Virtual Machine Mobility with Vmware VMotion and Cisco Data Center Interconnect Technologies

What You Will Learn

VMware has been the industry leader in virtualization technologies for the past decade and has brought to the data center several new features that enable faster and better provisioning of business-critical applications. One of the features is the VMware® VMotion™ technology that allows virtual machine mobility between two VMware vSphere™ servers instantaneously with no application downtime. The capability to migrate applications with no perceivable effect to the end user enables IT departments to develop new and improved methods for provisioning and maintaining data center infrastructure. IT departments can now perform hardware maintenance, consolidate CPU and memory resources, or migrate mission-critical applications from a data center when necessary without affecting the service-level agreements (SLAs) of the applications.

A successful application migration through VMware VMotion heavily relies on the underlying network infrastructure. Therefore it is extremely important that IP network be resilient, robust, and highly available.. The IP network becomes more important when the applications have to be mobile across data centers. Cisco has been the industry leader in IP network and routing technologies and has been providing data center IP network extensions since the 1980s. Cisco® switching and routing technologies provide the robust and redundant network that is essential for VMware VMotion to succeed.

This document discusses the VMware VMotion feature and Cisco networking technologies essential for application mobility across data centers.

VMware and Cisco Migration Solution

The VMware and Cisco solution shown in Figure 1 enables customers to perform live application migration across data centers. The components used are a VMware vSphere 4.0 server cluster enabled with VMware VMotion in each data center, a VMware vCenter server, and a data center interconnect (DCI) WAN. The applications provisioned on the VMware vSphere server can be migrated across the data centers or a private cloud with no application downtime. The solution itself does not need any new software or hardware to perform these migrations.

Figure 1. VMware VMotion Across Data Center

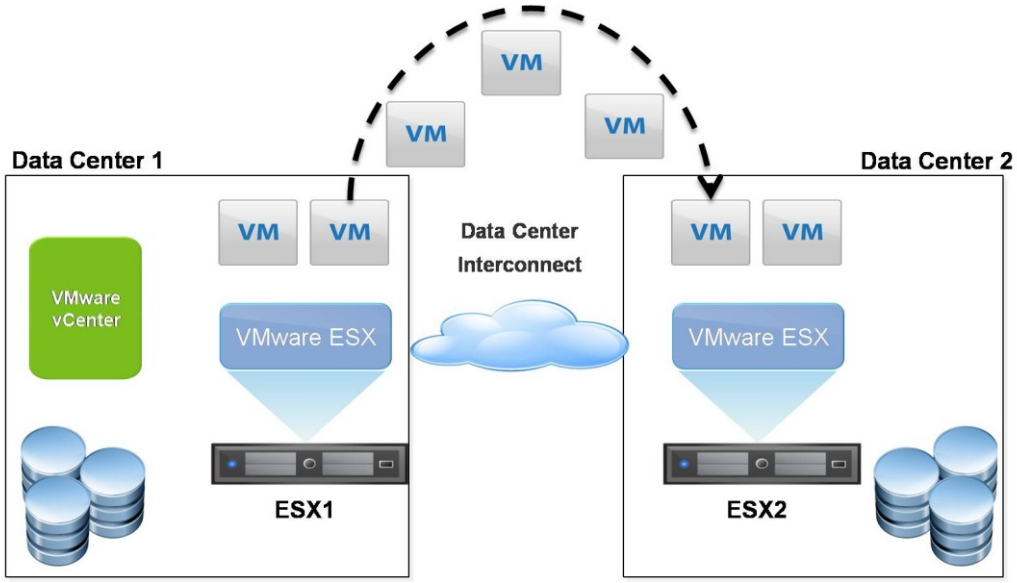

Need for Virtual Machine Mobility Across the Data Center

The changing model of data center management and provisioning allows VMware VMotion to be used for several purposes without violating the application SLAs.

- **Data center maintenance without downtime:** Applications on a server or data center infrastructure requiring maintenance can be migrated offsite without downtime.
- **Disaster avoidance:** Data centers in the path of natural calamities (such as hurricanes) can proactively migrate the mission-critical application environment to another data center.
- **Data center migration or consolidation:** Migrate applications from one data center to another without business downtime as part of a data center migration or consolidation effort.
- **Data center expansion:** Migrate virtual machines to a secondary data center as part of data center expansion to address power, cooling, and space constraints in the primary data center.
- **Workload balancing across multiple sites:** Migrate virtual machines between data centers to provide compute power from data centers closer to the clients (“follow the sun”) or to load-balance across multiple sites. Enterprises with multiple sites can also conserve power and reduce cooling costs by dynamically consolidating virtual machines into fewer data centers (automated by VMware Dynamic Power Management [DPM]), another feature enabling the green data center of the future.

The application mobility discussed in this document provides the foundation necessary to enable cloud computing—for example, cloud import and export—providing the flexibility to move virtual machines into the cloud from an enterprise data center, to move them between different clouds, and to move them back into the enterprise data center.

