

Enhance Business Continuance with Application Mobility Across Data Centers

What You Will Learn

With the continuing adoption of server virtualization technology, customers want to use the flexibility provided by virtual machines (VMs) to improve application availability. How to move applications within a server cluster or data center is well understood today. However, transparent mobility across data centers requires architectures that balance workload mobility, disaster avoidance, and disaster recovery objectives while coordinating workload and storage movements across a routed network core using flexible LAN extensions.

This document describes an application mobility architecture that brings together critical products and technologies from Cisco, NetApp and VMWare to deliver a robust, scalable, and resilient solution.

VMware has been the industry leader in virtualization technologies for the past decade and has brought to the data center critical technologies such as VMware vSphere VMotion for virtual machine mobility and Site Recovery Manager (SRM) for disaster recovery.

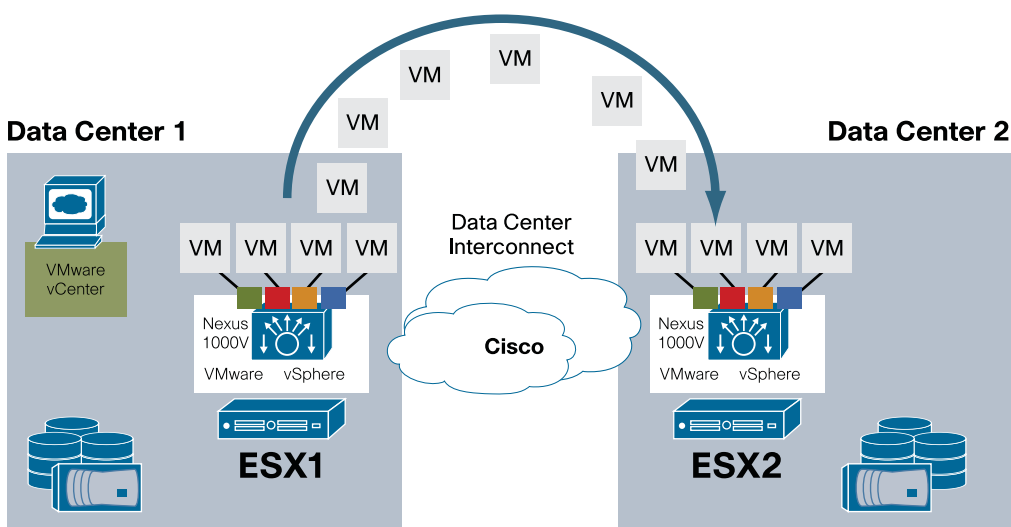
Cisco has been the industry leader in networking technologies. The Cisco® switching and routing portfolio enables a robust foundation for intelligent network connectivity within and across data centers. Cisco is offering a range of data center interconnect (DCI) technologies to address the need for LAN extensions across a variety of transport networks.

NetApp is the industry leader in IP-based and multiprotocol storage technologies. The NetApp family of unified storage products provides highly efficient, available, and resilient shared storage for VMware environments.

Cisco, NetApp and VMWare Migration Solution

The 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 and a Cisco Nexus® 1000V Switch in each data center; a VMware vCenter server to manage all virtual machine migrations; and Cisco products and technologies, in particular, Cisco Overlay Transport Virtualization (OTV) technology, to enable data center interconnect (DCI) across an IP network, connecting the data centers. NetApp storage provides data availability at data centers 1 and 2. In this configuration, applications provisioned on the VMware vSphere server can be migrated across the data centers or a private cloud with no application downtime. Migration is granular at the virtual machine level.

Figure 1. VMware VMotion Across Data Centers





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 in 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.

NOTE: This document does not address disaster recovery in the event of a data center outage. Cisco, NetApp and VMware provide comprehensive disaster recovery solutions, and these solutions are discussed in other documents.

