




Enhanced Business Resilience Using Capacity Expansion and Workload Portability

Last Updated: August 2, 2011



About the Authors



Brian Howard

Brian Howard, Software/Quality Assurance Engineer, Systems Architecture & Strategy Unit (SASU), Cisco Systems

Brian Howard is a Software/QA test engineer in SASU focusing on new and innovative Data Center Interconnect (DCI) technologies. Recent DCI design and test efforts include OTV, A-VPLS, Nexus 1000v, and the Virtual Security Gateway, all components utilized in DCI systems. Brian has been providing quality initiatives and testing in Advanced Services and the Cisco Corporate Development Office for 11 years, focusing primarily on routing and switching and most recently on data center virtualization using DCI.



Jake Howering

Jake Howering, Product Manager, Systems Architecture & Strategy Unit (SASU), Cisco Systems

Jake Howering, a 12 year Cisco veteran, is the Data Center Interconnect (DCI) Product Manager in the Systems Architecture & Strategy Unit. Jake's primary focus includes developing use case solutions leveraging DCI-enabled architectures and driving innovation within DCI with new technologies and products. Previously, Jake developed system architectures in Data Center Metro Ethernet services and core ISP routing as both a Systems Engineer and Advanced Services Network Consulting Engineer.



John Kennedy

John Kennedy, Technical Marketing Engineer, Systems Architecture & Strategy Unit (SASU), Cisco Systems

John Kennedy is a Technical Marketing Engineer in the Systems Architecture & Strategy Unit. Previously, John worked at VMware Inc. for 8-1/2 years as a Senior Systems Engineer supporting channel partners outside the U.S. and serving on the HP Alliance team. He is certified VCP on every version of VMware's ESX/ESXi, vCenter, and Virtual Infrastructure including vSphere 4. He has presented at VMworld, EMC, and CiscoLive in the U.S. and Europe. Currently, John is researching methods of implementing long distance vMotion for use in the Data Center Interconnect Cisco Validated Design.



Kenny Lei

Kenny Lei, Technical Marketing Engineer, Systems Architecture & Strategy Unit (SASU), Cisco Systems

Kenny Lei is a Technical Marketing Engineer in the Systems Architecture & Strategy Unit focusing on data center systems. Kenny has been at Cisco since 1997 and has worked in hardware development, software development, and marketing.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at <http://www.cisco.com/go/trademarks>. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Enhanced Business Resilience Using Capacity Expansion and Workload Portability

© 2011 Cisco Systems, Inc. All rights reserved.



Enhanced Business Resilience Using Capacity Expansion and Workload Portability

Data center administrators are faced with the increasing challenge of doing more with less. Data centers are under pressure from the ever-increasing demand for more compute and storage capacity, while economic pressures mean fewer dollars for staff and equipment. These pressures have led to the widespread adoption of virtualization as a technique for maximizing the efficiency of physical resources for both compute and storage. This trend has supported the recent explosive growth and adoption of cloud computing and has changed the approach that data center administrators are taking to address their growth demands.

The cloud computing and service provider model offers the data center administrator a means of satisfying dynamic bursts in demand without purchasing and building out the physical equipment within their facilities needed to satisfy peak demands. Additionally, the data center administrator needs to be able to migrate workloads within the enterprise private cloud or in a hybrid cloud solution between enterprise and service provider during both planned and unplanned outages.

This document addresses the [Capacity Expansion Use Case](#) and [Virtualized Workload Portability Use Case](#) with combined technologies from EMC[®], Cisco[®], and VMware[®].

Executive Summary

Capacity Expansion

Business Challenge

- Local data center virtual machines are overtaxed and exceeding acceptable resource utilization levels.

Business Solution

- Bursts new transactions to available virtual machines using Cisco Dynamic Workload Scaling.
- Data continuously available with EMC VPLEX[®] Geo replication and caching technologies.

Business Results

- New traffic is burst to new data center, stabilizing local data center resources.
 - Storage content available locally and remotely using EMC VPLEX Geo.
-



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

Copyright © 2011 Cisco Systems, Inc. All rights reserved.

Workload Portability

Business Challenge

- Workload portability is time-consuming and complex.

Business Solution

- UCS Service Profiles simplify server identity management.
- LISP optimizes traffic flow after migration.

Business Results

- UCS Service Profiles and VPLEX Geo minimize new configuration requirements for migration.
- Application downtime is reduced with simplified process and new optimized traffic redirection.

EMC, VMware, and Cisco—Joint Solution at EMC World 2011

EMC, VMware, and Cisco worked together to develop two solutions to address specific use case requirements. This proof of concept effort resulted in a jointly-developed architecture and test effort that integrates new EMC and Cisco technologies while leveraging existing VMware components.

Table 1 Partner Technologies per Use Case

Partner	Capacity Expansion Use Case	Virtualized Workload Portability Use Case
EMC	<ul style="list-style-type: none"> • VPLEX[®] Geo <ul style="list-style-type: none"> – Continuous data availability – Extends distributed volume support up 50 ms. – Asynchronous replication with distributed data mirroring 	<ul style="list-style-type: none"> • VPLEX Geo <ul style="list-style-type: none"> – Continuous data availability – Extends distributed volume support up 50 ms. – Asynchronous replication with distributed data mirroring
VMware	<ul style="list-style-type: none"> • VMware ESXi Host • vCenter[™] 	<ul style="list-style-type: none"> • VMware ESXi Host • vCenter
Cisco	<ul style="list-style-type: none"> • Dynamic Workload Scaling (DWS) • Overlay Transport Virtualization (OTV) 	<ul style="list-style-type: none"> • Cisco Unified Computing System[™] (UCS) Service Profiles • Locator/ID Separation Protocol (LISP)

Capacity Expansion Use Case

The capacity expansion use case supports Dynamic Workload Scaling (DWS), a new Cisco technology that integrates the Cisco Application Control Engine (ACE) with Cisco’s Dynamic MAC-in-IP encapsulation technology, Overlay Transport Virtualization (OTV). The ACE creates a table of local MAC addresses and periodically polls the OTV-enabled Nexus[®] 7000 to check for any MAC entries that are considered “remote”, that is only reachable over the OTV link. In any situation where the user-defined threshold has been met, new traffic flows burst through the OTV link to be processed in the remote data center. In essence, any event that increases the data center utilization requirements can be a candidate for capacity expansion. [Table 2](#) provides some examples.

Table 2 Capacity Expansion Use Case and Benefits

Capacity Expansion Benefits	Capacity Expansion Details
Increased business resiliency	Maximize availability and uptime—DWS enables resource distribution and can dynamically burst transactions to a remote data center in the event of unplanned events.
Flexibility	Customize to your environment—DWS has user-configurable thresholds and load balancing distribution algorithms that can be customized to an existing architecture and service level requirements.
Real-time scalability	Dynamically resize your virtual data center—DWS continuously monitors thresholds levels to add new virtual data centers and also remove virtual data centers when no longer needed.

Virtualized Workload Portability

The virtualized workload portability use case integrates several technologies to enable end user migration activities. This use case may help with migration events, systems maintenance, and unplanned events that require immediate action. Virtualized workload portability uniquely combines EMC's VPLEX Geo technology, VMware's vSphere™, and Cisco's UCS Service Profiles and LISP technologies to allow simplified provisioning of a new virtualized environment as well as minimize the downtime associated with re-provisioning a new server. [Table 3](#) provides some examples.

Table 3 Virtualized Workload Portability Use Case and Benefits

Workload Portability Benefit	Workload Portability Details
Increased business resiliency	Minimize traffic disruption—LISP provides dynamic traffic redirection for existing and new client-server transactions as the remote site becomes active.
Ease of use	Simplified management—UCS Service Profile identities are replicated to the remote location, minimizing the configuration changes needed to enable the remote site to become active.
Minimize migration time	Migration duration—The stateless UCS Service Profiles and continuous data availability with VPLEX Geo minimize configuration steps and reduce overall time for migration.

