interconnect. The diagram below depicts the use of two separate input sources (grids) to provide grid redundancy.

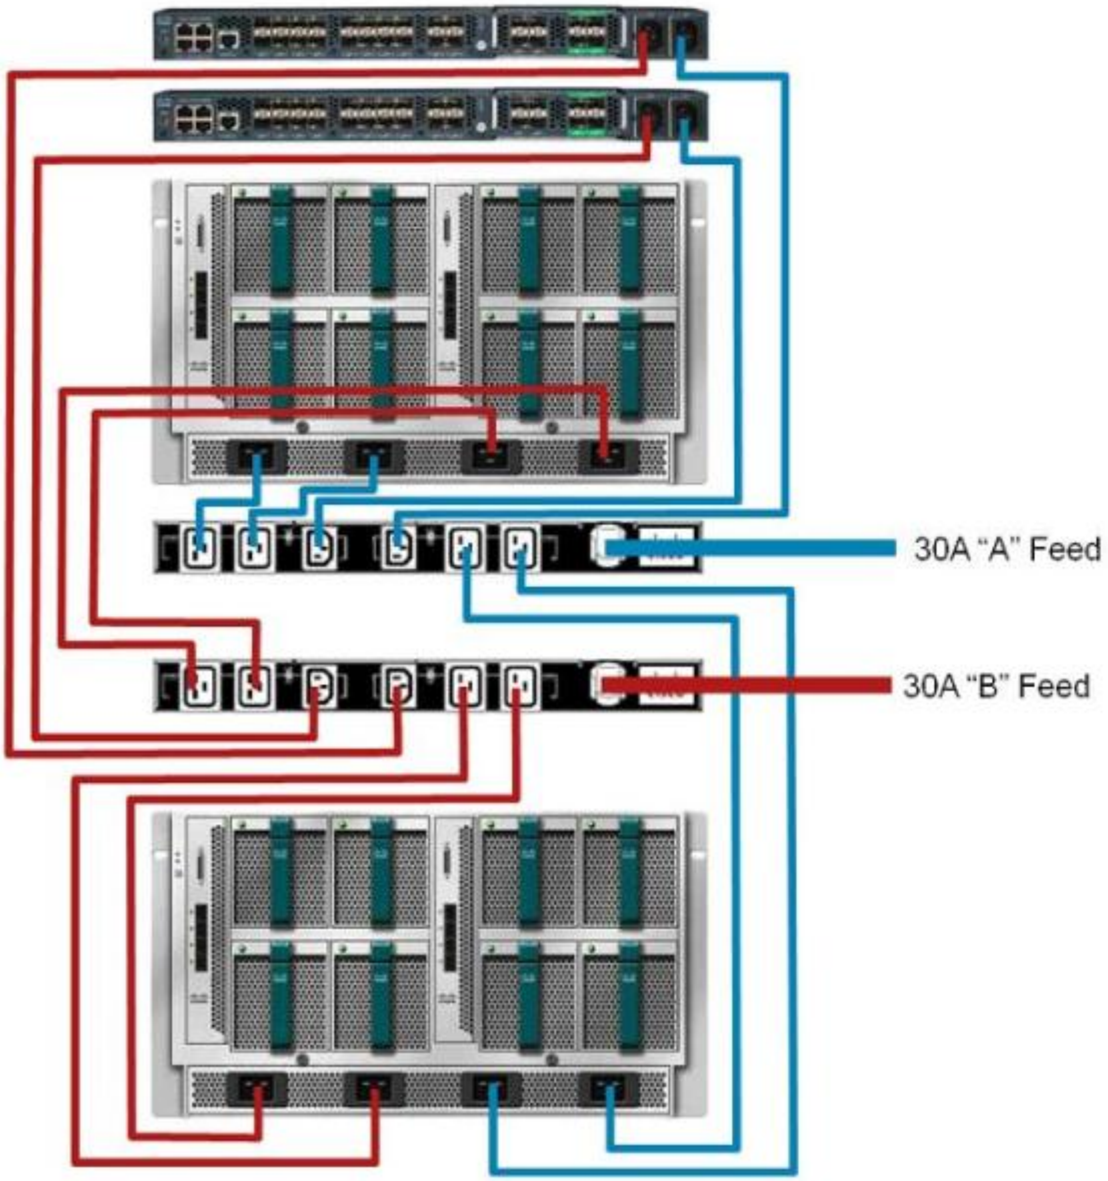

Additional details about the different power modes can be found in the 5108 Blade Server chassis Installation Guide
Cisco recommends implementing the highest level of power redundancy whenever possible.

Uninterrupted Power Supply (UPS) systems

A facility component such as an uninterruptible power supply (UPS) can be an important tool in adding an extra layer of availability to a UC on UCS solution. When a server, chassis, or fabric interconnect is connected to a UPS system, the effect of a power outage can be mitigated.

Redundant LAN links

Bare-metal appliance deployments frequently made use of “NIC teaming” for LAN link redundancy.

On a virtualized C-Series rack-mount server, LAN link redundancy is accomplished by leveraging both the LAN On Motherboard (LOM) Ethernet ports as well as Ethernet PCIe NICs (such as the Intel i350 used in some UC on UCS Tested Reference Configurations which are documented here: http://docwiki.cisco.com/wiki/UC_Virtualization_Supported_Hardware)

On servers with both LOM and PCIe NIC ports, Cisco recommends using a port from the LOM and another port from the PCIe when creating a NIC teaming pair. Splitting the team across the motherboard and PCIe card provides protection in the event of a problem with either the PCIe card or the motherboard ports.

On a virtualized B-Series blade server, all LAN connectivity is via the Fabric Interconnect switches which support multiple modules and ports for LAN interconnect. Configure redundant ports on redundant modules for link redundancy.

Disk Redundancy

With virtualized deployment of Cisco UC applications, the storage (disk) component of the solution plays a critical role in ensuring the availability of the applications. When deploying UC applications, ensure that the storage component in use is highly available and provides redundant capabilities. Where possible, it is also important to separate the virtual machine from the application files used by the virtual machine. If this is done, then in the event of a disk or other hardware problem, it will be a trivial task to boot up your virtual machine on another blade or rack server as long as the redundant physical storage for the virtual machine is intact.