VMware VMotion Requirements

VMware VMotion application mobility is based on certain infrastructure requirements:

- An IP network with a minimum bandwidth of 622 Mbps is required.
- The maximum latency between the two VMware vSphere servers cannot exceed 5 milliseconds (ms).
- The source and destination VMware ESX servers must have a private VMware VMotion network on the same IP subnet and broadcast domain.
- The IP subnet on which the virtual machine resides must be accessible from both the source and destination VMware ESX servers. This requirement is very important because a virtual machine retains its IP address when it moves to the destination VMware ESX server to help ensure that its communication with the outside world (for example, with TCP clients) continues smoothly after the move.
- The data storage location including the boot device used by the virtual machine must be active and accessible by both the source and destination VMware ESX servers at all times.
- Access from VMware vCenter, the VMware Virtual Infrastructure (VI) management GUI, to both the VMware ESX servers must be available to accomplish the migration.

Challenges

Figure 2. Infrastructure Challenges to Vmware Vmotion Across Data Centers

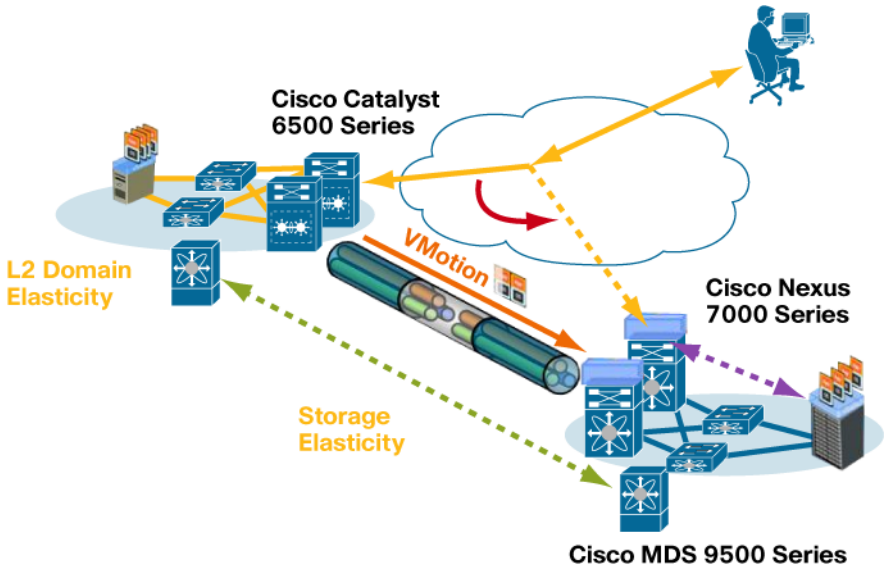

Facilitating VMware VMotion migration across physical infrastructure boundaries and data centers using the WAN poses specific challenges for the data center network and storage designs as shown in Figure 2. In particular, the LAN and SAN designs have to be addressed.

- **LAN (Layer 2 Domain Elasticity):** The LAN must be extended across the physical locations or data centers without compromising the availability, resilience, and security that exists within a single physical location or data center. The best practices for the current network

deployment require separate virtual LANs (VLANs) for the different physical locations. The extension of the VLAN across data centers must be implemented in a way that preserves the configuration and security constraints of a single data center.

- **Data availability (Storage Elasticity):** The availability of data to the two VMware ESX servers is critical to a successful migration of the application. The SAN design should encompass all the properties and parameters of the data within the same data center. This requirement poses a heightened security risk, since data now has to traverse the WAN. Data network designs should take into account these parameters to help ensure that data is not only available, but secured and available with I/O latencies that will not affect the SLAs of the applications.
- **IP localization:** A primary requirement for application mobility is that the migrated virtual machine maintains all its existing network connections after it has been moved to the secondary data center. Traffic routing to and from the virtual machine needs to be optimized so that any traffic flows in an optimized way to the virtual machine's new location. If the traffic to the virtual machine originates in the same Layer 2 domain, then the Layer 2 extension will suffice. However, if the traffic to the virtual machine is traversing a Layer 3 network or the Internet, then granular routes needs to advertised by the secondary data center for migrated virtual machines. If these changes are not provisioned, suboptimal routing may result in additional delay, which may or may not be acceptable depending on the specific virtual machine or application. To help ensure that the traffic from the virtual machine is optimally routed, the default gateway IP addresses on the subnets in both the primary and secondary data centers need to be identical, a goal achieved through active-active Hot Standby Router Protocol (HSRP) configuration.
- **Services localization:** Another critical requirement is helping ensure that any applicable services (firewall, load balancer, etc.) policies are retained as the virtual machine moves from one data center to another data center.

In addition, WAN characteristics such as available bandwidth for VMware VMotion and SAN extension and its latency are crucial factors in the eventual success of application mobility.

Solution Options

The challenges that arise from migrating applications across data centers are addressed by the joint VMware and Cisco solution. The solution uses best practices for the LAN, data network, and WAN to overcome the challenges and provide the IT staff with a tool for migrating applications between data centers without downtime.

The solution can be designed in multiple ways based on the topologies used for the LAN and storage extensions, as shown in Table 1.