VMware VMotion Requirements

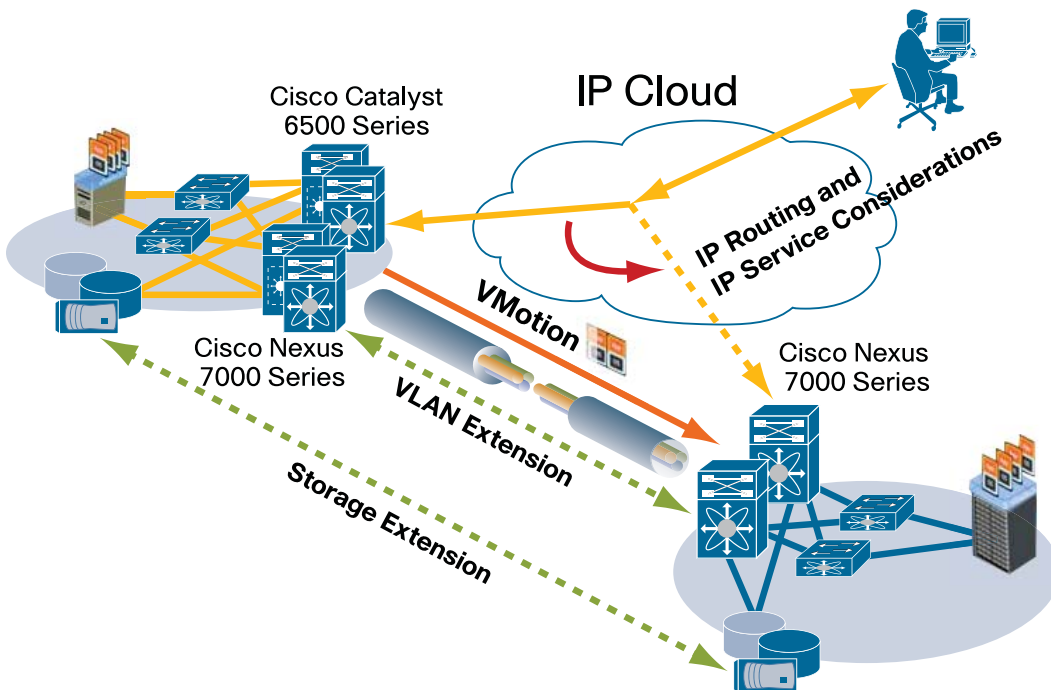
VMware VMotion application mobility is based on certain infrastructure requirements:

- A minimum bandwidth of 622 Mbps is required between data centers and at least 1-Gbps links within a data center. The source and destination VMware ESX servers must be on the same IP subnet and broadcast domain.
- The maximum round-trip latency between the source and destination VMware ESX servers cannot exceed 5 milliseconds. Based on the speed of light over fiber and certain guard bands for network delays, a maximum distance of 400 km is supported today.
- 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. If servers are present in two distinct locations, the sets of data must be identical.
- Access from VMware vCenter, the VMware Virtual Infrastructure (VI) management GUI, to both the VMware ESX servers must be available to accomplish the migration. This implies that a single VMware vCenter server should span both data centers.

Challenges with Virtual Machine (VM) Mobility Across Data Centers

Successful VMware VMotion migration across data centers in different physical locations poses certain challenges. VMware VMotion migration across long distances requires careful evaluation of data center network and storage design as shown in Figure 2.

Figure 2. Infrastructure Challenges with VM Mobility Across Data Centers



These challenges are:

- VLAN extension:** Layer 2 VLANs must be extended across data centers without compromising the availability, resilience, and security that exists within a single physical location. Layer 2 domain elasticity should be possible over different connection media such as dark fiber or IP and Multiprotocol Label Switching (IP/MPLS)-based infrastructure. The obvious challenge in extending VLANs across data centers is the possibility of increased risk from expanded spanning-tree domains, the possibility of loops if spanning tree is isolated across data centers, and packet broadcast and flooding in multiple sites. However, some broadcast traffic is essential for application-level communication. Layer 2 domain elasticity should thus help ensure a loop-free topology while isolating spanning tree across data centers, the scalability to connect multiple data center sites, and optimal use of WAN bandwidth.
- Storage extension:** The availability of virtual machine data on storage devices with read and write access to the two VMware ESX servers is critical to a successful application migration. The WAN design should intelligently and optimally manage the large data sets associated with applications. Storage network designs and storage systems should take into account these parameters to help ensure that data is not only available, but that it is secure, with I/O latencies and performance that will not affect the SLAs of the applications. Application bandwidth and latency requirements will partially determine the storage architecture: for example, workflows that consist primarily of data read operations are particularly easy to facilitate using caching mechanisms.



- **IP routing and IP service considerations:** An application migrated across data centers using VMware VMotion maintains its existing IP and MAC addresses. If the traffic to the virtual machine originates in the same Layer 2 domain, then the Layer 2 extension will suffice. If the traffic to the virtual machine is traversing a Layer 3 network, such as an IP cloud or the Internet, then the traffic needs to be rerouted to the new data center location. Existing application sessions may continue to be routed through the existing data center due to specific or existing IP service requirements, such as firewalls. Because of this behavior, the following IP routing considerations are required:
 - **Routing from remote clients to application servers:** Requires intelligent routing-based or Domain Name System (DNS)-based mechanisms to adapt to IP mobility
 - **Routing from application servers to remote clients:** Requires forwarding of application traffic to the appropriate default gateway (preferably in the local data center pod) to achieve optimal routing as well as symmetrical routing for IP services such as firewalls
- **WAN characteristics:** The WAN bandwidth and latency requirements for VMware VMotion are critical factors in a successful application VMotion migration across data centers.