RAID with Redundant Disks

Latest UCS C-Series Tested Reference Configurations all use DAS storage and make use of RAID arrays with redundant disks. Configuration details and requirements can be found here: http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/virtual/servers.html

Latest UCS B-Series Tested Reference Configurations are diskless and presume use of Fibre Channel SAN for all software (UC applications and VMware vSphere). The storage array provides disk redundancy.

UCS Specs-based allows many other storage options for DAS, SAN or NAS.

RAID “rules” and allowed storage options for UC can be found here: http://docwiki.cisco.com/wiki/UC_Virtualization_Supported_Hardware#Storage

RAID on Array + Redundant SAN Switching and Links

B-Series or C-Series deployments with SAN/NAS storage leverage RAID in the storage array as well as redundant storage access network links (Fibre Channel, Ethernet or 10Gb Ethernet depending on the storage transport to the array).

Other storage best practices for UC are documented here: http://docwiki.cisco.com/wiki/UC_Virtualization_Storage_System_Design_Requirements

Application-level redundancy: Cisco UC Application approaches vs. VMware approaches

Cisco Application Layer Approach

All Cisco UC applications provide native capabilities for implementing redundancy and high availability. Cisco does not recommend relying solely on VMware approaches for UC redundancy, but rather a complementary approach using both application approaches and VMware approaches together. The following are some of the methods by which application redundancy may be implemented:

- Clustering – Publisher node working in conjunction with one or more subscriber nodes, and sharing a common database. With clustering, multiple nodes may be deployed, with secondary nodes providing backup for primary nodes in the cluster. Examples of clustering as a means of providing high availability and redundancy include clusters in Unified Communications Manager, Unity Connection, or Unified Contact Center Express.

