

Cisco UCS Outperforms Dell on East-West Network Latency



Performance Brief
April 2014

Highlights

Well designed test methodology

- Cisco designs tests to mimic real-world workloads as much as possible
- Cisco runs thousands of iterations, not just best-case scenarios, to gain insight into test variances and averages

Better performance

- The Cisco Unified Computing System™ (Cisco UCS®) delivers consistent low-latency and high-workload throughput.

Faster transaction times

- Cisco UCS demonstrated lower latency than the Dell solution for every test group and every packet size (User Datagram Protocol [UDP], TCP, and round-trip time [RTT] TCP).

Faster virtual environments

- Cisco UCS delivered better performance (faster virtual machine migration times) than the Dell solution for every virtual machine migration group size tested

Test methodology matters. As any IT administrator knows, the performance of applications over time varies as a result of changes in load in response to business needs. The innovative architecture of the Cisco Unified Computing System™ (Cisco UCS®) delivers consistently faster applications and virtualized environments to businesses with more dependable latency.

For a test to meter performance reliably, it must be designed both to stress the system components being tested and to run multiple iterations of the test. Otherwise, you fail to gain important insight into how the system will perform in the both best cases and the worst. Cisco designed these tests to, as much as possible, reflect the real ways in which our customers and their applications use the system and associated components (rather than collecting just the best-case results). Then we run many test iterations, with different workloads, to discover performance trends and variances.

The tests we ran revealed that latency in Cisco UCS, compared to latency in a Dell PowerEdge solution, is the opposite of what Dell has been communicating. The tests showed that Dell's claims about Cisco UCS performance—that Cisco UCS is slower than the Dell solution because all blade traffic has to go through the fabric interconnect—are wrong: Cisco UCS demonstrated lower latency and faster virtual machine migration timing in every use-case test Cisco performed against Dell PowerEdge M620 blade servers with the Dell PowerEdge M I/O Aggregator. This document explains how the Cisco UCS unified fabric works and presents the results of comprehensive latency tests using Cisco UCS and Dell PowerEdge blade servers.

Cisco UCS: Innovative Architecture

Cisco UCS and Dell PowerEdge have significantly different architectures. These differences help Cisco UCS deliver consistency, visibility, and portability across servers, regardless of whether they are physical or virtual. Cisco's innovation in building the end-to-end fabric at the core of Cisco UCS enables business applications to run better and faster. Cisco UCS is designed to function as a single large virtual chassis that can support up to 160 blade or rack servers in a single management and connectivity domain. This unique architecture provides the flexibility to place workloads anywhere within Cisco UCS with consistent network

Cisco UCS Outperforms Dell on East-West Network Latency

performance. Important benefits of the Cisco UCS architecture include reduced latency across all I/O operations and consistent I/O latency between physical servers.

How Cisco UCS Unified Fabric Works

Cisco UCS uses a unified fabric to create a centrally managed but physically distributed virtual chassis that supports both blade and rack servers. Each Cisco UCS instance unifies computing, networking, management, virtualization, and storage access into a radically simplified, integrated architecture. Unlike most blade configurations, in Cisco UCS all switching (Ethernet and Fibre Channel) has been removed from each blade chassis to reduce cost and simplify configuration and management. Instead of containing network switches, each Cisco UCS blade chassis supports up to two Cisco UCS fabric extenders that communicate with up to two Cisco UCS fabric interconnects (smart switches).

Cisco UCS Fabric Extenders

Cisco UCS fabric extenders bring the I/O fabric from the system's fabric interconnects to the Cisco UCS blade server chassis. They support a lossless and deterministic Ethernet and Fibre Channel over Ethernet (FCoE) fabric to connect all blades and chassis in Cisco UCS. Each fabric extender is logically a distributed line card that does not perform any switching and is managed as an extension of the fabric interconnects. This approach removes switching from within the chassis, reducing overall infrastructure and management complexity and the total number of switches needed. It

allows a single Cisco UCS platform to scale to many chassis without increasing the number of hops when moving from chassis to chassis. This radically simplified architecture delivers consistency and leadership in network performance.