Use Case Technologies

Table 4 Use Case Technologies

Partner	Technology
EMC	VPLEX Geo

Table 4 Use Case Technologies

Partner	Technology
VMware	vSphere
Cisco Systems	Dynamic Workload Scaling
	UCS Service Profiles
	Overlay Transport Virtualization
	Locator/ID Separation Protocol

EMC VPLEX Family Overview

The EMC VPLEX family today consists of:

- VPLEX Local for managing data mobility and access within the data center using a single VPLEX cluster.
- VPLEX Metro for mobility and access across two locations separated by inter-site round-trip time (RTT) of up to 5 ms. VPLEX Metro uses two VPLEX clusters and includes the unique capability in which a remote VPLEX Metro cluster can present LUNs without the need for physical storage for those LUNs at the remote cluster. It also supports synchronous distributed volumes that mirror data between the two clusters using write-through caching.
- VPLEX Geo, which also uses two VPLEX clusters, adds access between two sites over extended asynchronous distances with RTT latencies up to 50 ms. VPLEX Geo distributed volumes support AccessAnywhere distributed mirroring using write-back caching.

The VPLEX family provides the data center with the benefits of mobility, availability, and collaboration through VPLEX AccessAnywhere™ virtual storage, the breakthrough block-storage technology that enables a single copy of data to be shared, accessed, and relocated over distances. VPLEX 5.0 extends and enhances the capabilities of VPLEX 4.0 virtual storage, adding VPLEX Geo to join VPLEX Local and VPLEX Metro in the VPLEX product family. VPLEX Geo extends the reach of VPLEX AccessAnywhere storage supporting RTT inter-site latencies of up to 50 ms. through asynchronous communication.

The EMC VPLEX family removes physical barriers within, across, and between data centers. VPLEX is the first platform in the world that delivers both Local and Distributed Federation. Local Federation provides the transparent cooperation of physical elements within a site. Distributed Federation extends access between two locations across distances. At the highest level, VPLEX has unique capabilities that customers value. First, VPLEX is distributed, because it can connect multiple sites together over distances, allowing secure and consistent collaboration across distributed users. Next, VPLEX is dynamic, because it is a single interface for multi-vendor storage and it delivers dynamic data mobility, which is being able to move applications and data in real time with no outage required. And finally, VPLEX is smart, because its unique AccessAnywhere technology can present and keep the same data consistent within and between sites, even across distances.

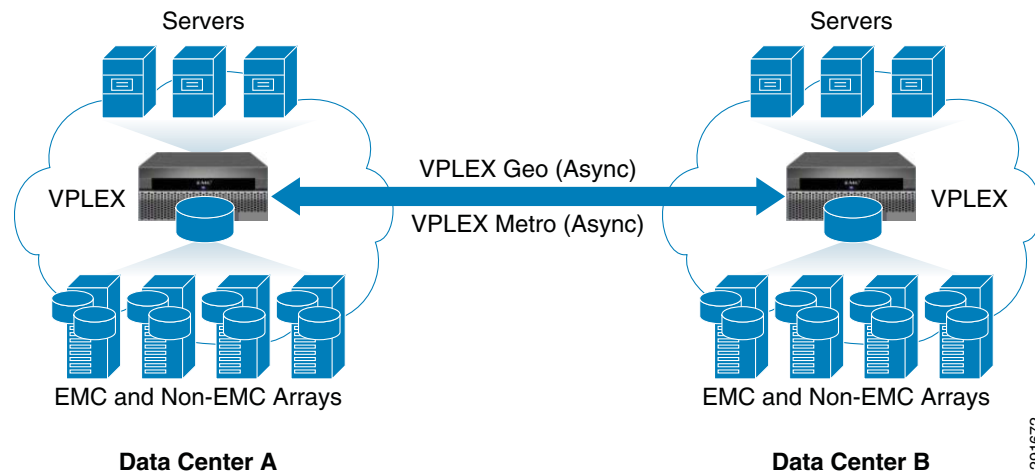
EMC VPLEX represents the next-generation architecture for data mobility and information access. This architecture is based on EMC's 20-plus years of expertise in designing, implementing, and perfecting enterprise-class intelligent cache and distributed data protection solutions.

Finally, VPLEX is a solution for federating EMC and non-EMC storage. VPLEX resides between the servers and heterogeneous storage assets and has unique characteristics in its architecture:

- Scale-out clustering hardware lets you start small and grow big with predictable service levels.

- Advanced data caching utilizes large-scale SDRAM cache to improve performance and reduce I/O latency and array contention.
- Distributed cache coherence provides automatic sharing, balancing, and failover of I/O within and between VPLEX clusters.
- A consistent view of one or more LUNs between VPLEX clusters separated either by a few feet within a data center or across asynchronous RTT distances enables new models of high availability, workload mobility, and collaboration.

Figure 1 A Multi-Cluster Deployment of VPLEX



The combination of virtual servers and EMC Virtual Storage enables entirely new ways to solve IT problems and introduces new models of computing, allowing customers to seamlessly integrate solutions built with VMware, Cisco, and EMC products.

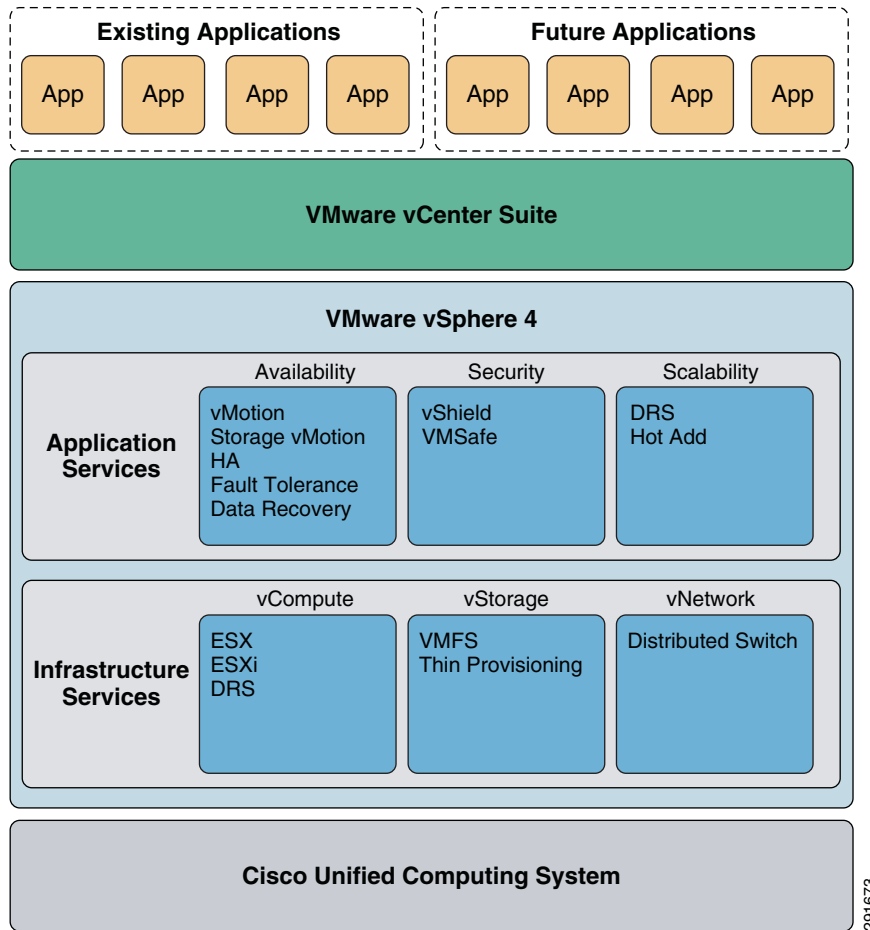
VMware vSphere Overview

VMware vSphere brings the power of virtualization to IT infrastructure—the next evolutionary step in IT computing—and a highly trusted virtualization platform. With a proven virtualization solution as the foundation, VMware vSphere and Cisco Unified Computing System can quickly respond to evolving business needs. The technology underpinnings of the solution include (see [Figure 2](#)):

- VMware vSphere vCompute services, which efficiently virtualize server resources and aggregate them into logical pools that can be precisely allocated to applications.
- VMware vStorage services, which remove the complexity of back-end storage systems and enable highly efficient utilization of storage in virtual environments.
- VMware vNetwork services, which enable optimal administration and management of networking in virtual environments.

This powerful combination of services transforms data centers with a dramatically simplified infrastructure and enables the next generation of flexible, reliable IT services. Application services can be enabled simply and uniformly for any application running in VMware virtual machines, providing IT with simple built-in controls over application service levels.

Figure 2 VMware vSphere 4



vSphere enables IT organizations to delay costly and disruptive data center expansion projects by making it possible to consolidate 15 or more virtual machines on a single physical server without sacrificing performance or throughput. Customers also use vSphere to create powerful multicore virtual machines and virtual machine clusters that span multiple physical servers to support even the most demanding applications. In addition, vSphere reduces the complexity of hardware management through comprehensive virtualization of server, storage, and networking hardware. vSphere customers can slash IT capital expenses by an average of more than 70 percent and IT operating costs by more than 30 percent.