  With Unified Communications Manager, a cluster with multiple nodes can be configured using redundancy groups which comprise a collection of three Unified Communications Manager nodes configured in a prioritized order. Each redundancy group is associated with a pool of devices which would register with and get call control and call processing functionality from the Unified Communications Managers in the group based on the prioritized order. If the primary Unified Communications Manager node in a group fails, the devices registered to it will automatically register to the second Unified Communications Manager in the list. If there is a tertiary server in the group, it takes control in the event that the primary and secondary servers in the list fail. In effect, the Unified Communications Manager cluster with Unified Communications Manager group configuration provides a means of implementing redundancy.

  When a Unified Communications Manager cluster has a pair of redundant subscribers, an outage of the primary subscriber will result in the secondary subscriber taking over after a predetermined period of time, typically defined by a keepalive timer. For example, with keepalives being exchanged every 15 seconds, a secondary subscriber may take over after 45 seconds (after 3 missed keepalives).

- High Availability for applications such as Unified Contact Center Express and Unity Connection. In this example, the high availability feature works in conjunction with the clustering mechanism. When the high availability feature is enabled in Unified Contact Center Express and Unity Connection, the first node configured is automatically the publisher, and the second node becomes the subscriber. These publisher-subscriber pairs provide seamless, fast, and automatic failover from one to the other in the event of a problem with the primary node.

In addition to the native methods for increasing availability, most UC applications also provide additional means for recovering in the event of an outage. The Cisco Disaster Recovery System (DRS) is one such method. With Cisco DRS, customers can perform regular backups of the UC applications in question, and have a means to provide full data recovery of the cluster in the event of a problem.

When performing a system recovery of an application using Cisco DRS, one needs to first reinstall the application, then perform a system restore using the backup files that were generated during the Cisco DRS backup. This allows for the system to be restored to the same state that it was in at the time of the backup.

Cisco recommends that customers use Cisco DRS (where applicable) as the primary means to backup their UC applications.

Cisco provides guidelines and recommendations in various documents such as the Solution Reference Network Design (SRND) guides (www.cisco.com/go/srnd) for how clustering may be implemented. Additionally, such design documents provide rules for how to distribute application roles across multiple virtual machines. Such rules and guidelines are intended to ensure a balanced approach to implementing high availability, as well as to ensure that common sense approaches such as distributing virtual machines in order to avoid single points of failure or to otherwise expand the failure domain are observed.

As an example, the Unified Communication Manager SRND describes (depending on size) the deployment of publisher/subscriber pairs or the deployment of the following call processing redundancy configuration options:

1. A two to one (2:1) redundancy option where there is a shared secondary or backup call processing subscriber for every two primary call processing subscribers
2. A one to one (1:1) redundancy option where there is a secondary or backup subscriber for every primary call processing subscriber. This is 1:1 with active/standby. You can also have 1:1 with load balancing, where device registration is split between the two subscriber pairs.

Guidance can also be found in the SRND recommending breaking out various servers that play certain roles such as TFTP or Music on Hold servers in large deployments and having dedicated servers performing those roles. VMware features for implementing high availability and redundancy do not address or replace these application server capabilities.

When deciding how to distribute applications to achieve redundancy and high availability goals, consider distributing virtual machines across multiple physical servers in order to provide server or hardware redundancy. At the same time, consider the customer's objectives and specific needs and constraints, and balance the goal of a reduced server footprint versus the need to minimize the impact of a host failure.

An additional consideration when determining how applications may be deployed to achieve highest levels or redundancy is the placement logic. A well thought out placement logic helps minimize the impact of hardware failure or unavailability. When planning the placement logic to use, consider the following:

- Place the primary and secondary call processing subscribers on separate servers, chassis or sites.
- When using redundancy groups, place active servers on separate servers, chassis or sites.
- Distribute virtual machines with the same role across different servers, chassis or sites.