Infrastructure Architectures for Application Mobility

The application mobility solution can be designed in multiple ways based on LAN and storage extensions. Some of the possible ways are:

Extended VLAN with Simple Storage Extension (VMware VMotion with Shared Storage)

An extended VLAN and shared storage architecture with simple VMware VMotion migration extends the VLAN between the two sites, but with storage remaining at the original location (a single array). 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. This design requires consideration to help ensure that the application in the remote data center will be able to tolerate the I/O latency of the remote storage.

Extended VLAN with VMware VMotion and Storage VMotion (Active-Passive Storage)

An extended VLAN with VMware Storage VMotion topology requires migration of data to the remote data center prior to migration of the virtual machine itself to the remote data center. 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.

Extended VLAN with Intelligent Data Management (Active-Cached Storage)

To mitigate the performance problems associated with simple shared storage and active-passive architectures, a cached architecture can be used. With a cached model, performance concerns that arise when an application accesses remote storage are partially or fully mitigated. In addition, cached architectures can allow administrators to migrate applications and their corresponding virtual machines using VMware VMotion prior to a VMware Storage VMotion event. Caching is on a per-file basis and provides virtual machine-level granularity. Caches are of zero size until files are accessed and the cache is filled.

Cached architectures allow applications to read data from the cache on a local device while the canonical data copy is maintained on a remote-storage device. This approach is particularly useful in cases in which long-distance VMware VMotion migration is a temporary measure and a virtual machine is scheduled to be returned to its original data center shortly after transfer.



Cisco, NetApp and VMWare Joint Solution

The Cisco, NetApp and VMWare solution, jointly engineered by the three companies, addresses a VMware VMotion migration with intelligent data management. The components of the solution are:

- VMware VMotion technology
- Cisco OTV LAN extension technology
- Distributed virtual switching
- Routing optimization
- NetApp FlexCache storage caching technology for seamless transition of virtual machine mobility

VMware VMotion Technology

VMware VMotion enables the live migration of running virtual machines from one physical server to another with zero downtime, continuous service availability, and complete transaction integrity.

VMware VMotion migration is achieved when the active memory and precise execution state of a virtual machine is rapidly transmitted over a high-speed network from one physical server to another and access to the virtual machines disk storage is instantly switched to the new physical host. Since the network is also virtualized by VMware vSphere, the virtual machine retains its network identity and connections, helping ensure a transparent migration process. VMware VMotion is a crucial enabling technology for the creation of a single highly available virtual data center spanning geographically disparate data centers.

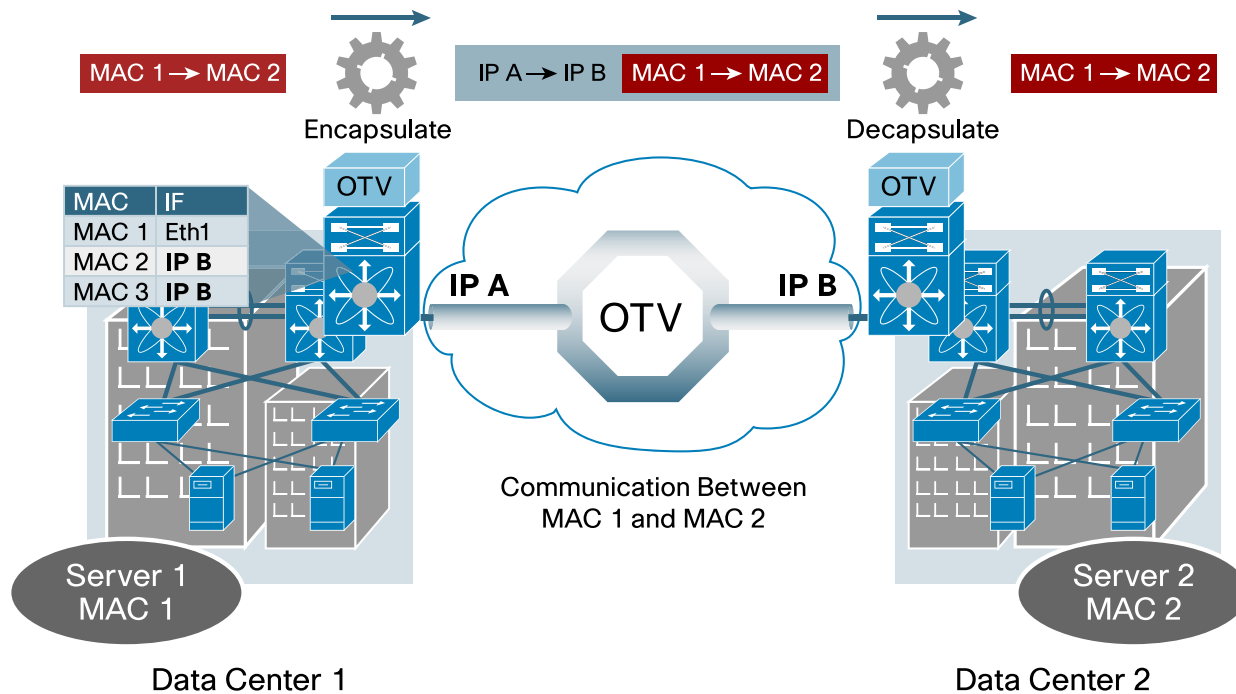
Cisco Overlay Transport Virtualization LAN Extension Technology