Dynamic Workload Scaling Overview

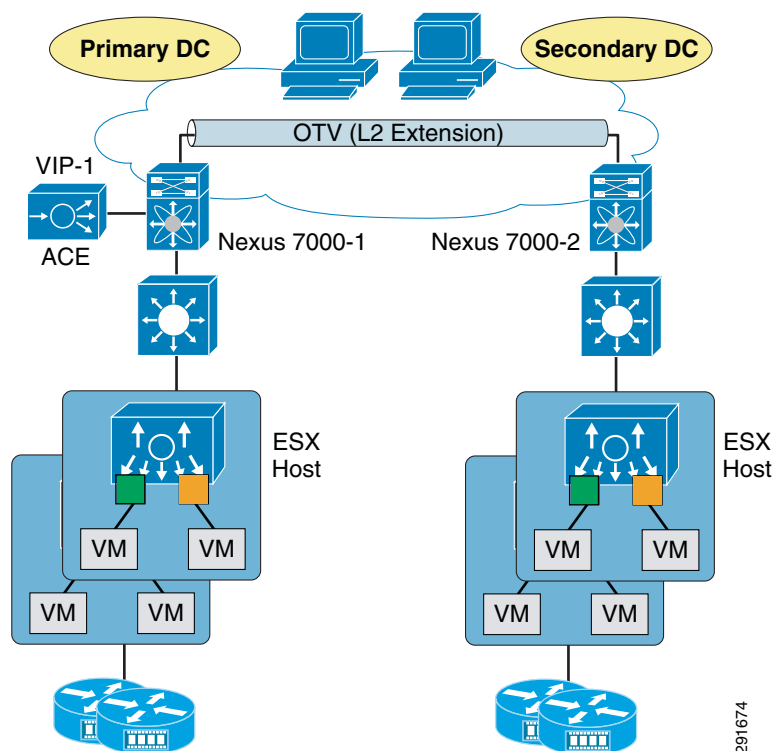
DWS integrates Cisco ACE load-balancing technology with Nexus 7000 OTV technology and VMware virtualization to deliver flexible workload mobility and application resiliency in distributed environments. DWS monitors VM capacity for an application and expands the application overload to a backup server resource pool in a secondary data center during periods of peak usage.

DWS delivers the following benefits to the ACE customer:

- Ability to scale application delivery on-demand to meet higher Service Level Agreements (SLA).
- Lower capital investment in server infrastructure through better utilization of server capacity across data center(s).

- Simplified operations through intelligent integration of application delivery with the network.
- Optimal use of OTV L2 Extension for services delivery.

Figure 3 *Dynamic Workload Scaling*



Cisco ACE, along with the Cisco Nexus 7000 and VMware vCenter, provides a complete solution for private cloud capacity expansion for data centers. In this solution, Cisco ACE actively monitors the CPU and memory information of the local VMs and computes the average load of the local data center. During normal operations, when the average load is below a pre-configured threshold, Cisco ACE load balances the incoming traffic to only local VMs. However, during peak hours, local VMs may be fully loaded and additional capacity may be required to service incoming requests. When the average load of the local data center crosses a configured threshold, Cisco ACE adds the remote VMs to its load balancing rotation pool, adding more compute resources to service the increased load.

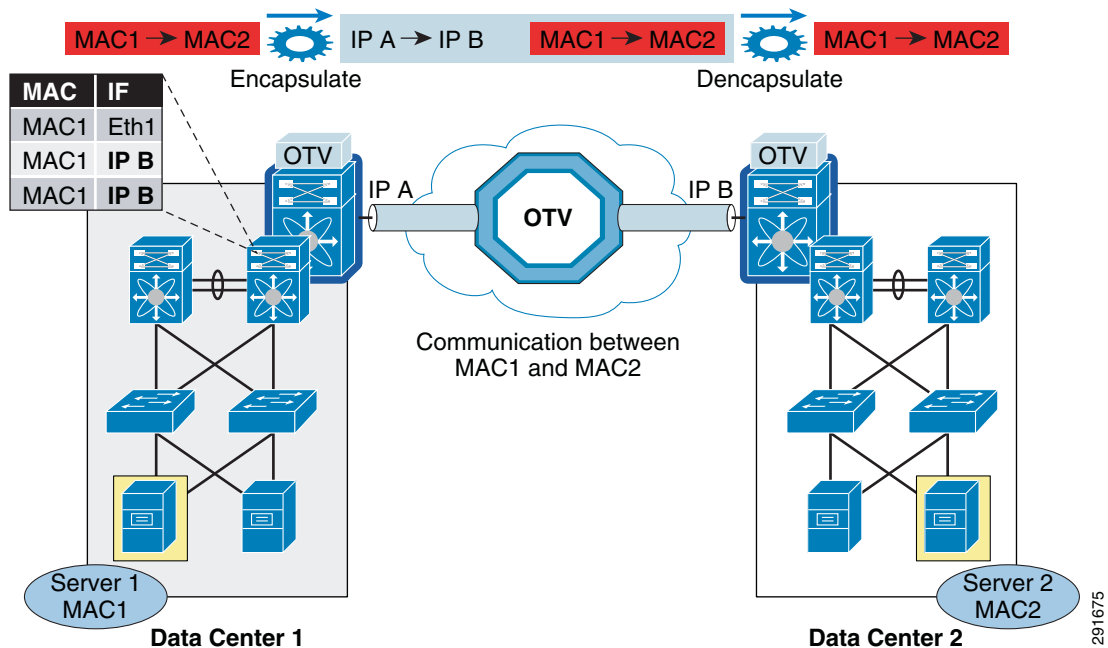
Overlay Transport Virtualization Overview

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 delivers optimal traffic load balancing, multicast traffic replication, and fast failover just like the core would. Figure 4 illustrates this dynamic encapsulation mechanism.

Figure 4 *Overlay Transport Virtualization*



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.

UCS Service Profiles Overview

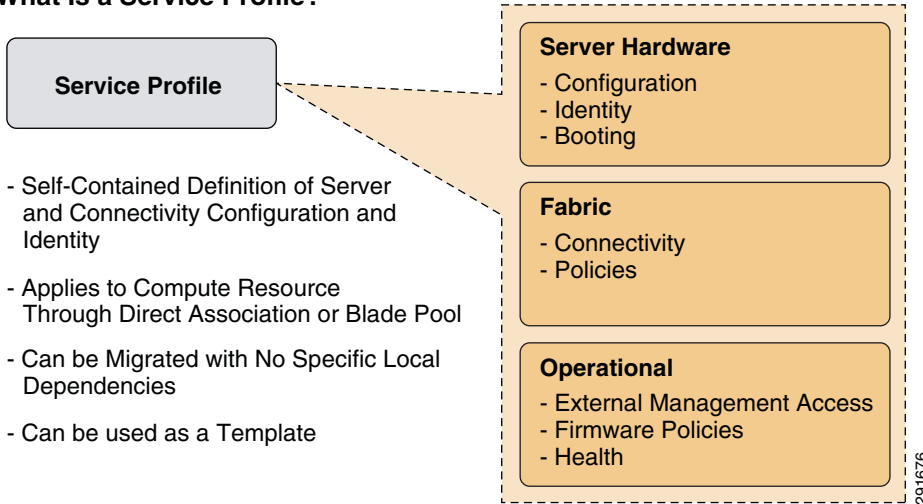
The approach of the Cisco Unified Computing System enables data center servers to become stateless and fungible, where the server's identity (using MAC or WWN addressing or UIDs) as well as build and operational policy information, such as firmware and BIOS revisions and network and storage connectivity profiles, can be dynamically provisioned or migrated to any physical server in the system. The Cisco Unified Computing System integrates server management with network and storage resources to meet the rapidly changing needs in today's data centers. New computing resources can be deployed "just in time." Traditional physical and virtual workloads can be easily migrated between servers through remote management, regardless of physical connectivity. The Cisco Unified Computing System directly improves capital utilization and operational cost and enables gains in availability, security, agility, and performance through an integrated architecture.