The depiction below shows a deployment of several UC virtual machines on a single large capacity server. While this provides least hardware footprint and cost, from an availability perspective this is an example of an "all eggs in one basket" approach. A hardware failure of the physical server would affect all applications in the deployment, including negating application redundancy.

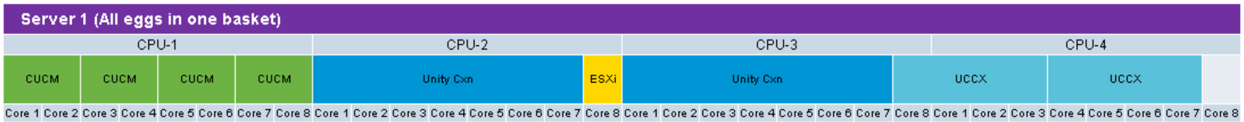
Server 1 (All eggs in one basket)
CPU-1
CPU-2
CPU-3
CPU-4
CUCM
CUCM
CUCM
CUCM
Unity Cxn
ESXi
Unity Cxn
UCCX
UCCX
Core 1 Core 2 Core 3 Core 4 Core 5 Core 6 Core 7 Core 8 Core 1 Core 2 Core 3 Core 4 Core 5 Core 6 Core 7 Core 8 Core 1 Core 2 Core 3 Core 4 Core 5 Core 6 Core 7 Core 8 Core 1 Core 2 Core 3 Core 4 Core 5 Core 6 Core 7 Core 8

Resource Usage: 31 Cores, 56GB RAM

The depiction below shows the same UC deployment (i.e. same application mix and virtual machine type / quantity capacity points as in the previous depiction), but on multiple physical servers to provide layers of redundancy. Application VMs are distributed across physical servers. E.g. Servers 1 and 3 could be running the "primary" VMs with Servers 2 and 4 running the "secondary" VMs. A hardware failure of any one of these physical servers will not impact the applications due to combined application and hardware redundancy.

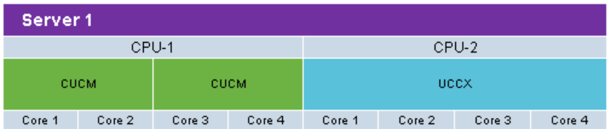

Resource Usage: 8 Cores, 20GB RAM

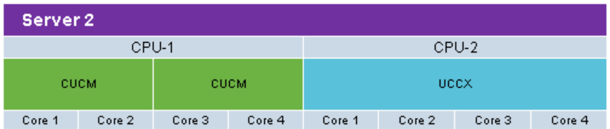

Resource Usage: 8 Cores, 20GB RAM

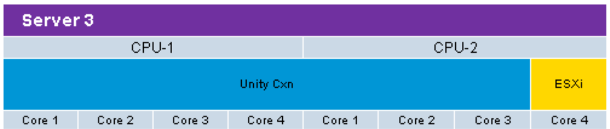

Resource Usage: 8 Cores, 8GB RAM

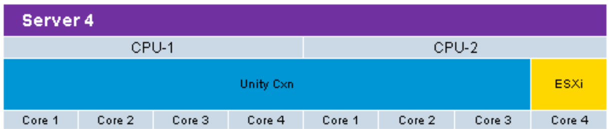

Resource Usage: 8 Cores, 8GB RAM

When considered carefully and implemented according to design guidelines, native application features for providing high availability and redundancy provide the fastest time to failover in the event of an outage.

VMware Approaches

Various VMware features provide methods for achieving high availability and survivability of virtualized applications. Some of the VMware features designed to provide redundancy and high availability include features such as the following:

- VMware High Availability
- VMware Site Recovery Manager
- VMware Fault Tolerance

VMware High Availability (HA)

VMware vSphere High Availability provides high availability for virtual machines and the applications running within them by pooling the ESXi hosts they reside on into a cluster. Hosts in the cluster are continuously monitored. In the event of application, OS or VM failure, the virtual machines on a failed host are attempted to be restarted on alternate hosts.