Cisco UCS fabric extenders also eliminate the need for management modules. With Cisco® SingleConnect technology, Cisco UCS fabric extenders bring the management network into the blade chassis and distribute its connections to each blade server's integrated management controller and environmental management components.

Cisco UCS blade chassis typically are configured with two fabric extenders, bringing up to 160 Gbps of I/O bandwidth to each 8-slot chassis along with high availability and resource utilization through active-active uplinks. The Cisco unified fabric is built on industry standards, with Cisco innovations embodied in silicon to accelerate performance. This end-to-end solution helps business applications run better and faster.

Cisco UCS Fabric Interconnects

The Cisco UCS fabric interconnects provide a single point of management and connectivity for the entire system. All rack servers and blade chassis are attached to fabric interconnects through fabric extenders, making them part of a single, highly available management domain. Cisco UCS fabric interconnects provide uniform access to both LAN and SAN connectivity for all servers that can be dynamically configured.

Traffic Flow Within a Chassis and Between Chassis

When trying to reduce latency between servers, the best-case scenario for any vendor is a reduced hop count for data communicated between servers. For the Dell PowerEdge with the 10-Gbps immediately available ("plug-and-play") Dell PowerEdge M I/O Aggregator, that best case can be achieved only by keeping communication within a single 16-blade chassis, as shown by path X in Figure 1. This requirement limits workload placement flexibility when the performance of data exchanges between workloads, such as high-availability clustering, is a concern. Although Dell calls the Dell PowerEdge M I/O Aggregator a module, it contains a switch application-specific integrated circuit (ASIC) that performs switching functions; therefore, it is a switch.

With Cisco UCS, a single virtual blade chassis is created, with consistent, deterministic latency regardless of where the actual blade servers are located. This unique fabric solution creates the largest flat, Layer 2 network for intrablade communication with support for 160 blade servers, with optimized traffic regardless of physical chassis location. The tests performed validate this claim by **only** testing the Cisco UCS chassis-to-chassis latency and performance as shown in Figure 1, path A, rather than testing an intrachassis configuration like Dell's **best-case** configuration, in which all traffic stays within a single chassis and single switch module. As explained before, the Cisco UCS fabric extender is not a switch, but instead it is an aggregator that passes traffic to the

Cisco UCS Outperforms Dell on East-West Network Latency

fabric interconnect. Please contact your Cisco sales representative for additional test information.

Test Configurations and Results

Two distinct test cases were measured—for traffic within a single Dell chassis and for traffic across multiple Cisco UCS chassis—to determine actual blade-to-blade performance. In the first set of tests, raw latency was measured through a series of transport protocol tests (User Datagram Protocol [UDP], TCP, and round-trip time [RTT] TCP), with increasing payload sizes. In the second set of tests, virtual machine migration time was measured. This measurement was accomplished by moving loaded virtual machines from one blade to another, with all migrations performed at the same time to simulate rolling updates. To help ensure that the latency and virtual machine migration tests were relevant and equitable between manufacturers, each server was configured identically. Table 1 shows the Cisco UCS and Dell PowerEdge blade solution configurations.

All virtual machine migration tests were performed by connecting all hosts to the exact same Small Computer System Interface over IP (iSCSI)-based SAN storage logical unit number (LUN) for the VMware ESXi data store. A separate, dedicated Ethernet data path and VMware ESXi virtual switch (vSwitch) were used for iSCSI storage traffic (Figure 2). Cisco’s other east-