Cisco OTV technology provides an operationally optimized solution for the extension of Layer 2 connectivity across any transport. OTV is therefore critical to the effective deployment of distributed data centers to support application availability and flexible workload mobility. OTV is a “MAC in IP” technique. By using the principles of MAC address routing, OTV provides an overlay that enables Layer 2 connectivity between separate Layer 2 domains while keeping these domains independent and preserving the fault-isolation, resiliency, and load-balancing benefits of an IP-based interconnection.

OTV uses a control protocol to map MAC address destinations to IP next hops that are reachable through a routed network core. OTV can be thought of as MAC address routing, in which the destination is a MAC address, the next hop is an IP address, and traffic is encapsulated in IP so it can simply be carried to its MAC address routing next hop over the core IP network. Thus, a flow between source and destination host MAC addresses is translated in the overlay into an IP flow between the source and destination IP addresses of the relevant OTV edge devices. This process is referred to as encapsulation rather than tunneling because the encapsulation is imposed dynamically and tunnels are not maintained. Since traffic is IP forwarded, OTV is as efficient as the core IP network and will deliver optimal traffic load balancing, multicast traffic replication, and fast failover just like the core would.

Figure 3 illustrates this dynamic encapsulation mechanism.

Figure 3. OTV Operation



OTV provides the following benefits:

- **Transport agnostic:** OTV is IP encapsulated and can therefore use any core capable of forwarding IP traffic. OTV therefore does not pose any requirements for the core transport.
- **High availability:** OTV preserves the failure boundary and site independence: OTV does not rely on traffic flooding to propagate reachability information for MAC addresses. Instead, a control protocol is used to distribute such information. Thus, flooding of unknown traffic is suppressed on the OTV overlay, Address Resolution Protocol (ARP) traffic is forwarded only in a controlled manner, and broadcasts can be forwarded based on specific policies. Spanning-tree Bridge Protocol Data Units (BPDUs) are not forwarded at all on the overlay. The result is failure containment comparable to that achieved using a Layer 3 boundary at the Layer 2 domain edge. Sites remain independent of each other, and failures do not propagate beyond the OTV edge device. The loop prevention mechanisms in OTV prevent loops from forming on the overlay, and also prevent loops from being induced by sites when these are multihomed to the overlay
- **Full WAN bandwidth utilization and optimal multicast replication:** When sites are multihomed, OTV provides the capability to actively use multiple paths over multiple edge devices. This capability is crucial to keeping all edge devices active and thus optimizes the use of available bandwidth. OTV uses the IP-multicast capabilities of the core to provide optimal multicast traffic replication to multiple sites and avoid head-end replication that leads to suboptimal bandwidth utilization.
- **Transparent to the sites:** OTV extensions do not affect the design or protocols of the Layer 2 sites they interconnect. Interconnection is as transparent as connection of a router to the Layer 2 domain and therefore does not affect the local spanning tree or topology.



Distributed Virtual Switching

VMware VMotion moves virtual machines running applications across physical ports and data centers; however, the network provisioning, security, and management policies as well as visibility at the virtual machine level must be available. With current hardware switches, it is impossible to view or apply network policy to locally switched traffic between virtual machines at the virtual network interface card (vNIC) level. For example, it is not possible to correlate traffic on the same physical link from multiple virtual machines.