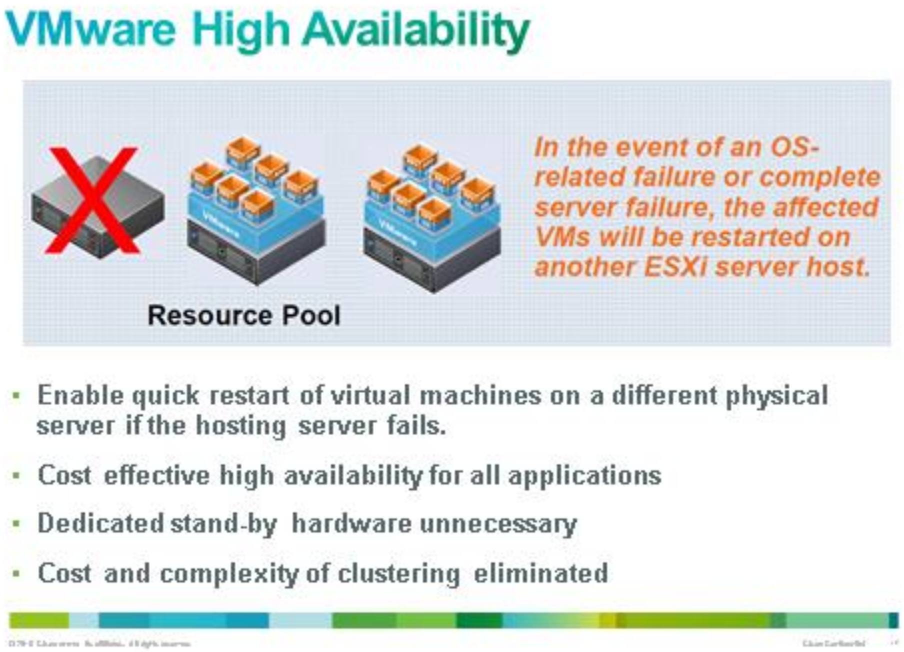

VMware High Availability

In the event of an OS-related failure or complete server failure, the affected VMs will be restarted on another ESXi server host.

- Enable quick restart of virtual machines on a different physical server if the hosting server fails.
- Cost effective high availability for all applications
- Dedicated stand-by hardware unnecessary
- Cost and complexity of clustering eliminated

VMware HA can allow for a UC application on a failed host to be restarted on a separate physical host, allowing for fast recovery. Failovers to other physical servers must not result in an unsupported UC deployment model. For example, the destination server must align with supported co-residency rules after failover occurs. Sizing rules such as no support for oversubscription must also not be violated. The destination server must additionally align with supported hardware policies.

Cisco recommends regular backups via mechanisms like Cisco Disaster Recovery System (DRS) even if the customer plans to deploy VMware HA.

The following caveats exist when using VMware HA with UC with UCS:

- VMware HA does not protect from VMFS corruption
- VMs running on DAS cannot be restarted via VMware HA if there is DAS storage failure.

- VMs running on SAN/NAS can be restarted via VMware HA if the alternate server is on same SAN as the failed server.
- VMware HA without complementary application redundancy can increase service outage duration as the recovery time is based on the time to boot up a virtual machine and for all services to start (typically about 15 minutes for Unified Communications Manager). VMware HA is a good solution to protect from host failure, but application approaches provide faster recovery from software issues.

VMware Fault Tolerance

VMware Fault Tolerance is a feature that allows for fast failover of VMs with no interruption of services on the VM in question. Failures are detected and fast and seamless failover takes place, thereby minimizing service interruptions.