Cisco UCS Service Profiles provide the following benefits:

- **Service profile abstraction**—A service profile is an extension of the virtual machine abstraction applied to physical servers. The definition has been expanded to include elements of the environment that span the entire data center, encapsulating the server identity (LAN and SAN addressing, I/O configurations, firmware versions, boot order, network VLAN, physical port, and quality-of-service [QoS] policies) in logical "service profiles" that can be dynamically created and associated with any physical server in the system within minutes rather than hours or days. The association of service profiles with physical servers is performed as a simple, single operation. It enables migration of identities between servers in the environment without requiring any physical configuration changes and facilitates rapid bare-metal provisioning of replacements for failed servers. Service profiles also include operational policy information, such as information about firmware versions.
- **Service profile templates**—Since service profiles can be abstracted from the specifics of a given server to create a service profile template, templates can define policies that can be applied any number of times to provision any number of servers. Service profile templates help enable large-scale operations in which many servers are provisioned as easily as a single server.
- **Service profile logical groups**—In addition, using service profiles, Cisco UCS Manager (UCSM) provides logical grouping capabilities for both physical servers and service profiles and their associated templates. This pooling or grouping, combined with fine-grained role-based access, allows businesses to treat a farm of compute blades as a flexible resource pool that can be reallocated in real time to meet their changing needs, while maintaining any organizational overlay on the environment that they want.

Figure 5 Service Profile

What Is a Service Profile?



Another way to understand the concept of service profiles is by looking at their configuration points—the aspects of the Cisco Unified Computing System that they control.

Figure 6 Service Profile Component to Physical Device Mapping

Physical Layers Configured	Policies Created by Server Administrator Role	Policies Created by Network Administrator Role	Policies Created by Storage Administrator Role	Service Profile
Fabric Interconnects		Uplink and downlink ports, pin groups, EtherChannel definitions, and Qos policies	Fibre Channel Uplink ports, pin groups, and Qos policies	Uplink port selection, VLAN and VSAN values, MAC address and WWN pinning and Qos policies
Fabric Extenders		VN-Link parameters determined by port profiles		VN-Link maps virtual links to physical ports
Network Adapters		Mapping of physical ports to chassis		Physical port configuration including Cisco DCE and FCoE Setting
Server Resources	Server pools and assignment, local disk policy, and blade firmware version policy.	NC Adapter profiles, port profiles, service classes, VLAN policies and MAC pools	HBA Adapter profiles, service profiles, classes, VLAN policies and WWN pools	Fabric extender parameters determined by switch configuration and set by switch
	Discovery policies, pool definitions, and membership			Specify which profiles, service classes, VLAN and VSAN policies to use, which MAC address and WWN s to consume
				UUD from "database" pool
				Use server from "large memory" pool

291677

The service profile components define the server environment, including the local server settings plus storage (SAN) settings such as VSAN specifications and complete network settings such as uplink, VLAN, and QoS settings.

Locator/ID Separation Protocol Overview

Cisco LISP enables enterprises and service providers to build reliable multi-homing environments, enable mobility, and reduce operational complexities.

The current Internet routing and addressing architecture uses a single numbering space, the IP address, to simultaneously express two functions about a device: its identity and how it is attached to the network. One very visible and detrimental result of this single numbering space has been manifested in the rapid growth of the Internet's DFZ (default-free zone) as a consequence of multi-homing, traffic engineering, non-aggregatable address allocations, and business events such as mergers and acquisitions.

LISP creates a new paradigm by splitting the device identity, known as an Endpoint Identifier (EID), and its location, known as its Routing Locator (RLOC), into two different numbering spaces. Splitting EID and RLOC functions yields several advantages, including improved scalability of the routing system through greater aggregation of RLOCs and improved multi-homing efficiency and ingress traffic engineering.

LISP is a special case of tunneling that uses a dynamic encapsulation approach rather than requiring the pre-configuration of tunnel endpoints. It is designed to work in a multi-homing environment and supports communications between LISP and non-LISP sites for simple interworking. A LISP-enabled network includes some or all of the following components:

- LISP name spaces:
 - End-point Identifier (EID) addresses—Consist of the IP addresses and prefixes identifying the end-points. EID reachability across LISP sites is achieved by resolving EID-to-RLOC mappings.
 - Route Locator (RLOC) addresses—Consist of the IP addresses and prefixes identifying the different routers in the IP network. Reachability within the RLOC space is achieved by traditional routing methods.
- LISP site devices:
 - Ingress Tunnel Router (ITR)—A LISP site edge device that receives packets from site-facing interfaces (internal hosts) and encapsulates them to remote LISP sites or natively forwards them to non-LISP sites.
 - Egress Tunnel Router (ETR)—A LISP site edge device that receives packets from core-facing interfaces (the Internet) and decapsulates LISP packets and delivers them to local EIDs at the site.



Note Customer Edge (CE) devices typically implement ITR and ETR functions at the same time. When this is the case, the device is referred to as an xTR.

- LISP infrastructure devices:
 - Map-Server (MS)—A LISP infrastructure device that LISP site ETRs register to with their EID prefixes. The MS advertises aggregates for the registered EID prefixes into the LISP mapping system. All LISP sites use the LISP mapping system to resolve EID-to-RLOC mappings.
 - Map-Resolver (MR)—A LISP infrastructure device to which LISP site ITRs send LISP Map-Request queries when resolving EID-to-RLOC mappings.
 - Proxy ITR (PITR)—A LISP infrastructure device that provides connectivity between non-LISP sites and LISP sites by attracting non-LISP traffic destined to LISP sites and encapsulating this traffic to LISP sites. In the IPv6 transition case, the PITR can attract IPv6 non-LISP traffic and forward it to a LISP site using IPv4 as the transport.

- Proxy ETR (PETR)—A LISP infrastructure device that allows IPv6 LISP sites that have only IPv4 RLOC connectivity to reach LISP and non-LISP sites that have only IPv6 RLOC connectivity.


EID namespace is used within the LISP sites for end-site addressing for hosts and routers. These EID addresses go in DNS records, just as they do today. Generally, EID namespace is not globally routed in the underlying Internet. RLOC namespace, however, is used in the (Internet) core. RLOCs are used as infrastructure addresses for LISP routers and core (service provider) routers and are globally routed in the underlying infrastructure, just as they are today. Hosts do not know about RLOCs and RLOCs do not know about hosts.

The LISP solution allows any IP addressable device to move off its original subnet while keeping its IP address (its EID). Being able to preserve the IP address upon moves implies that IP connections which involve socket state can also be preserved as IP end-points move.

LISP decouples the IP addresses of end-points (EIDs) from their topological location address (RLOCs) and maintains mappings between these addresses in order to achieve routing to EIDs over the locator (RLOC) topology. LISP is a map and encapsulate routing model in which traffic destined to an EID is encapsulated and sent to an authoritative RLOC, rather than directly to the destination EID, based on the results of a lookup in a mapping database. This core operating principle of LISP is what makes it well suited to support IP mobility and IP migrations.

In a LISP enabled infrastructure, achieving mobility of end-points (EIDs) is fundamentally a matter of updating the mapping between the end-point addresses (EIDs) and their locations (RLOCs) when a move occurs. As EIDs change locations, mappings to RLOCs representative of the new location must be included in the database. Dynamically changing the RLOC to which traffic is sent for a specific EID address allows end-points (EIDs) to continue to receive traffic as they seamlessly move behind different RLOCs (i.e., to different locations). Since the EIDs remain unchanged, any stateful connections to the end-point are preserved and, at the same time, the RLOC topology is not altered, so no routing reconvergence is required in the RLOC space. The only update necessary is that of the mappings between EIDs and RLOCs, but the addresses themselves and the routing in the different spaces remain untouched. See [Figure 7](#) for additional details.

Figure 7 Locator/ID Separation Protocol



- **EID (Endpoint Identifier)** is the host IP address
- **RLOC (Routing Locator)** is the infrastructure IP address of the LISP router
- **Mapping Database (M-DB)** is the distributed database and policy repository

291674

Use Case Proof of Concept Testing

Capacity Expansion Use Case

Capacity expansion using dynamic workload scaling and continuous data availability.

Capacity Expansion

Business Challenge

- Need for on-demand server availability
 - Monitor and distribute service traffic at user-defined threshold criteria
 - Scale data center and minimize expenses
 - Continuous data availability
-

Network Solution

- Cisco Dynamic Workload Scaling feature for on-demand cloud bursting
- EMC VPLEX Geo for continuous data availability
- Layer 2 Extension using OTV

Business Results

- User defined trigger bursts new transaction to remote Data Center
 - Tested distance of 3,000 kilometers
 - Local virtual machine servers resources maintained
 - DWS feature dynamically bursts traffic over the OTV Layer 2 extension
-

Business Challenge

Traffic within the enterprise private cloud or enterprise-service provider hybrid cloud may be unpredictable. Unforeseeable real-time events may dramatically increase the load beyond the expected traffic utilization levels. Conversely, traffic patterns may predictably and temporarily exceed system utilization levels that are considered thresholds for adding additional system capacity. In both cases, additional on-demand system capacity should be available for applications and transactions.