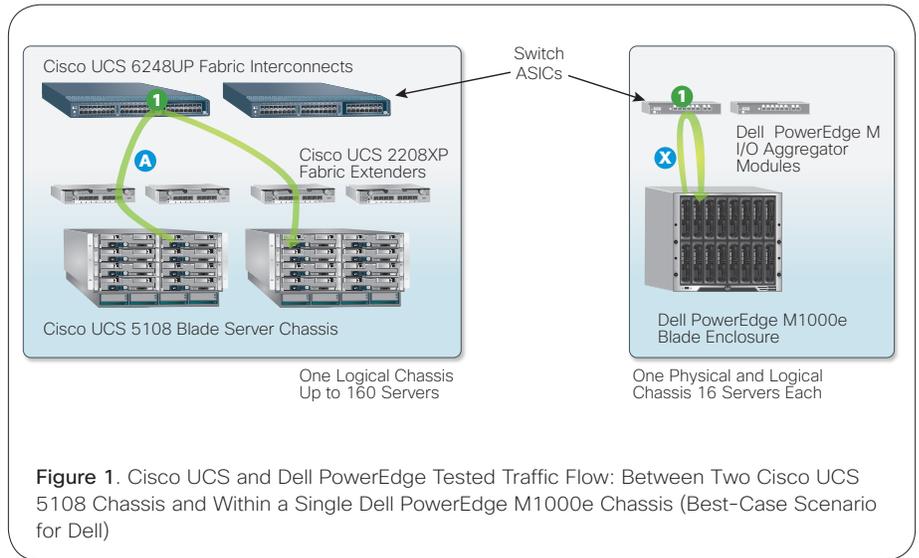


Figure 1. Cisco UCS and Dell PowerEdge Tested Traffic Flow: Between Two Cisco UCS 5108 Chassis and Within a Single Dell PowerEdge M1000e Chassis (Best-Case Scenario for Dell)

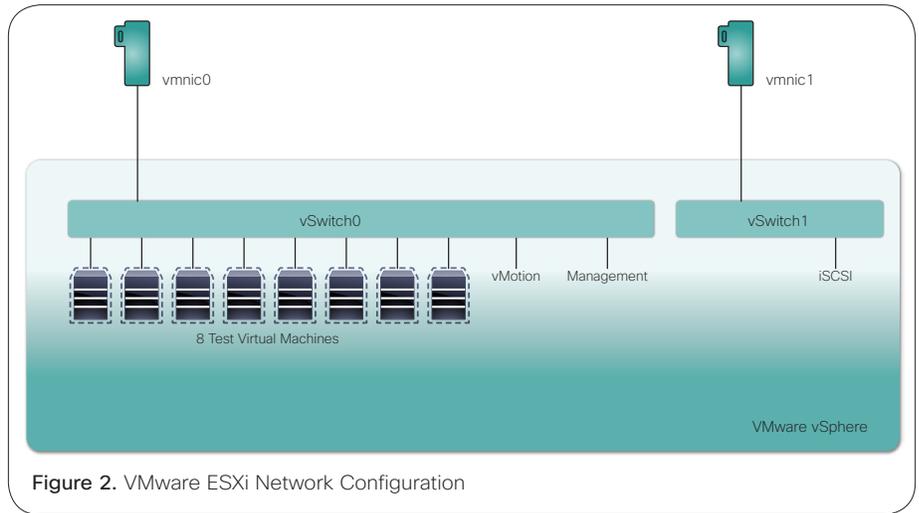


Figure 2. VMware ESXi Network Configuration

west virtual machine migration tests used multihop FCoE. Cisco tried to mimic that configuration with the Dell PowerEdge but was unable to get it to work correctly.

The same latency measurement tests and virtual machine workloads were run on each blade server solution operating with the same environmental conditions and physical configurations. Thousands of samples were collected in a variety