Table 1. Application Mobility Using Vmware Vmotion Solution Options

Network and Storage Topologies	Shared Storage	Active-Passive Storage	Active-Active Storage
Extended or Stretched VLAN	Storage remains at original location	Storage is migrated before virtual machine migrates	Storage is actively available at both locations

Extended VLAN with Shared Storage

An extended VLAN and shared storage architecture extends the VLAN between the two sites, but with storage remaining at the original location. When the virtual machine migrates to the remote data center, the application will access the storage from the primary site. Storage is not provisioned for the application at the remote data center; hence, there is only one copy of the storage at any given point in time. This design can be appropriate when the distance between the data centers is not great, since I/O latency will affect application performance.

Extended VLAN with Active-Passive Storage

An extended VLAN with VMware Storage VMotion topology requires migration of storage to the remote data center prior to migration of the virtual machine itself to the remote data center. Storage is migrated to remote data center prior to the virtual machine migration using VMware Storage VMotion. VMware Storage VMotion migrates the data space associated with a virtual machine to the secondary storage location and enables a virtual machine to access this new storage after the VMware Storage VMotion migration is complete.

If storage replication were in place between the data centers, the volume(s) containing the virtual machine data could be readily available in real time at the secondary data center. The existing Active-Passive storage replication techniques require a set of explicit control operations to make the storage replica accessible to the servers in the secondary data center. Consequently, at present, this is not a supported technology to perform virtual machine vMotion.

Extended VLAN and Active-Active Storage

An extended VLAN and active-active storage solution incorporates technologies that make data actively available at both the local and remote data centers at all times. The LAN extends across the data centers, and storage is provisioned in both data centers. Data is replicated across data centers using synchronous replication technology and rendered in an active-active state by the storage manufacturer. Normally when data is replicated, the secondary storage is locked by the replication process and is available to the remote server only in a read-only state. In contrast, active-active storage allows both servers to mount the data with read and write permissions as dictated by the VMware VMotion requirements.

VMware VMotion Across Very Long Distances (Routed or Disparate IP Subnets)

Deploying VMware VMotion across data centers that are dispersed over very long distances (500 miles or more) potentially involves moving the virtual machine to an entirely new subnet, but the goal continues to be to help ensure that the IP address of the virtual machine as well as the existing client connections are not disrupted. This type of VMware VMotion migration is not possible with existing technologies. Special hardware and software features will be required to route the TCP connections

to the virtual machine in its new location without terminating the sessions. This approach will require the redesign of the IP network between the data centers involving the Internet. Technologies are being developed by Cisco, VMware, and standards organizations to address this network scenario in the future.

VMware and Cisco Solution

The VMware and Cisco solution, jointly engineered by the two companies, addresses two of the available topologies: extended VLAN with shared storage and extended VLAN with active-passive storage. The two components of this solution are LAN extension technologies and SAN extension or storage availability technologies.

LAN Extension Technologies

To make the same LAN to be available across the data center, Cisco has solutions to suit the type of DCI available in your network. Any solution that extends the Layer 2 subnet across data centers needs to meet the following requirements:

- **High availability:** The solution must help ensure that no link or device failure will cause the Layer 2 extension to be disconnected and thus affect the ability to perform VMware VMotion migration across the data center.
- **Load balancing:** The solution should fully utilize cross-sectional bandwidth between the data centers; DCI connections are usually more expensive than the LAN, making effective use of the available connection critical.
- **Spanning Tree Protocol isolation and loop and broadcast storm prevention:** The solution must fully contain and isolate Spanning Tree Protocol within each data center with Bridge Protocol data units (BPDUs) filtered at the boundary of each edge switch facing the core. Network loop and broadcast storm avoidance features need to be available to prevent disruption of applications.
- **Scalability:** The solution must be able to handle multiple VMware VMotion migrations concurrently. To meet this requirement, the network must be able to scale. The available bandwidth, the number of VLANs, and the number of data centers connected through the solution should all be capable of expansion as needed.

Additional services may be required in many cases:

- **Encryption:** IP Security (IPsec) or Layer 2 link encryption (IEEE 802.1AE) may be needed to help ensure the privacy and confidentiality of the data traversing between the data centers. This requirement is particularly important if the interconnection crosses a public network. Encryption may also be a requirement for regulatory compliance.
- **Hierarchical quality of service (HQoS):** HQoS may be required to help ensure quality of service (QoS) for VMware VMotion, particularly on WAN links acquired through a service provider. HQoS is important on interconnecting devices when an enterprise subscribes to a subrate service provider service.

Table 2 lists the LAN extension options for each type of DCI.