There is a balance between maximizing service availability and efficiently allowing capital expenses, therefore virtualizing the data center to share assets and resources may help distribute and reduce expenses while increasing service availability and uptime.


In the virtualized data center, the physical data centers are geographically distributed and connected using Data Center Interconnect (DCI). DCI connects data centers using either Layer 2 or Layer 3 protocols and features. The result is that the virtualized data center can now share the physical infrastructure and maximize its utilization. Workloads can migrate between data centers or resources can be share by redirecting traffic over a Layer 2 extension.

The capacity expansion use case addresses the need for on-demand workload capacity expansion triggered by defined threshold criteria, whether due to a planned or unplanned event.

Capacity Expansion Test Results

The capacity expansion test results are expressed through the graph shown in [Figure 8](#).

Figure 8 Capacity Expansion Test Results



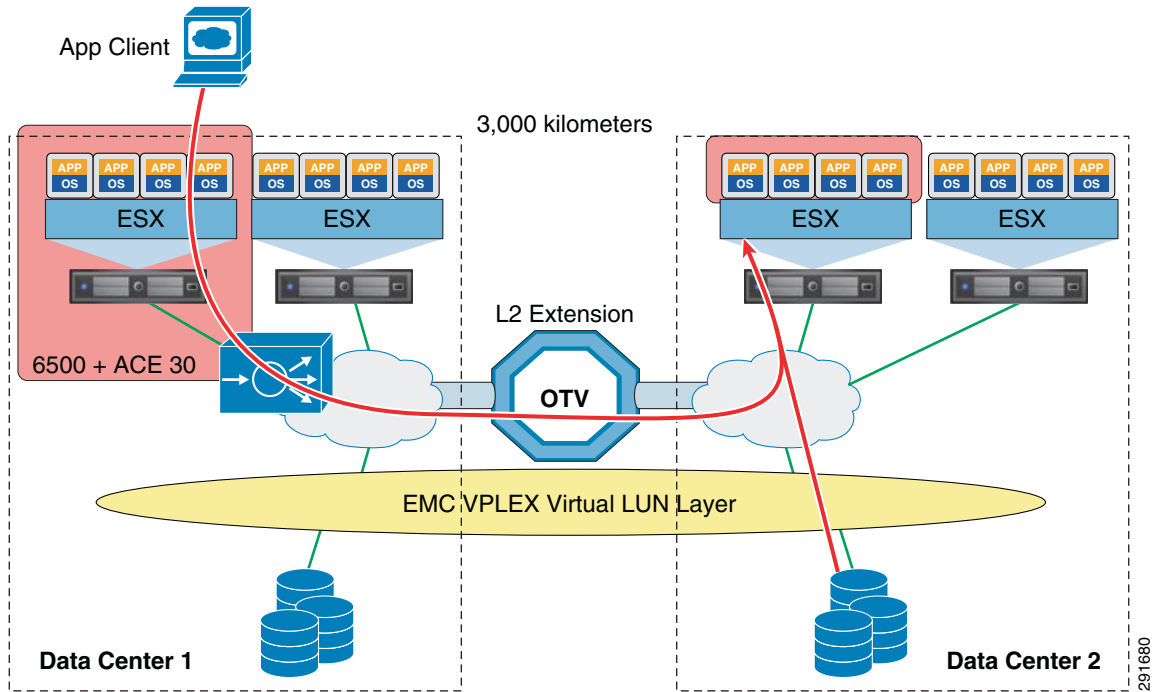
Results Discussion

The client-server traffic is processed on the local virtual machine servers until the local virtual machine servers reach a user-defined threshold level, in this case 50%. Once this trigger is exceeded, new connections will burst over the Layer 2 extension, OTV, using the Dynamic Workload Scaling feature and be served by the virtual machines approximately 3,000 kilometers away. As a result, new transactions burst, on-demand, through OTV Layer 2 extension to the remote data center and the local data center virtual machine CPUs stabilize at approximately 60%. The EMC VPLEX Geo technology provides stored content availability with continuous asynchronous replication and caching capabilities.

Capacity Expansion Logical View

Client traffic enters data center 1 and is forwarded through the OTV L2 Extension if the defined threshold trigger is exceeded. This on-demand burst, temporary in nature, allows additional capacity to be successfully processed by distributing the resource requirement to the remote data center.

Figure 9 Capacity Expansion Logical View



Capacity Expansion Test Details

Table 5 Capacity Expansion Test Setup

Application	Server Configuration	Stress-Generation Tool	Application Performance Metrics	Description
Microsoft SQL Server 2005 (64-bit)	CPU—2 virtual CPUs (vCPUs) Memory—4 GB Storage—EMC VPLEX, EMC VMAX, EMC DMX3 OS—Microsoft 2008 64-bit server	Dell DVD Store open source benchmark	Orders per minute (OPMs) Server CPU processor utilization	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.

Capacity Expansion Test Methodology

1. Reinitialize the Microsoft SQL Server by copying in the initial version of the SQL database.
2. All clients perform a trace route to the public IP address of the VM SQL server.
3. Start the first Dell DVD Store client on a virtual machine that has IP connectivity to the VMware ESX servers.

4. Each SQL clients create 49 TCP threads consisting of database inquiries, new orders, and order updates.
5. Wait 10 minutes to allow the servers in the local data center to ramp up and normalize.
6. Start the second Dell DVD store client on a virtual machine that has IP connectivity to the VMware ESX servers.
7. Wait 10 minutes to allow the servers in the local data center to ramp up and normalize.
8. Start the third Dell DVD store client on a virtual machine that has IP connectivity to the VMware ESX servers. After the third client is started, 144 threads are being established to the VM Servers in DC1 and the aggregate CPU percentage load on the servers in the local data center average 50% or greater.
9. Start the fourth Dell DVD store client on a virtual machine that has IP connectivity to the VMware ESX servers.
10. All new threads from the fourth client get established to the remote data center. CPU percentage load on the local data center servers has stabilized.
11. Collect OPM and CPU statistics and evaluate the ACE30 operations and the VM Server CPU utilizations.

Capacity Expansion Test Overview

The goal of joint testing is to demonstrate capacity expansion. Capacity expansion is the ability to monitor local virtual machine server resources and to provide additional computing power in remote data centers during peak periods. Capacity expansion leverages EMC's VPLEX Geo products over long distances and Cisco's ACE30 DWS feature along with the Nexus 7000 and OTV. Capacity expansion is a valuable tool which allows data center administrators to dynamically allocate resources to remote data centers, which may be minimally used at peak periods.

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 monitored in each data center and performance of the application is measured in OPMs versus the actual CPU percentage utilization of the host server.

EMC VPLEX Geo Configuration

The EMC VPLEX Geo provides distributed virtual volumes across the data centers for the SQL database. Inter-cluster connectivity was established between the clusters using 10 Gigabit Ethernet across the Cisco OTV cloud. VPLEX consistency groups running in asynchronous cache mode were used to leverage write back caching capabilities at each cluster. To increase tolerance to fluctuations in the write workload and to smooth overall performance, the consistency group's maximum queue depth was increased to 64 and jumbo frames (MTU set to 9000) were enabled on the 10 Gigabit Ethernet interfaces.

The following two configuration updates were made on the VPLEX Geo.

Example 1 *Maximum Queue Depth*

```
VPlxcli:/clusters/cluster-1/consistency-groups/CG-DC1/advanced>ll
Name                               Value
-----
auto-resume-at-loser                false
current-queue-depth                 1
current-rollback-data               7.6M
```

```

default-closeout-time      0.5min
delta-size                 16M
local-read-override       true
max-possible-rollback-data 1.97G
maximum-queue-depth      64 (default = 6 )
potential-winner           -
write-pacing               inactive
    
```

Example 2 Maximum Transmission Unit

```

Vplexcli:/> ll /engines/**/ports/[AB]*-X*

/engines/engine-1-1/directors/director-1-1-A/hardware/ports/A2-XG00:
Name                               Value
-----
address                             10.38.21.35
config-status                       applied
enabled                             true
mtu                               9000 (default = 1500 )
operational-status                 ok
port-status                         up
protocols                           [udt]
role                                 wan-com
speed                               10000
    
```

An additional EMC VPLEX configuration included consistency groups. Consistency groups are logical sets of virtual volumes that provide a consistent point in time across all members. They provide write consistency for applications such as databases, clusters, and applications groups. Each of these use cases require the same point in time for all virtual volumes to properly recover. By grouping virtual volumes in this fashion, the VPLEX administrator can also ensure that the I/O behavior, cache mode, and cluster detach rule set for all volumes are the same.