VMware and Cisco jointly developed the concept of a distributed virtual switch (DVS), which essentially decouples the control and data planes of the embedded switch and allows multiple, independent virtual switches (data planes), called virtual Ethernet modules (VEMs) to be managed by a centralized management system (control plane), called a virtual switch module (VSM).

The Cisco Nexus 1000V virtual switch (vSwitch) has the following strengths:

- It is the industry's first DVS for VMware.
- It overcomes network challenges and accelerates server virtualization.
- It is compatible with all switching platforms. The Cisco Nexus 1000V maintains the existing VMware vCenter provisioning model for server administration while allowing network administration of the virtual network using the Cisco NX-OS Software command-line interface (CLI).
- It allows server teams to offload vSwitch responsibility to the network teams, helping ensure proper network connectivity and security. Network teams get virtual machine-level visibility, NetFlow, Encapsulated Remote Switched Port Analyzer (ERSPAN), and port statistics that continue through VMware VMotion migration.

Routing Optimization

Routing optimally to a virtual machine migrated by VMware VMotion is critical because the virtual machine maintains its existing IP and MAC addresses as it is moved across data centers. The Layer 3 and 2 reachability design of the data center should accommodate this behavior.

If the traffic to the virtual machine originates in the same Layer 2 domain, then the Layer 2 extensions will suffice for connectivity across pods or data centers. OTV facilitates this design. In the example in Figure 4 later in this document, in Data Center 1, Layer 2 traffic requiring reachability from the Cisco 6500 Catalyst Series pod to the Cisco Nexus 7000 Series pod will be switched using Layer 2. No additional configuration is needed.

If the traffic to the virtual machine is traversing a Layer 3 network, such as an IP cloud or the Internet, then the traffic needs to be rerouted to the new data center location. Existing application sessions may continue to be routed through the existing data center due to specific or existing IP service requirements, such as firewalls.

The following IP routing considerations are required:

- **Routing from remote clients to application servers:** This consideration can be addressed using the Cisco Global Site Selector (GSS) and the Cisco Application Control Engine (ACE). The remote client will perform a DNS-based lookup for the application server (Microsoft SQL Server in this case) on the Cisco GSS in the IP WAN. The Cisco GSS will respond to the DNS lookup with the virtual IP address of the Cisco ACE in the data center containing the Microsoft SQL Server after VMware VMotion migration. This capability is made possible through the workflow integration and tight coupling of Cisco GSS and ACE and VMware vCenter.



- **Routing from application servers to remote clients:** Application traffic must be forwarded to the appropriate default gateway preferably in the local data center pod to achieve optimal routing as well as symmetrical routing for IP services such as firewalls. The Hot Standby Router Protocol (HSRP) default gateway pointing from the Microsoft SQL Server to the aggregation-layer switch can be the local switch. This design is achieved by having an identical HSRP default gateway address in both data centers. HSRP localization techniques can be used to filter HSRP keepalives across the Layer 2 VLAN extension between data centers. The HSRP default gateway will forward traffic to the appropriate Cisco ACE device for processing of source Network Address Translation (NAT) or firewall services. This capability can be achieved by deploying the Cisco ACE in Layer 3 one-arm mode with source NAT. This design maintains symmetrical routing from an IP services perspective.

NetApp FlexCache Storage Caching Technology for Seamless Transition of Virtual Machine Mobility

Application mobility through VMware VMotion migration poses inherent challenges on the performance of storage data in a location remote from the applications that access the data. Because of WAN latency and bandwidth constraints, performance typically degrades when data is accessed over long distances. Even in extraordinarily high-bandwidth deployments, latency can significantly degrade storage performance.

Caching data using technologies such as NetApp FlexCache software mitigates WAN latency and bandwidth constraints by creating locally available copies of remote data. With a storage cache, frequently accessed data is cached on a storage system that is local to the application virtual machine, reducing latency. In NetApp storage systems, NetApp FlexCache software is a component of the unified storage architecture that allows ease of management across all storage systems, regardless of protocol. Table 1 shows NetApp's elegant and efficient storage caching.