Table 2. Cisco LAN Extension Solutions and Platforms for Various Transport Options

Type of Interconnect (Transport Option)	Cisco LAN Extension Solutions	Cisco Platform
Dark Fiber or Dense Wave-Division Multiplexing (DWDM)	Virtual switching system (VSS)	Cisco Catalyst® 6500 Series Switches
	Virtual PortChannel (vPC)	Cisco Nexus™ 7000 Series Switches
	Crossponder	Cisco ONS 15454 crossponder
Multiprotocol Label Switching (MPLS)	Ethernet over MPLS (two data centers)	• Cisco Catalyst 6500 Series Shared Port Adapter (SPA) Interface Processor 400 (SIP-400) and SIP-600 • Cisco ASR 1000 Series Aggregation Services Routers
	Virtual Private LAN Service (VPLS) (multiple data centers)	• Cisco Catalyst 6500 Series SIP-400 and SIP-600
IP	Ethernet over MPLS (EoMPLS) over Generic Routing Encapsulation (GRE) (2 data centers)	• Cisco Catalyst 6500 Series SIP-400 • Cisco ASR 1000 Series
	VPLS over GRE (multiple data centers)	• Cisco Catalyst 6500 Series SIP-400

Additional services such as encryption and HQoS can be implemented on the Cisco hardware listed in Table 3.

Table 3. Cisco Solutions and Platforms for Additional Services for LAN Extension Schedules

Service	Solution	Platform
Encryption	IEEE 802.1AE	Cisco Nexus 7000 Series
	IPsec	• Cisco Catalyst 6500 Series SPA Services Card 600 (SSC-600) and VPN Services Port Adapter (VSPA) • Cisco ASR 1000 Series
Multilevel QoS	HQoS	• Cisco Catalyst 6500 Series SIP-400 and SIP-600 • Cisco ASR 1000 Series

For more information about Cisco LAN extension solutions for data center interconnection, refer to http://www.cisco.com/en/US/netsol/ns975/index.html.

Storage Extension Technologies

The availability, scalability, security, and performance of the storage subsystem are of utmost importance to any enterprise. The task of ensuring that all these factors are addressed in a single data center is a daunting task for any storage administrator; extending them across data centers is an even greater challenge, requiring implementation of storage best practices. These factors directly affect application performance, in turn affecting the SLAs of business-critical applications. The Cisco MDS 9000 Family of SAN switches is especially suited to these SAN topologies. Table 4 summarizes the features that can be used to address the requirements for storage across data centers.

Table 4. Cisco SAN Extension Solutions

Feature	Requirements	Functions
Virtual SAN	Isolation and security	The VSAN technology provides secure hardware-based network segmentation, similar to the VLAN technology that is widely deployed in LANs. Fabric services such as zoning and routing are independent per VSAN. In this validated solution, the nodes in each VMware ESX cluster are placed in a dedicated VSAN, to use a consolidated physical infrastructure and to be isolated with respect to security threat and fabricwide errors.
	Management and access control	Cisco MDS 9000 NX-OS Software management offers several levels of role-based access control (RBAC). This feature allows an administrator to be in charge of a specific VSAN without having any visibility into other VSANs. The administrator can map the roles defined in the VMware vCenter; for instance, an administrator may be able to access a specific VSAN and the corresponding VMware ESX cluster and nothing else.
Inter-VSAN Routing (IVR)	Isolation and security	In a DCI solution, each data center can implement independent VSANs, preserving the fabric services segmentation, data isolation, and administration independence. IVR allows selected devices from different VSANs, even across different data centers, to communicate without any fabric merging. In this validated solution IVR provides connectivity between the VMware ESX servers located in the secondary data center and the storage located in the primary data center (shared storage). IVR can also provide connectivity to execute VMware Storage VMotion across data centers and to perform primary-array-to-secondary-array storage replication.
SAN	Integrated solution	The capability to plug long-wave and Coarse Wavelength

Feature	Requirements	Functions
extension with dark fiber		Division Multiplexing (CWDM) optics into the Cisco MDS 9000 Series Switches simplifies SAN extension over dark fiber. The performances are guaranteed by the extended buffer-to-buffer credits available with the Cisco MDS 9000 Series.
	Security	Cisco MDS 9000 Series Switches provide Cisco TrustSec Fibre Channel Link Encryption to secure SAN extension data across native Fibre Channel links
SAN extension with FCIP	Integrated solution	Cisco MDS 9000 Series Switches provide Gigabit Ethernet interfaces and support the FCIP protocol, to transparently extend the SAN over an IP network.
	Security	The Cisco MDS 9000 Series provides native IP Security (IPsec) encryption to secure FCIP links.
Port channeling	Availability	Cisco MDS 9000 Series PortChannels are the aggregation of multiple physical Fibre Channel or FCIP links into one logical link, to provide higher aggregated bandwidth, load balancing, and link redundancy.
I/O acceleration (IOA)	Application performances	IOA is an intelligent distributed fabric service built into Cisco MDS 9000 Series Switches. IOA accelerates I/O performance across distances. This feature helps the overall application performance remain relatively the same, even when the application server and the storage are separated by considerable distance. In this validated solution, I/O performance has been enhanced over the FCIP link.

All the features listed in Table 4 make a Cisco MDS 9000 Family SAN resilient and highly available. More information about the Cisco MDS 9000 Series Switches can be obtained from http://www.cisco.com/en/US/products/hw/ps4159/ps4358/prod_white_papers_list.html.

Solution Reference Architecture and Validation

To prove the validity of the solution, VMware and Cisco configured the solution as shown in Figure 3, simulating the WAN and migrating a live application across the data centers without any downtime.The configurations of the VMware vSphere server, the LAN within the data centers, the SAN, and the WAN have been designed to enable VMware VMotion across data centers while adhering to the VMware VMotion requirements.