Dynamic Workload Scaling Configuration

The ACE30 module was configured for DWS in the local data center. DWS uses application programmable interface (API) functionality to monitor the server load information polled from the VMware’s VCenter application. DWS also uses an API to poll OTV information from the Nexus 7000 and determines which hosts are in the local data center and which hosts reside in the remote data center.

Example 3 ACE30 DWS Configuration

```

probe vm VCENTER-VPLEX-GEO
  load cpu burst-threshold max 50 min 25
  vm-controller VCENTER-VPLEX-GEO
serverfarm host SRV1
  predictor leastconns
  dws burst probe VCENTER-VPLEX-GEO
  rserver VM1
    inservice
  rserver VM2
    inservice
  rserver VM3
    inservice
  rserver VM4
    inservice
    
```

The test was performed using (three) VMware clients simulating Internet users and performing SQL updates and reads to the SQL servers. All client connections were established to the SQL servers in data center 1. These connections increased the local servers CPU to exceed the aggregate 50% threshold configured on the ACE30. At this point, another SQL client was started. The ACE30 then forwarded those new connections to the SQL servers in the remote data center.

The following information on the ACE30 was captured after the initial three clients had established SQL sessions:

```
dclc-ace-s9/cap-exp# show probe VCENTER-VPLEX-GEO detail
```

```
probe      : VCENTER-VPLEX-GEO
type       : VM
state      : ACTIVE
description :
-----
interval   : 10      vm-controller : VCENTER-VPLEX-GEO
cpu-load:
  burst-threshold:
    max threshold : 50      min threshold : 25
mem-load:
  burst-threshold:
    max threshold : 99      min threshold : 99
----- probe results -----
associations ip-address  cpu-load mem-load  health
-----+-----+-----+-----+-----
serverfarm : SRV1
aggregate-stats  59      69      BURST_REMOTE

No. Passed probe      :      501      No. Failed probe : 0
No. Probes skipped   :      0      Last status code : 0
Last probe time      : Tue Apr 26 10:37:50 2011
Last fail time       : Never

  real      : VM1[0]
             10.25.1.11      60      67      SUCCESS

  real      : VM2[0]
             10.25.1.12      59      72      SUCCESS
```

```
dclc-ace-s9/cap-exp# show serverfarm SRV1
```

Codes: L - local, R - remote

```
serverfarm      : SRV1, type: HOST
total rservers  : 6
state           : ACTIVE
DWS state      : ENABLED_REMOTE_LB(Bursting traffic to local and remote VMs)

-----connections-----
  real      weight state      current  total  failures
-----+-----+-----+-----+-----+-----
rserver: VM1
10.25.1.11:0      8  OPERATIONAL [L] 74      959      0
rserver: VM2
10.25.1.12:0      8  OPERATIONAL [L] 73      952      0
rserver: VM3
10.25.1.21:0      8  OPERATIONAL [R] 0      134      0
rserver: VM4
10.25.1.22:0      8  OPERATIONAL [R] 0      130      0
```

The ACE30 determined that VM3 and VM4 are located in the remote data center by polling the Nexus 7000 and OTV information for those servers through the API (for brevity, the local server information is not shown):

```
dc1c-agg-7k1# show ip arp | i 2501
10.25.1.21      00:15:58 0050.568f.0021 Vlan2501
10.25.1.22      00:15:54 0050.568f.0023 Vlan2501
```

```
dc1c-agg-7k1-otv# show otv route vlan 2501
OTV Unicast MAC Routing Table For Overlay200
```

VLAN	MAC-Address	Metric	Uptime	Owner	Next-hop(s)
2501	0050.568f.0021	42	1w1d	overlay	dc2c-agg-7k1-otv
2501	0050.568f.0023	42	1w1d	overlay	dc2c-agg-7k1-otv

After starting the fourth client, SQL connections are now forwarded to the remote data center and the CPU percentage threshold remains steady in the local data center:

```
dc1c-ace-s9/cap-exp# show probe VCENTER-VPLEX-GEO detail
```

```
probe      : VCENTER-VPLEX-GEO
type       : VM
state      : ACTIVE
description :
-----
interval   : 10      vm-controller : VCENTER-VPLEX-GEO
cpu-load:
  burst-threshold:
    max threshold : 50      min threshold : 25
mem-load:
  burst-threshold:
    max threshold : 99      min threshold : 99
----- probe results -----
associations ip-address  cpu-load  mem-load  health
-----+-----+-----+-----+-----
serverfarm  : SRV1
              aggregate-stats  59      69      BURST_REMOTE

No. Passed probe      :      512      No. Failed probe : 0
No. Probes skipped    :      0      Last status code : 0
Last probe time       : Tue Apr 26 10:39:40 2011
Last fail time        : Never

real      : VM1[0]
          10.25.1.11      60      67      SUCCESS

real      : VM2[0]
          10.25.1.12      59      72      SUCCESS
```

```
dc1c-ace-s9/cap-exp# show serverfarm SRV1
```

```
Codes: L - local, R - remote

serverfarm : SRV1, type: HOST
total rservers : 6
state       : ACTIVE
DWS state   : ENABLED_REMOTE_LB(Bursting traffic to local and remote VMs)

real      weight state      -----connections-----
current  total  failures
-----+-----+-----+-----+-----
rserver: VM1
10.25.1.11:0      8  OPERATIONAL [L] 74      959      0
```

rserver: VM2	10.25.1.12:0	8	OPERATIONAL [L]	73	952	0
rserver: VM3	10.25.1.21:0	8	OPERATIONAL [R]	25	158	0
rserver: VM4	10.25.1.22:0	8	OPERATIONAL [R]	24	154	0

Virtualized Workload Portability Use Case

Virtualized workload portability uses UCS Service Profile abstraction, LISP, and continuous data availability.

Workload Portability

Business Challenge

- Workload migration complexity
- Optimizing traffic forwarding to new, remote location
- Minimize traffic disruption or downtime during and post migration

Network Solution

- UCS Service Profiles for simplified deployment
- LISP for optimized traffic redirection
- VPLEX Geo for continuous data availability

Business Results

- Service Profile migration steps are simplified with duplication
 - Stored content continuously available in local and remote locations.
 - LISP redirects traffic to new location
-

Business Challenge

Cloud computing requires accessibility, availability, and scalability. Customers who desire scalability have, in the past, found themselves constrained by accessibility. Scaling out workloads to a new data center can cause availability outage times measured in days or weeks, due to shutting down servers, loading on trucks, shipping, reinstallation, and repair of any damage caused by the move process. Reconfiguring workloads to accommodate new networking can also cause unforeseen problems, with IP conflicts, remapping of DNS systems, and clients that, for whatever reason, connect via IP address rather than DNS name. Centralized storage, generally seen as a benefit for its simplification of backup and data availability features, adds another level of complication, as data that moved servers rely on can now be separated by high latency connections, drastically reducing system performance. In the event that the workload needs to return to the original data center, all of the same costs in time and resources apply, further reducing scalability and accessibility.

Solution Description

To achieve accessibility while utilizing scalability, customers require workload mobility that provides ease of management, rapidity, and physical range. The challenges of obtaining manageable, rapid, and long distance workload mobility stem from data center designs that do not take mobility into consideration.

Virtualization has been a key enabler to mobility by freeing the application and operating system from a physical server, allowing the application to be moved with relative ease throughout a data center. Virtualization frees applications from the requirement to operate on only one particular server. But applications are tethered to their permanent storage system and network identity, which limits the range of mobility to short distances accessible via SAN connectivity, typically a single data center or one within a few kilometers of SAN fabric path. IP addresses not only confer identity, but locality as well; a server node needs to remain on the subnet where routers expect it to reside. Otherwise, connectivity via TCP/IP is disrupted. Creating subnets with large amounts of nodes is not an effective remedy as packet storms can ensue.