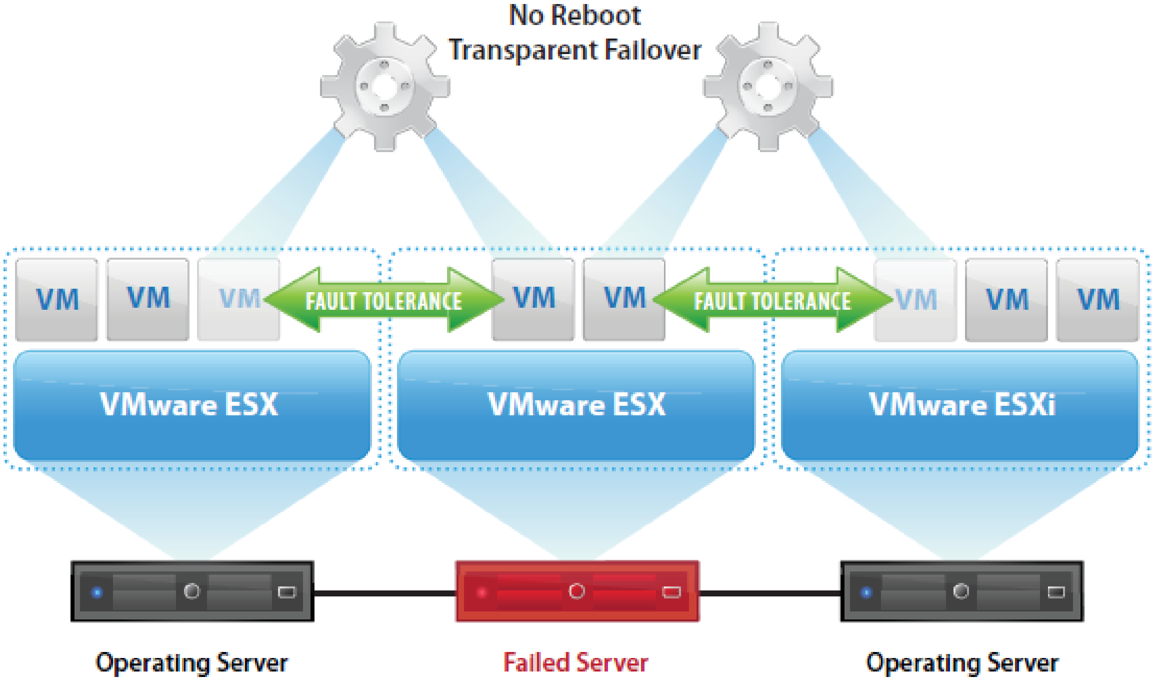

VMware Fault Tolerance enables a transparent failover with no disruption of service in the event of hardware failures

VMware Fault Tolerance is generally not needed or supported with Cisco UC applications, as the feature provides minimal value beyond the native Cisco approaches to redundancy and high availability described earlier. Also note the following:

- VMware Fault Tolerance is essentially a ‘mirroring’ technology with virtual machines on different servers running in lockstep (the “active” VM and the “shadow” VM). This does not provide any incremental value over redundant application instances.
- Implementing the VMware Fault Tolerance feature does not provide savings in servers or hardware. VMs are duplicated (due to “active” vs. “shadow”), and since the shadow virtual machine resides on a different host physical servers, the result is duplicated hardware, virtual machines, cores and memory.

- The VMware Fault Tolerance feature only works with 1vcpu VMs. There is no official support by VMware for virtual machines with more than 1vcpu. This means that this feature cannot be used for most UC virtual machine templates (“OVAs”).
- If the primary virtual machine crashes, then by virtue of the fact that the shadow virtual machine is running in lockstep, then the shadow virtual machine will also crash. Contrast this with redundant Unified Communications Manager primary and secondary pairs providing failover where the virtual machines are unaffected by each other.
- VMware Fault Tolerance only protects against local host failure. The feature does not help to mitigate SAN or LUN storage failure.
- The feature requires the active virtual machine and the shadow virtual machine to be on the same SAN and same LUN. This restriction impacts geo-redundancy feasibility. This restriction does not exist with redundant Unified Communications Manager primary and secondary pairs.