The solutions described here have been jointly validated in the VMware and Cisco Joint Solutions Lab. The validated topology used for the testing is shown in Figure 3.

Figure 3. Jointly Validated Architecture

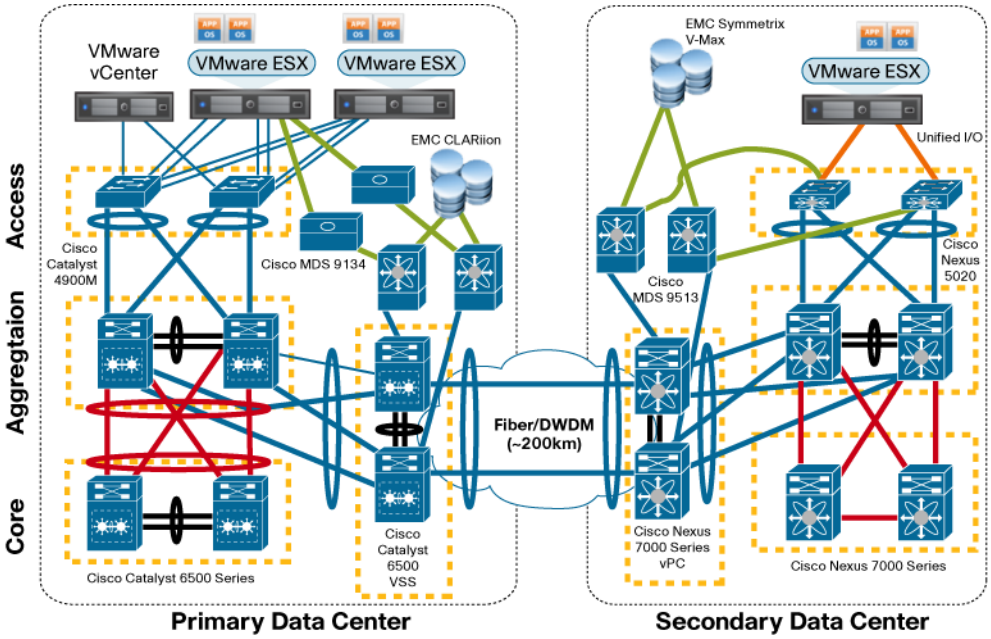

The network topology used in the joint solution test simulates two data centers extended over different types of DCIs. Data center 1 is a Cisco best-practices three-tier architecture with Cisco Catalyst 6500 Series Switches forming the core and aggregation layers and a pair of Cisco Catalyst 4900M Switches forming the access layer. VSS technology is used to provide the extended (or stretched) VLAN. Data center 2 is also designed using the Cisco three-tier architecture, but with the Cisco Nexus 7000 Series Switches and the Nexus 5000 Series Switches. The functional architecture at both data centers is similar, with VSS and vPCs used to provide the extended VLAN.

The storage for the solution is provisioned using either of two methods depending on the test being performed:

- **Shared storage:** Storage is located in data center 1, and the SAN is extended to data center 2 using FCIP SAN extension. FCIP IOA is enabled to help ensure that the application performance does not suffer when the application is accessing its storage across the DCI switches.
- **Active-passive storage using VMware Storage VMotion:** Storage is provisioned at both locations. The storage capacity is provisioned identically at both data centers, and the storage is presented to the VMware ESX servers as unique data stores. VMware Storage VMotion migration is performed from one data store to the other; later, the VMware VMotion virtual machine migration is performed. The storage must be available at all times to both the source and destination VMware vSphere servers for VMware VMotion migration to be successful. The SAN extension method used is again FCIP with IOA enabled.

Test Topology and Tools

The solution was validated using real-life application servers migrating across data centers while clients were accessing these applications. The applications used were Microsoft SQL Server and Microsoft Exchange Server. Table 5 lists the configurations and the test tools used.

Table 5. Cisco and Vmware Vmotion Migration Across Data Center: Test Topology Summary

Application	Server Configuration	Stress-Generation Tool	Application Performance Metrics	Description
Microsoft SQL Server 2005 (64-bit)	CPU: 4 virtual CPUs (vCPUs) Memory: 8 GB Storage: Approximately 300 GB (EMC CLARiiON) OS: Microsoft 2008 64-bit server	Dell DVD Store open source benchmark	Orders per minute (OPM)	The DVD Store benchmark is an online transaction processing (OLTP) benchmark that simulates the operation of a DVD store. Performance is measured in OPMs, indicating the number of orders successfully inserted into the database per minute.
Microsoft Exchange Server 2007	CPU: 4 vCPUs Memory: 8 GB Storage: Approximately 300 GB (EMC Symmetrix V-Max) OS: Microsoft 2003 32-bit server	Microsoft LoadGen	Mail messages being handled by Microsoft Exchange Server	LoadGen is a Microsoft tool developed to closely mimic the operation performed by Microsoft Outlook clients on a Microsoft Exchange Server. LoadGen is a standard test tool used to perform Microsoft Exchange Server load and stress tests.