The innovative products incorporated in the design below remedy these restraining issues, allowing applications to be migrated long distances to remote data centers with minimal, if any, availability impact.

EMC VPLEX Geo is a storage active/active virtualization solution that replicates stored data automatically between data centers. This does not remove the tether between workload and data, but lengthens it to asynchronous network distances, potentially up to 5000 kilometers. Workloads can access the data they rely on locally, whether they reside in California or Massachusetts. Since the storage is virtualized, no reconfiguration of the workload is required, easing management and minimizing risk of configuration errors.

Cisco OTV and LISP technologies contained in the Nexus 7000 product line allow IP subnets to be extended across data centers. This lengthens the IP network tether and improves workload accessibility.

By combining these technologies, customers can enjoy the flexibility of migrating workloads to where they are needed the most. Capacity can be scaled to accommodate demand with minimal cost in time and resources. And the two-way nature of the design allows workloads to be returned as easily as they were deployed. This design provides scalability with flexibility, freeing resources that would have otherwise been dedicated to workload migration across distance.

Figure 10 Workload Migration Solution Logical View



201681

Table 6 Virtualized Workload Test Setup

Application	Server Configuration	Stress-Generation Tool	Application Performance Metrics	Description
Microsoft SQL Server 2005 (64-bit)	CPU—8 virtual CPUs (vCPUs) Memory—8 GB Storage—EMC VPLEX, EMC VMAX, EMC DMX3 OS—Microsoft 2003 64-bit server	Dell DVD Store open source benchmark	Orders per minute (OPMs) Server CPU processor utilization	The DVD Store benchmark is an 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.

Virtualized Workload Portability Test Overview

The migration use case explores the movement of one application, the Dell DVD Store 2.0, from one data center to another, capturing Time To Recover (TTR) as the primary measurement of success. Of secondary, though not unimportant, consequence are the OPM values registered by the DVD Store itself. These are used to verify application availability at levels previously found acceptable prior to the migration. As distances between data centers increase, network latency increases. This increase in latency results in an increasing TTR because the access time of the data storage increases.

In a data center interconnect solution, a land extension is needed when applications running in the VMs use non-routable traffic, such as Node Discover and Heartbeats in clustered applications, and those VMs are distributed across data centers. In a virtualized environment, it is likely that the application may be distributed across data centers. So it is recommended to implement a land extension technology such as OTV in a data center interconnected solution. In this solution, OTV is optional and not required as we were moving all the VMs in the application. LISP is implemented to provide optimal routing between client and host. As the application departs one location and joins the other, LISP updates routing so that clients are not required to pass through older routes to connect to the application.

Storage networking between data centers is connected via FCoE through IP network. Maintaining permanent storage in one data center only would greatly impact the application response time and place the system in a potentially compromised state in the event of disconnection between data centers. EMC VPLEX Geo was chosen for data store replication. VPLEX Geo provides asynchronous storage replication between sites for latencies up to 50 ms. round trip. This roughly equates to a distance of 5,000 km, though there is no accurate direct correlation between latency and physical distance. VPLEX Geo provides consistent WWNNs for each distributed volume, minimizing chances of configuration errors between data centers.

Cisco UCS servers were used to host the applications. More than just a grouping of blades, the Cisco UCS system utilizes a new concept in server administration called Service Profiles. Each server in a Cisco UCS system has its own identifying addresses, such as network card MAC addresses, HBA WWPN and WWNN, and UUID. On any other server, these addresses would be embedded permanently in the server, unchangeable and difficult to mask. On Cisco UCS, however, these embedded addresses can be overridden, replaced by administrator-defined addresses acquired from “address pools”. These pools can be set to any value the administrator of the UCS system sees fit. As such, they create a perfect candidate for this scenario, since the “identity” of each of the blades can be completely and accurately replicated to each data center well in advance. (Figure 13 and Figure 14) depict the Cisco UCSM view of the Service Profiles on each data center. As can be seen, each Profile has the same sets of MAC and WWPN addresses as its corresponding twin in the other data center. This allows the servers to be migrated in a simple, step-wise fashion: disassociate the Service Profiles in the source data center, associate the Service Profiles in the destination, and verify that the applications are running.

EMC VPLEX Geo provides distributed virtual volumes across the data centers for the SQL database. Inter-cluster connectivity was established between the clusters using Cisco MDS FCIP. Fiber Channel Write Acceleration feature is enabled on MDS for improved storage access time. A single consistency group in asynchronous mode is used for the distributed volumes. The consistency group ensures that I/O to all volumes in the group is coordinated across both clusters and all directors in each cluster. The interconnection between the two data centers is a 1 Gb Ethernet and the IP network is enabled with jumbo frames (MTU set to 9000) on all the network interfaces.

The test bed was configured with two UCS chassis, one in each data center, with each chassis including four blades. Duplicate Service Profiles were created on each UCS system with duplicate pools of MAC and WWPN addresses. EMC Vplex Geo provided the storage replication between the two data centers. Vplex Geo operates in asynchronous synchronization mode and this enhanced the storage access time from the computing devices. This ensured that the same WWPN in one data center had access to the same

distributed volume, regardless of data center in which it was active. The network VLAN was stretched between data centers using OTV and routing changes were updated using LISP to ensure clients used the fastest route to access the server in each case.

Virtualized Workload Portability Test Methodology

1. Connect the DVD store client (SQL client) to the DVD store SQL server and run for five minutes with 48 threads. At the end of 5 minutes, note OPM.
2. Shutdown or suspend the DVD store SQL server in vCenter.
3. Disassociate the Service Profiles in the current data center in UCSM.
4. Associate the duplicate Service Profiles in the target data center.
5. Set the power state of the target data center Service Profiles to Boot Server.
6. Once ESXi servers are pingable, power on the DVD store SQL server VM.
7. Once the DVD store SQL server is pingable, connect the SQL Client.
8. Run the SQL Client for five minutes, taking note of the steady state OPM.
9. Collect test statistics to evaluate the total elapsed time.

Virtualized Workload Portability Test Results

The scenario explored by these tests is to migrate workloads easily and with virtually no re-configuration of network or storage equipment using UCS Service Profiles.

Two different migration methods were used: shutdown and suspend. In the shutdown method, the SQL server VM Guest OS was shutdown. In the suspend method, the VM was suspended.

The goal of the joint testing is to measure the completion time taken for the overall migration and the impact to application performance in terms of OPMs due to the migration of the workload between the data centers. The overall migration time is an important measure and it becomes critical when attempting to determine the feasibility of migrating to a new data center using Service Profiles. The duration of a migration largely depends on the latency between two data centers, the amount of memory configured on the virtual machine, and the amount of bandwidth available between the data centers.

[Figure 11](#) and [Figure 12](#) are graphs that show migration times versus distance. Distance is measured in term of round trip latency. Typically, 1 ms. of round trip latency is equivalent to 100 Kilometers.

- Distance:
 - 10ms = 1,000 Kilometers
 - 20ms = 2,000 Kilometers
 - 30ms = 3,000 Kilometers
 - 40ms = 4,000 Kilometers
- Migration time:
 - Shutdown/Suspend—Time that vCenter takes to shutdown or suspend the VMs.
 - Dis-associate and associate—Time to dis-associate the service profile in the current data center and then associated the service profile in the target data center
 - ESX server up—Time for the ESX server to boot up
 - Workload re-established IP connectivity—Time for the SQL Server VM to boot up and establish IP connectivity with the client

Figure 11 Virtualized Workload Portability—VM Shutdown





Figure 12 Virtualized Workload Portability—VM Suspend



Results Discussion

These results indicate some important findings.

First, the amount of time to recovery of the application depended on the amount of the data that was written to the disk, as we assumed. This is indicated by the time required for the workload to re-establish IP connectivity.

Second, the difference between shutting down the SQL server and suspending it was negligible. Each method required some recovery time and each method showed that stored data on the SAN had been replicated to the new data center in time for a quick recovery, in some cases at or below 20 ms.

In summary, the testing indicates that it is feasible to migrate at least one heavily loaded server from one data center to another, quickly and easily and with minimal reconfiguration, using VPLEX Geo, OTV, LISP, and Service Profiles across high latency links. When compared to the traditional methods of migrating an application, which takes days rather than minutes, this solution provides significant benefits in saved resources.