Table 1. NetApp Storage Caching and Performance Enhancement Technologies

Feature	Requirements	Functions
NetApp FlexCache Software	Performance enhancement	<ul style="list-style-type: none"> • Enables application mobility by providing data availability over long distances with minor performance degradation.
	Efficiency	<ul style="list-style-type: none"> • NetApp FlexCache software provides storage efficiency by transferring to the cache only the data that is actually accessed by the application virtual machine. As a result, the caching system uses less storage than full replication and provides significant cost savings when compared to alternatives that use full data copies to relieve storage bottlenecks.
	Data coherency	<ul style="list-style-type: none"> • NetApp FlexCache software implements delegations to improve performance and help guarantee cache coherency to shared files. • A write-through policy helps ensure that data written through the cache is proxied to the source data, thereby helping ensure that subsequent accesses to that data from other virtual machines do not read stale data.
	Demand based	<ul style="list-style-type: none"> • NetApp FlexCache software automatically stores blocks of data read by clients and hosts on the cache volume. No data management policies are required to create copies of data at the remote VMware vSphere location.
NetApp Data ONTAP OS	Data management	<ul style="list-style-type: none"> • Unified management is provided in shared, active-passive, and active-cached environments.
	Thin provisioning	<ul style="list-style-type: none"> • Thin-provisioned storage for cache volumes are on the same NetApp storage systems already deployed in the data center and use capacity only when needed for VMware VMotion migration. • Cache volumes are available at both sides of a VMware vSphere deployment so virtual machine migration can happen either way at any time.

Reference Architecture and Validation

To prove the validity of the solution, Cisco, NetApp and VMware configured the solution as shown in Figure 4, simulating the WAN and migrating a live application across the data centers without any downtime to the application. The solutions described here have been jointly tested by Cisco, NetApp and VMware.

Figure 4. Cisco, NetApp and VMWare Jointly Validated Architecture



The network topology used in the joint solution test simulates two data centers that are about 400 km apart and connected through an IP network. Data Center 1 is a brownfield data center with one pod made up of Cisco Catalyst 6500 Series Switches and Cisco Catalyst 4900M Switches. The other pod in Data Center 1 is a new deployment with Cisco Nexus 7000 and 5000 Series Switches. Data Center 2 is a greenfield data center with Cisco Nexus 7000 and 5000 Series Switches. Both data centers are using OTV to extend Layer 2 across long distances. In Data Center 1, the pod built with Cisco Catalyst 6500 Series Switches connects to the Cisco Nexus 7000 Series Switches to forward DCI traffic. Both data centers are using virtual device contexts (VDCs) on the Cisco Nexus 7000 Series Switches.

Both data centers use Cisco Catalyst 6500 Series Switches in the data center core layer to connect to the internet. VSS technology is used between pairs of Cisco Catalyst 6500 Series Switches for high availability and to provide Multichassis EtherChannel (MEC) connectivity. Both data centers are built on Cisco's three-tier data center architecture. With OTV, the DCI layer has been combined into the aggregation layer. The data centers use the Cisco ACE Module on the Cisco Catalyst 6500 Series Switches in Data Center 1, and a Cisco ACE appliance for the Cisco Nexus 7000 Series Switches in Data Center 2. Cisco ACE is used in Layer 3 one-arm mode with source NAT.

Each data center pod has a VMWare ESX Server on which a Microsoft SQL Server can be run as a virtual machine. Initially the SQL Server is running in Data Center 1. This SQL Server is migrated between data centers using VMWare vMotion. A SQL client connects to the Microsoft SQL Server over the Internet, using the data center WAN core, to access e-commerce applications (see Test Tools section for application details).



The Cisco GSS cluster in the Data Center WAN is used for DNS based lookup, to get to the SQL Server. Cisco GSS will point to the Cisco ACE device in the data center where the SQL Server currently resides.

The storage for the solution is provisioned using NetApp FAS3100 Series storage systems with NetApp FlexCache intelligent data management.

Test Topology and Tools

The solution has been validated using real-life application servers migrating across data centers while clients access these applications. The applications used were Microsoft SQL Server and Microsoft Exchange Server. Table 2 lists the configurations and the test tools used.

Table 2. Test Tool Configuration

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 (NetApp FAS3100) OS: Microsoft 2008 64-bit server	Dell DVD Store open source benchmark	Orders per minute (OPMs)	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.

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 Dell DVD Store client on a virtual machine that has IP connectivity to both VMware ESX servers.
- Run the Dell DVD client and wait for 30 minutes for the client to attain a steady state; note the operations per minute (OPM) on that VMware ESX server.
- Migrate the system to the corresponding target.
- Wait 30 minutes for the client to attain steady state; note the OPMs 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.