Test Methodology

Microsoft SQL Server Test

- Reinitialize the Microsoft SQL Server by rebooting the VMware ESX server on which it is resident and the target VMware ESX server to reset the statistics data.
- Start the DVD Store client on a virtual machine that has IP connectivity to both VMware ESX servers.

- Run the DVD client and wait for 30 minutes for the client to attain a steady state; note the OPM on that VMware ESX server.
- Migrate the system to the corresponding target.
- Wait for 30 minutes for the client to attain steady state; note the OPM for that VMware ESX server.
- Perform 18 more migrations with a 10-minute wait between each migration.
- Collect test statistics to evaluate the total elapsed time.

Data Center Evacuation Test

1. Start the LoadGen and DVD Store clients on the respective client virtual machines.
2. Start the migration of the application servers in a sequential manner.
3. Collect test statistics to evaluate the total time required to migrate the data center.
4. Migrate the application servers simultaneously.
5. Collect statistics to evaluate the total elapsed time.

Test Results

The goal of the joint testing is to measure the duration of the overall VMware VMotion migration process during VMware VMotion migration. The tests measure the time taken for the overall VMware VMotion migration to be complete, beginning from the initiation of the VMware VMotion process. The overall migration time is an important measure, and it becomes critical when multiple VMware VMotion migrations are being performed. The amount of time for the overall migration depends on the duration of each VMware VMotion migration. The duration of a VMware VMotion migration largely depends on the distance between the source and destination VMware ESX servers and the amount of bandwidth available between the data centers.

The application used to validate the solution is an e-commerce suite hosted on a Microsoft SQL Server 2005. DVD Store Version 2 (DS2) is a complete online e-commerce test application with a back-end database component, a web application layer, and driver programs. The virtual machine hosting the back-end Microsoft SQL Server database is migrated across the data centers, and the performance of the application in OPMs is captured.

Figure 4 shows the VMware VMotion migration times as a virtual machine is migrated from one VMware ESX server to another, with the servers separated by different distances. The elapsed time increases with distance, but it is a directly related to the network latency and the amount of network bandwidth available for VMware VMotion as the distance increases. In the test scenario, the application client maintained all the sessions, and a momentary drop in performance was observed before the performance returned to steady-state values.

Figure 4. Duration of Intra– and Inter–Data Center Vmware Vmotion Migration

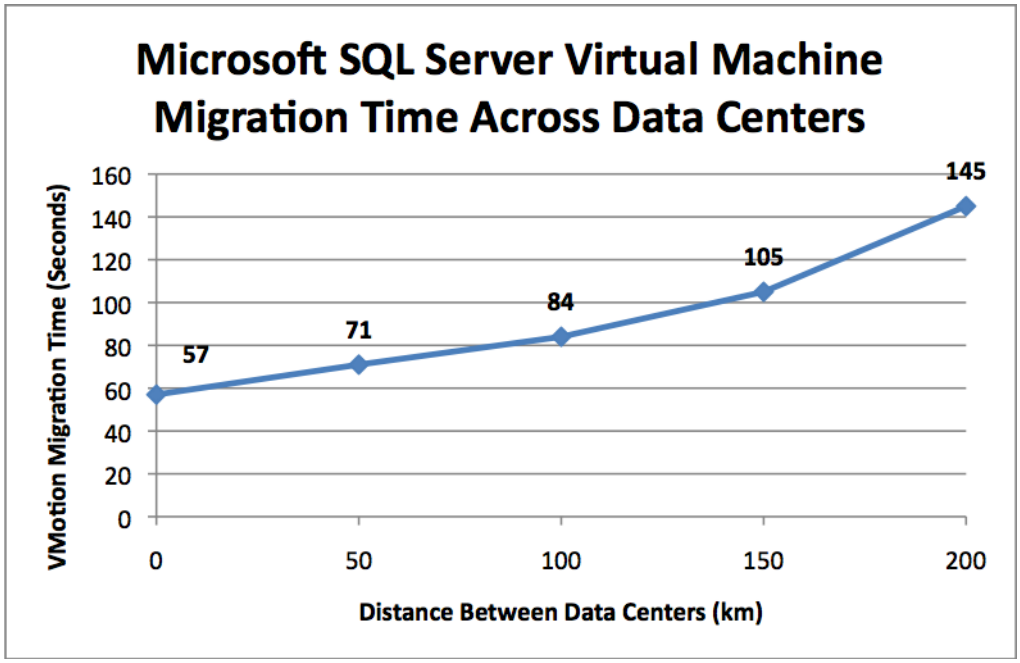

The performance of the application when the back-end database is moved away from the users to a data center 200 km away is shown in Figure 5. The results show that the migration to a remote data center is feasible not only in terms of the mobility of the application but also from a business perspective since degradation is less than 15 percent. Shared storage for the application therefore is a viable option; with advanced SAN extension technologies such as FCIP and IOA, application performance is maintained well within the SLA of the application.