VMware Site Recovery Manager

Site Recovery Manager (SRM) provides a means of replicating virtual machines at a secondary site. Virtual machines that are protected with Site Recovery Manager are replicated at the secondary site and made available when the primary site experiences an outage.

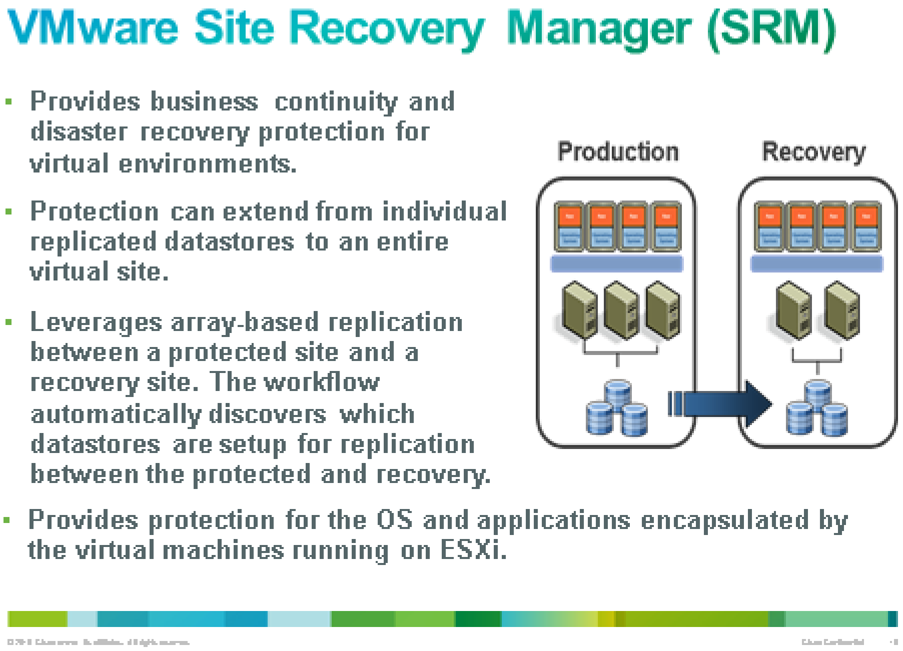

VMware Site Recovery Manager (SRM)

- Provides business continuity and disaster recovery protection for virtual environments.
- Protection can extend from individual replicated datastores to an entire virtual site.
- Leverages array-based replication between a protected site and a recovery site. The workflow automatically discovers which datastores are setup for replication between the protected and recovery.
- Provides protection for the OS and applications encapsulated by the virtual machines running on ESXi.

Some Cisco UC applications may be protected using SRM. The applications in question must be able to tolerate active/passive failure. Support varies by Cisco UC application and version, so one must verify support on www.cisco.com/go/uc-virtualized before implementing this feature. In a disaster recovery situation, UC applications can be restored at the recovery site to allow for business continuity.

Note that VMware SRM requires SAN replication to be implemented in the customer’s storage arrays across the production and recovery environments and sites.

Cisco recommends regular backups via mechanisms like Cisco Disaster Recovery System (DRS) even if the customer plans to deploy VMware SRM. Cisco also recommends deploying Clustering over the WAN as a primary site redundancy approach whenever the deployment can meet the Clustering over the WAN requirements, such as delay and bandwidth.

The following caveats exist when using VMware SRM with UC with UCS

- VMware SRM does not protect from VMFS corruption nor from a file system corruption of the UC application. If the file system of a UC application is corrupted, the copy of the file system in the alternate site will also have a corrupted file system
- VMware SRM without complementary application redundancy can increase service outage duration. VMware SRM is a good solution to protect from site failure, but the application redundancy Clustering over the WAN approaches provide faster recovery from software issues.