Example 4 LISP Configuration on the Aggregation Device

```
DC1-AGG1#
feature lisp
```

```

ip lisp itr-etr
ip lisp database-mapping 192.168.94.0/24 192.168.5.6 priority 1 weight 50
ip lisp database-mapping 192.168.94.0/24 192.168.5.10 priority 2 weight 50
ip lisp database-mapping 192.168.94.0/24 192.168.5.14 priority 1 weight 50
ip lisp database-mapping 192.168.94.0/24 192.168.5.18 priority 2 weight 50
ip lisp itr map-resolver 192.168.1.210
ip lisp etr map-server 192.168.1.210 key <datacenter>
lisp dynamic-eid emc_world_vlan94
  database-mapping 192.168.94.128/25 192.168.5.6 priority 1 weight 50
  database-mapping 192.168.94.128/25 192.168.5.10 priority 2 weight 50
  database-mapping 192.168.94.128/25 192.168.5.14 priority 1 weight 50
  database-mapping 192.168.94.128/25 192.168.5.18 priority 2 weight 50
  map-notify-group 224.0.94.11

interface Vlan94
  no shutdown
  no ip redirects
  ip address 192.168.94.2/24
  lisp mobility emc_world_vlan94
  lisp extended-subnet-mode
  hsrp 94
    preempt delay minimum 120
    priority 200
  ip 192.168.94.1

DC1-AGG1#

```

LISP Output

The LISP data shown below indicates that routing optimization does indeed occur on application migration. It is important to note that the client was not required to perform any action to cause this network reconfiguration to occur.

We examine the LISP mapping information in the branch's xTR and see how it was changed after migrating from data center 1 to data center 2.

The virtual machine in this case has an IP address of **192.168.94.150**.

The virtual machine is up in data center 1 and the xTR in data center 1 detects the VM and registers it in the mapping systems.

```

DC1-AGG1# show lisp dynamic-eid summary
LISP Dynamic EID Summary for VRF "default"
* = Dyn-EID learned by site-based Map-Notify
Dyn-EID Name   Dynamic-EID      Interface    Uptime      Last      Ping
                Packet          Count
emc_world_vlan 192.168.94.150 Vlan94      00:09:31   00:09:31   0
DC1-AGG1#

```

When a client from the branch wants to communicate with the server, the branch xTR device requests the routing information from the map-server. Once it receives the information, it stores them in its cache. In the following output, there are four locators in data center 1 that the branch can use to reach VM **192.168.94.150**. Those four locators (**192.168.5.x**) are IP addresses in the data center 1's aggregation devices.

```

BRANCH-xTR# sho ip lisp map-cache
LISP IP Mapping Cache for VRF "default" (iid 0), 1 entries

192.168.94.150/32, uptime: 2d15h, expires: 23:50:42, via map-reply, auth
  Locator      Uptime      State      Priority/  Data      Control
                Weight      in/out    in/out

```

```

192.168.5.6 00:09:17 up 1/50 0/0 3/2
192.168.5.10 00:09:17 up 2/50 0/0 0/0
192.168.5.14 00:09:17 up 1/50 0/0 0/0
192.168.5.18 00:09:17 up 2/50 0/0 0/0

```

BRANCH-xTR#

After the VM moves to data center 2, the Nexus 7000 in the aggregation layer detects the VM is now in its data center. It then registers the information with the map-server.

```

DC2-AGG2# show lisp dynamic-eid summary
LISP Dynamic EID Summary for VRF "default"
* = Dyn-EID learned by site-based Map-Notify
Dyn-EID Name    Dynamic-EID    Interface    Uptime    Last    Ping
                Dynamic-EID    Interface    Uptime    Last    Ping
                Dynamic-EID    Interface    Uptime    Last    Ping
                Dynamic-EID    Interface    Uptime    Last    Ping
emc_world_vlan 192.168.94.150 Vlan94       00:00:30  00:00:24  0
DC2-AGG2#

```

After the move, the communication between the branch device and the VM triggers an update on the branch xTR. The branch xTR updates its map-cache with new locators. In the following output, it indicated there are four locators in data center 2 that branch can use to reach VM **192.168.94.150**.

```

BRANCH-xTR# show ip lisp map-cache
LISP IP Mapping Cache for VRF "default" (iid 0), 1 entries

192.168.94.150/32, uptime: 2d11h, expires: 23:54:08, via map-reply, auth
Locator    Uptime    State    Priority/    Data    Control
           Uptime    State    Weight      in/out  in/out
10.1.5.6   11:35:51  up       1/50        0/0    2/0
10.1.5.10  11:35:51  up       2/50        0/0    0/0
10.1.5.14  11:35:51  up       1/50        0/0    0/0
10.1.5.18  11:35:51  up       2/50        0/0    0/0

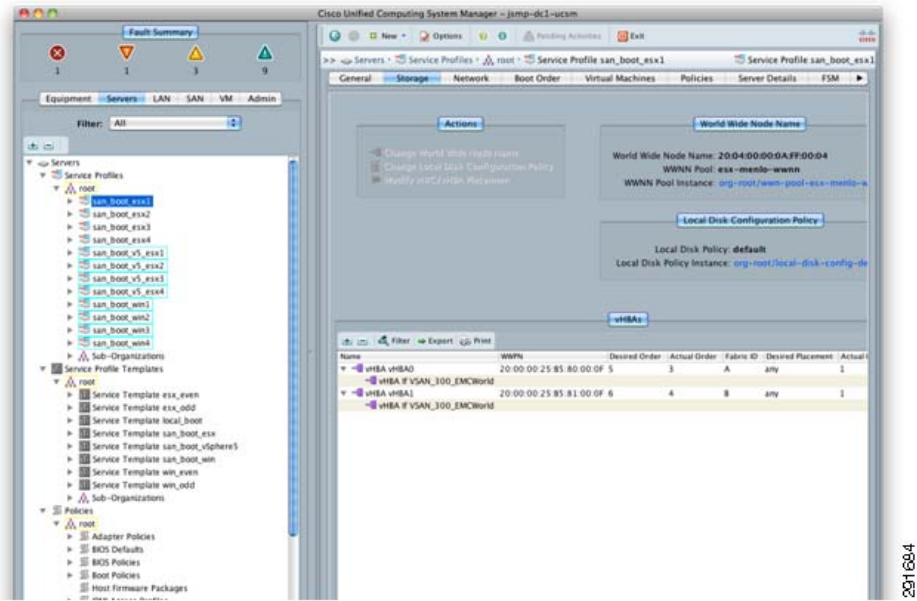
```

BRANCH-xTR#

UCSM Service Profile View

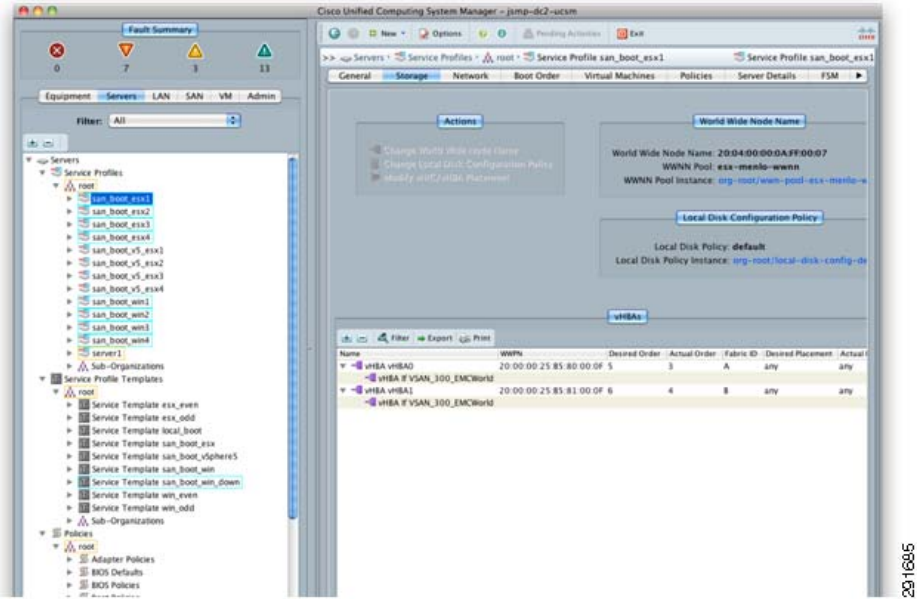
Figure 13 and Figure 14 depict the UCSM view of the data center 1 and 2 Service Profiles. Notice that each set of profiles is identical, which allows migration without re-configuration at migration event time.

Figure 13 UCSM Configuration—Data Center 1



201694

Figure 14 UCSM Configuration—Data Center 2



201695