The performance graph shows the performance of the application within the local data center, in which the storage is located. The graph then shows the performance of the application with increasing distance as it migrates to the remote data center, demonstrating that performance is within the range of the SLA, which makes the SAN topology of shared storage for application workload balancing a viable option.

Figure 5. Microsoft SQL Server DVD Store Performance

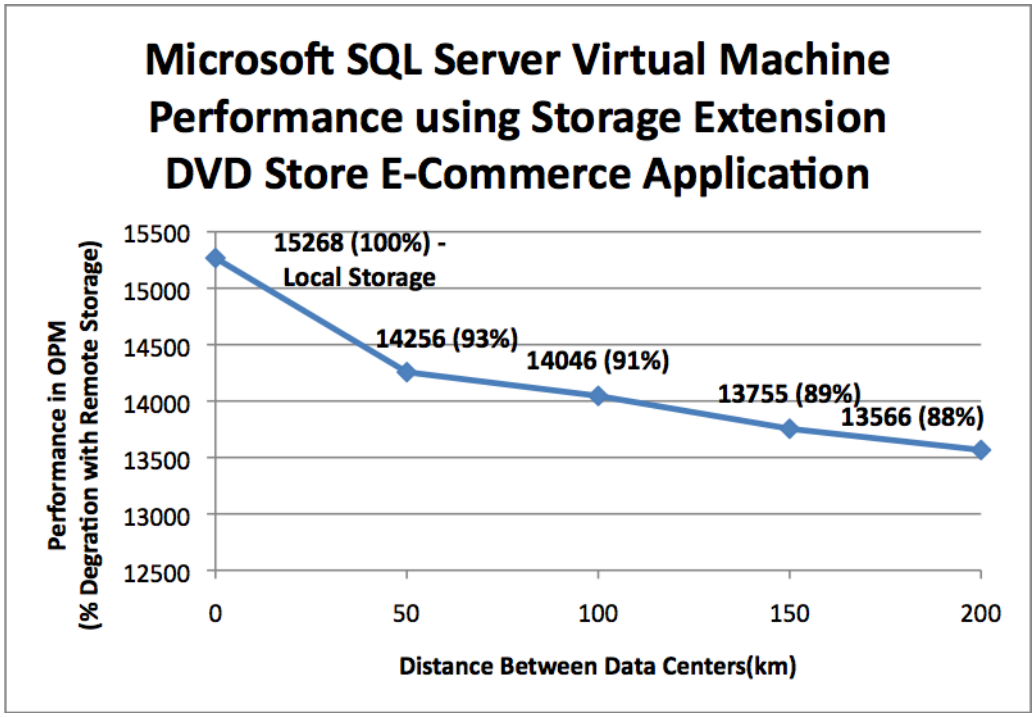

In addition to migrating one mission-critical application across the distance of 200 km, the joint testing migrated more applications in the data center to validate evacuation of the data center for disaster avoidance (Figure 6). The joint testing included a Microsoft Exchange Server 2007 with 1000 Microsoft Outlook 2007 LoadGen users. The application performance monitored by LoadGen is shown in Figure 7. The Microsoft Exchange Server is a four-vCPU virtual machine with 8 GB of memory; resource utilization was 80 to 90 percent for the CPU, with approximately 20 Mbps of data being read and written to the disks. Simultaneously, the Microsoft SQL Server running the DVD Store database was also migrated. The OPM values for the DVD Store application are shown in Figure 5. Figure 6 shows the migration times when the two applications are migrated simultaneously. The results clearly show that the elapsed time for the Microsoft SQL Server increases by a small amount, which is acceptable since the VMware VMotion network is now being shared with an extremely busy Microsoft Exchange Server workload. In spite of this increase, there was no perceivable effect on the clients performing the benchmark.