Other Geographic Redundancy Considerations

Where a customer has a 2-site strategy with data centers that are geographically dispersed, then either a “hot site” or a “cold site” may make sense. There is inherent cost and complexity associated with such designs, but geographic redundancy can be critical in ensuring that there is prompt recovery in the event of an outage that affects an entire site.

When planning for a high availability scheme, customers should consider:

- Building in redundancy to minimize outages and service disruptions
- Designing systems to ensure that they can recover quickly from an outage

When considering geographic redundancy, Cisco recommends always using application layer redundancy where possible. Guidelines such as minimum bandwidth and latency between sites should be observed.

In a “hot site” model, the VMware Site Recovery Manager feature (described above) implements SAN replication between the two data centers to protect from site, storage or host failure. Where SRM is used, customers should be aware that the feature is costly, and that there are no guarantees for proper functioning of UC applications.

Where cold standby methods of redundancy are used, keep in mind that recovery is a manual process and may therefore take time. Additionally, database replication may be needed before applications function correctly.

Unified Communications Manager

When a Unified Communications Manager cluster is deployed across two geographically dispersed sites with clustering over a WAN, the primary and secondary subscribers machines may be split across the sites (“Cluster over WAN”). This can provide true redundancy where service outage as a result of a problem at one site is prevented or minimized because there is a second site which is essentially in “hot standby”. Maintaining the “hot site” itself may, however, be expensive due to the required infrastructure.

Instead of maintaining a “hot site” some customers choose to maintain a “cold site” that only needs to be brought online in the event of an outage at the primary site. With a “cold standby” implementation, the customer has to go through the process of bringing up the applications once there is an outage at the primary site. This could mean “rack, stack and cable” work, powering on switches and routers, booting up servers and virtual machines and running restore procedures from backup, changing IP addresses and hostnames and updating routing tables. The recovery process may therefore take time before services are brought back up and fully operational. There is also a chance that at boot up, the application node databases are out of sync and require replication before full functionality is restored. However, since the cold standby site is only used in the event of an emergency, it may be less expensive to maintain than a “hot site”.

VMware Dynamic Resource Scheduler and Dynamic Power Management

Features such as VMware Dynamic Resource Scheduler and Dynamic Power Management are not supported for UC deployments. These features are primarily designed to dynamically balance computing resources across a collection of hardware resources. Because resource utilization requirements on UC applications are predefined, features which dynamically allocate resources to virtual machines do not provide any incremental value. Additionally, Dynamic Power Management makes decisions to consolidate virtual machines onto fewer hosts and to power hosts on and off to save energy. Since such decision making could result in violation of UC deployment rules such as co-residency and oversubscription, this feature is not supported for Cisco UC applications.

Clones, Snapshots and Templates

VMware supports capabilities such as the clone, snapshot and template features.

A clone is simply a copy of a virtual machine. When a clone is made, a copy of the virtual machine is taken, including all of its settings and installed applications. A clone also includes the contents of the original virtual machines disks.

A snapshot is a virtual machine that represents the original virtual machine's state at the moment the snapshot was taken. Snapshots are useful when there is a need to often revert to the original virtual machine state such as in a test or development environment. Snapshots are not intended for use as a backup/restore solution.

A template is a master copy taken of a virtual machine and used for the purpose of creating multiple similar virtual machines. For example, if you have a need to create several virtual machines, you can start by creating the first virtual machine, then creating a template from it and using the template to create the rest of the virtual machines, thus saving time.

While some of these features might be considered potentially useful for backing up and restoring applications, Cisco recommends using the Cisco Disaster Recovery System (DRS) as the primary method of backing up Cisco UC applications. Additionally, support for these features varies from application to application. Reference the Unified Communications VMware Requirements document located at the link below for which features are supported by which applications, and for other important caveats such as the need to shut down the application before using some features.
http://docwiki.cisco.com/wiki/Unified_Communications_VMware_Requirements

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.

Printed in USA