Test Results

The goal of the joint testing is to measure the completion time taken for the overall VMware VMotion migration and the application performance in terms of operations per minute (OPMs) due to VMotion. The overall migration time is an important measure, and it becomes critical when multiple VMware VMotion migrations are being performed. 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 Microsoft SQL Server 2005. Dell 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 5 shows the VMware VMotion migration times as a virtual machine is migrated from one VMware ESX server to another, with servers in data centers separated by different distances. The elapsed time increases with distance. Elapsed time is directly related to the network latency and the amount of network bandwidth available for VMware VMotion. In the test scenario, the SQL client maintained all sessions, and a momentary drop in performance was observed before the performance returned to steady-state values.

Figure 5. Duration of Intra- and Inter-Data Center VMware VMotion Migration

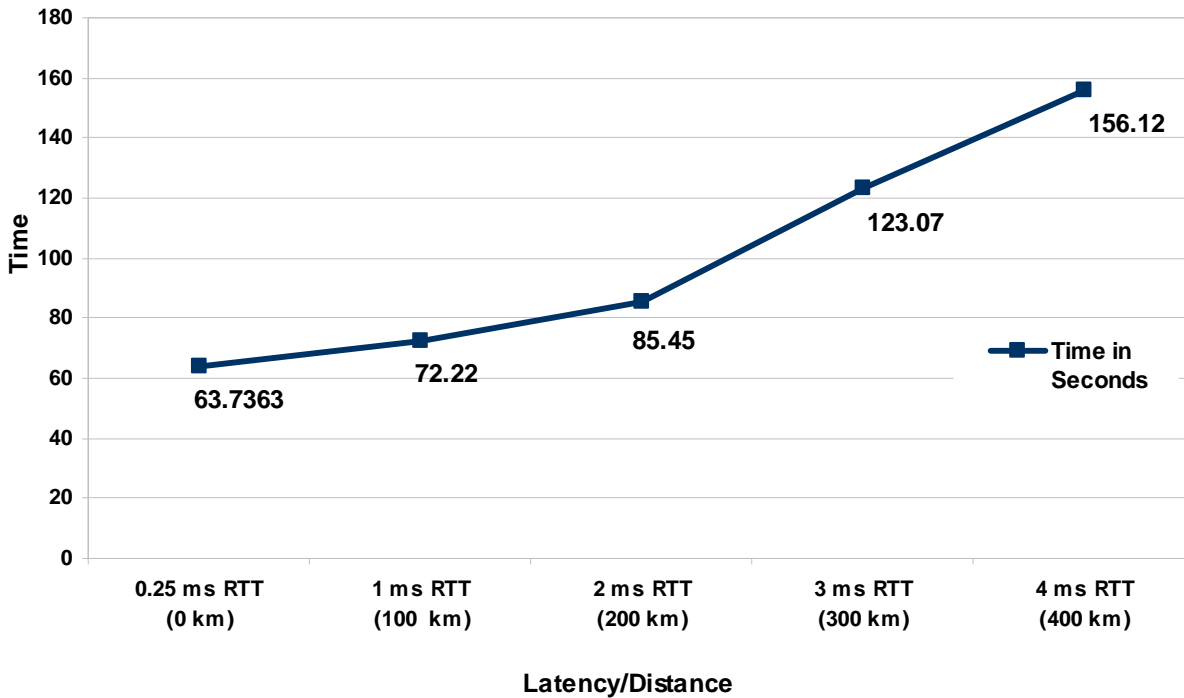
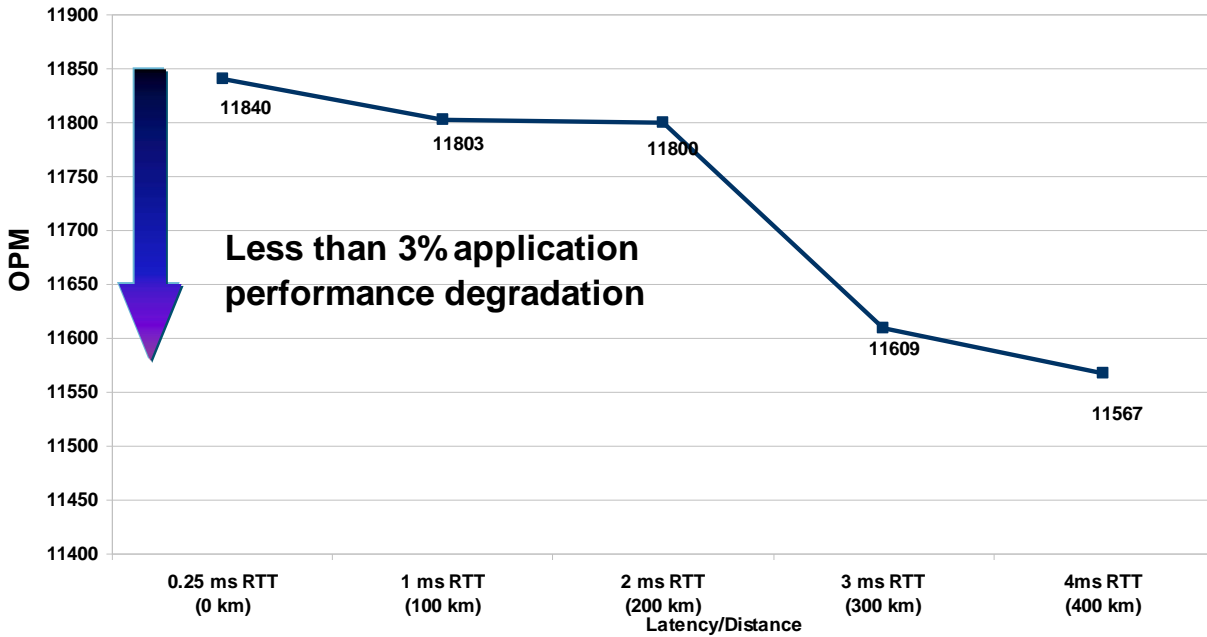