Cisco UCS Outperforms Dell
on East-West Network Latency

Table 1. Cisco and Dell Configurations Used for Latency Testing

System	Cisco UCS	Dell PowerEdge
Chassis used	2 Cisco UCS 5108 Blade Server Chassis	1 Dell PowerEdge M1000e Blade Enclosures
Enclosure management modules	2 Cisco UCS 6248UP 48-Port Fabric Interconnects	2 Dell Chassis Management Controllers per chassis
Internal I/O modules per chassis	2 Cisco UCS 2204XP Fabric Extenders	2 Dell PowerEdge M I/O Aggregator
Physical links from blade to chassis	1 x 10 Gigabit Ethernet uplink from fabric extender to fabric interconnect	Internal 10-Gbps switched enclosure
Maximum number of servers on a single switched network	160	16
Blade Model	Cisco UCS B200 M3 Blade Server	Dell PowerEdge M620
Form factor	Half width	Half-height
Processor	2 Intel® Xeon® processors E5-2680	2 Intel Xeon processors E5-2680
Memory	8 x 8-GB DDR3 RDIMM PC3L-12800	8 x 8-GB DDR3 RDIMM PC3L-12800
Hard disk drives	2 x 300-GB 10,000-rpm 6-Gbps RAID 1	2 x 300-GB 10,000-rpm 6-Gbps RAID 1
Network	Cisco UCS 1240 Virtual Interface Card (VIC)	Dell Broadcom NDC 57810S
Power and system profile	Maximum performance	Performance
Software	Cisco UCS B200 M3 Blade Server	Dell PowerEdge M620
Latency test host OS	SUSE Linux Enterprise Server (SLES) 11 SP2	SLES 11 SP2
Virtual machine migration test original equipment manufacturer (OEM) VMware ESXi build	VMware ESXi-5.1.0-799733-custom-Cisco-2.1.0.3	VMware ESXi-5.1.0-799733-Dell
Network VMware ESXi driver	Release 2.1.2.22	Release 1.76.54.v50.1
Migration test VMware ESXi guest OS	Microsoft Windows Server 2008 SP2	Microsoft Windows Server 2008 SP2

Cisco UCS Outperforms Dell on East-West Network Latency

of configurations. Please contact your Cisco sales representative for detailed test information.

Application Latency TCP Transactions (RTT)

Cisco loaded the SLES 11 SP2 operating system on each host, with a single network adapter configured with a single network fabric topology. The host BIOS and operating system kernel were optimized for performance on both the Cisco and Dell systems. The standard network analysis tool, Netperf, was used to measure both UDP and TCP latency and RTT TCP latency. This test mimics the HTTP protocol used by many web servers and represents typical transactional workloads. Rather than simply measuring the performance of all requests and responses on the same connection, the test establishes a new connection for every request-and-response pair. The tests measured actual blade-to-blade latency to obtain real-world results. The results shown in Figures 3, 4, and 5 are for Cisco UCS with one blade in one chassis and the other blade in another chassis (as shown by path A in Figure 1). Conversely, the Dell PowerEdge M620 with Dell PowerEdge M I/O Aggregator results are best-case scenarios with both blades located in the same chassis (as shown by path X in Figure 1).

Cisco tested a full range of payload sizes mimicking all application types, using a 9000-byte maximum transmission unit (MTU), to expose any port-buffering problems that might occur with large payload sizes (Dell