Figure 6. Migration Times with Multiple-Application VMware VMotion Migration

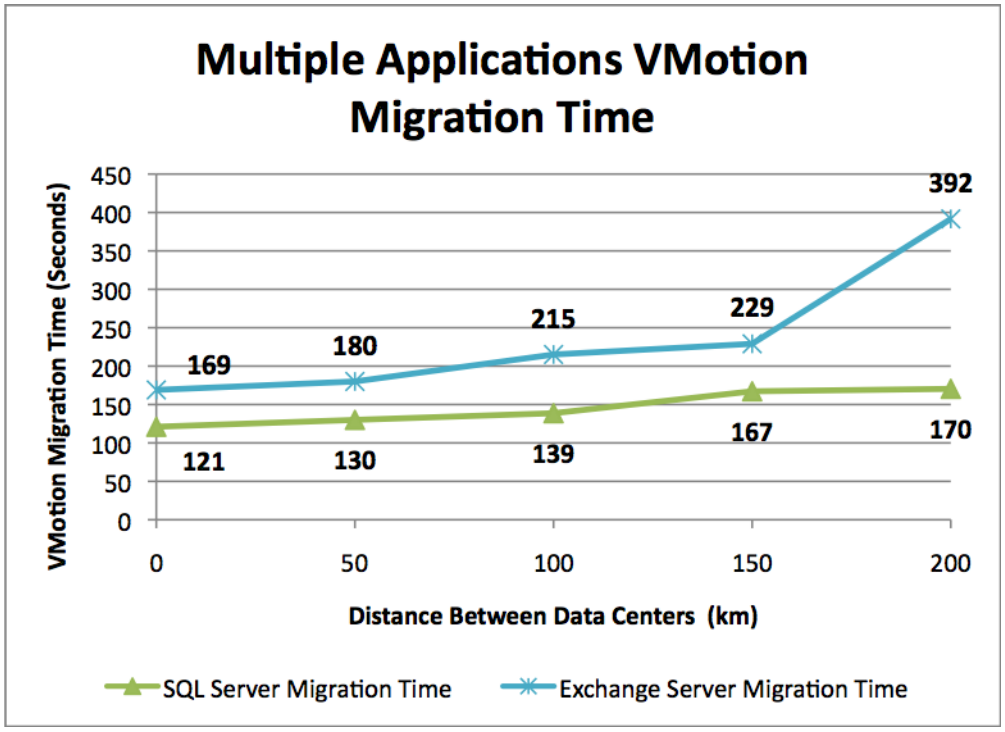

Figure 7. LoadGen Results Summary

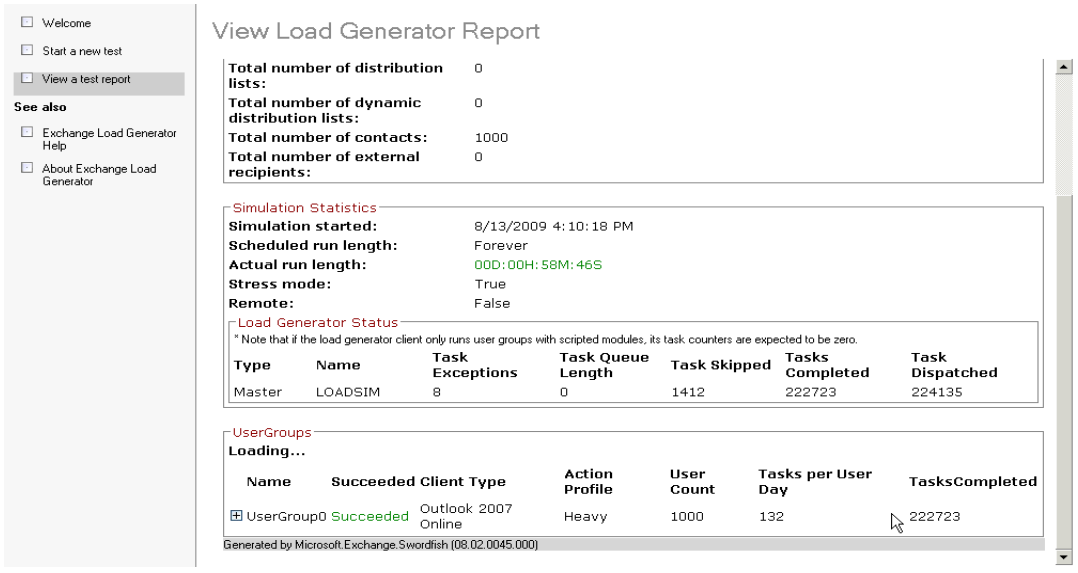

Recommended Operational Procedure

The recommended procedure for implementing the joint VMware and Cisco solution is to have the VMware vSphere high-availability clusters independent of each other in the two data centers. VMware VMotion migration across data centers should be a manually instantiated task to keep VMware Dynamic Resource Scheduling (DRS) from automatically moving virtual machines across data centers.

Conclusion

VMware VMotion enables data centers to transparently implement virtual machine mobility using the Cisco LAN and storage extension solutions. The VMware vSphere Virtual Data Center Operating System (vDCOS) with the suite of features bundled with vSphere allows customers to transparently migrate or evacuate data center applications with no downtime from a user perspective. This technological capability gives IT departments tools to redefine the business continuance and disaster recovery plans of the enterprise. The need for a more complex and expensive solution to meet the recovery-point objective (RPO) of the business continuance plan can be reduced with the use of VMware VMotion. The optimal architecture of the underlying transport infrastructure—the IP network and SAN—enhances the solution to bring the RPO to near zero, and also reduces the recovery time objective (RTO) to a very small number. The joint Cisco and VMware solution gives IT departments a very powerful tool for better provisioning, utilizing, and maintaining a virtualized data center with resources spread across multiple physical locations.

For More Information

- VMware VMotion: http://www.vmware.com/products/vi/vc/vmotion.html
- Data Center Interconnect (DCI): Layer 2 Extension Between Remote Data Centers: http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11_493718.html
- Cisco Catalyst 6500 Series Switches: http://www.cisco.com/go/6500
- Cisco Nexus 7000 Series Switches: http://www.cisco.com/go/nexus7000
- Cisco MDS 9000 Family: http://www.cisco.com/go/mds

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

VMware, Inc
3401 Hillview Ave
Palo Alto, CA 94304
USA
www.vmware.com
Tel: 1-877-486-9273 or 650-427-5000
Fax: 650-427-5001

(0812R) C11-557822-00 08/09