Figure 6 shows application performance, in OPMs, when the back-end database is moved to a remote data center up to 400 km away. The results show that the migration to a remote data center is feasible not only in terms of the mobility of the application but also from a business perspective since degradation is less than 3 percent. Intelligent data management with NetApp FlexCache and long-distance VMware VMotion migration for the application therefore are viable options.

Figure 6. Microsoft SQL Server and Dell DVD Store Performance



Recommended Operational Procedure

The recommended procedure for implementing the joint Cisco, NetApp and VMWare 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

Cisco, NetApp and VMWare jointly tested and verified an application mobility architecture that allows customers to move virtual machine based application workloads between data centers without affecting the application uptime. The described solution uses innovative technologies from the three companies to solve common problems that customers face when trying to extend Layer 2 domains between data centers.

VMware VMotion and SRM are the building blocks for achieving granular mobility of workloads and important recovery mechanisms for helping ensure business continuance. Cisco OTV technology provides a powerful mechanism for easily and flexibly extending the LAN across any type of transport network without requiring a network redesign. The integration of VMware vCenter with Cisco GSS and Cisco ACE delivers crucial route optimization functions. The NetApp FlexCache solution offers an elegant and efficient way to deliver storage access with nearly no performance degradation between data centers.

The joint Cisco, NetApp and VMWare solution gives IT departments a powerful tool for better provisioning, using, 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>
 - <http://www.vmware.com/products/vmotion/>
- VMware SRM: <http://www.vmware.com/products/site-recovery-manager/>
- 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
 - <http://www.cisco.com/en/US/netsol/ns975/index.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 Overlay Transport Virtualization (OTV): http://www.cisco.com/en/US/prod/switches/ps9441/nexus7000_promo.html
- Cisco GSS and ACE products: <http://www.cisco.com/go/ace>
- Cisco Nexus 1000V Switch:
 - <http://www.cisco.com/go/1000veval>
 - <http://www.cisco.com/en/US/products/ps9902/index.html>



Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA

www.cisco.com

Tel: 1-408-526-4000 or 1-800-553-NETS (6387)
Fax: 1-408-527-0883



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
USA

www.netapp.com

Tel: 1-408-822-6000
Fax: 1-408-822-4501



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

© 2010 Cisco, the Cisco logo, and Cisco Systems are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All other trademarks mentioned in this document or Website 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.

© Copyright 2010 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. NetApp, the NetApp logo, Go further, faster, Data ONTAP, and FlexCache are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft and SQL Server are registered trademarks of Microsoft Corporation. VMware and VMotion are registered trademarks of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Copyright © 2010. VMware, Inc. All rights reserved. Protected by one or more U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,944,699, 6,961,806, 6,961,941, 7,069,413, 7,082,598, 7,089,377, 7,111,086, 7,111,145, 7,117,481, 7,149,843, 7,155,558, 7,222,221, 7,260,815, 7,260,820, 7,269,683, 7,275,136, 7,277,998, 7,277,999, 7,278,030, 7,281,102, 7,290,253, 7,356,679 and patents pending.

(1002R)

C11-591960-00 03/10