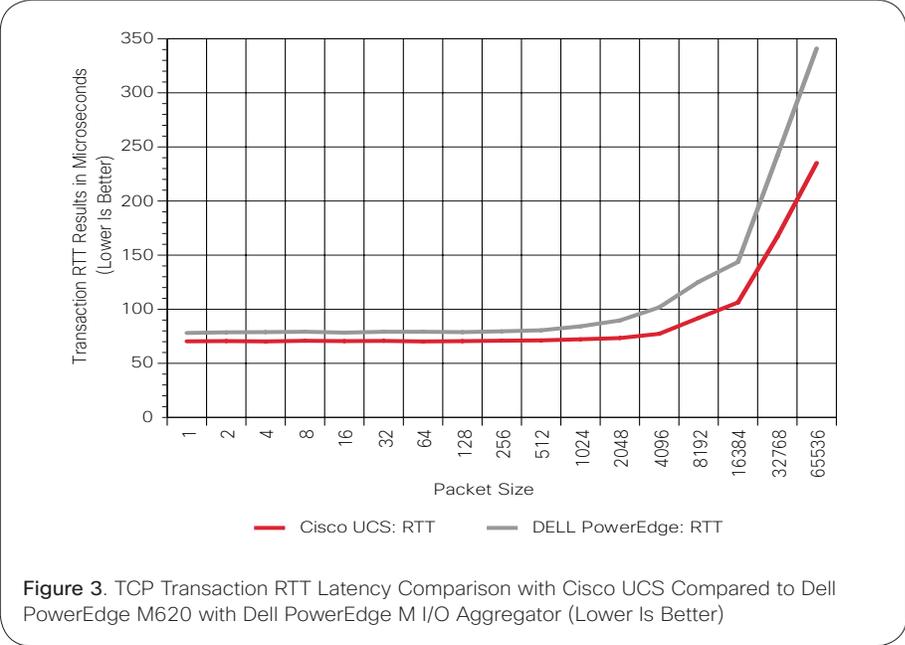


Figure 3. TCP Transaction RTT Latency Comparison with Cisco UCS Compared to Dell PowerEdge M620 with Dell PowerEdge M I/O Aggregator (Lower Is Better)

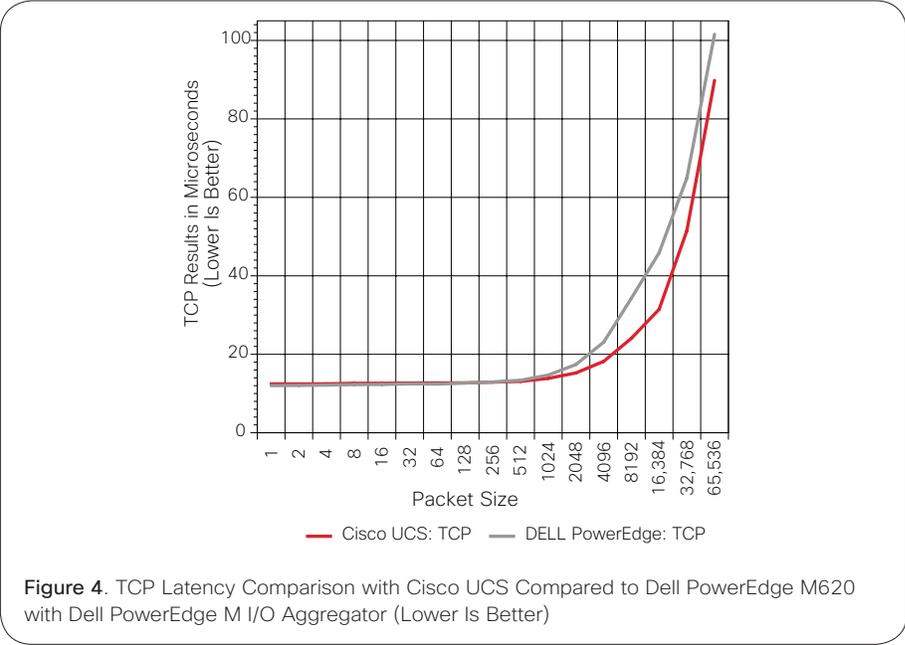


Figure 4. TCP Latency Comparison with Cisco UCS Compared to Dell PowerEdge M620 with Dell PowerEdge M I/O Aggregator (Lower Is Better)

Cisco UCS Outperforms Dell on East-West Network Latency

tested only small payloads of up to 512 bytes). Payload sizes greater than 1500 bytes tend to expose switch port-buffering problems. The lack of deep buffers results in higher latency and lower transactional performance times. The results show that Cisco has better network packet processing and buffering performance, on both hosts and switches, than Dell.

The results of the latency tests revealed the following:

- Cisco UCS demonstrated lower latency than Dell PowerEdge for **every** test group and every packet size.
- As packet sizes increased in each test, the Dell PowerEdge disadvantage also increased compared to Cisco UCS.
- As Figure 3 shows, Dell PowerEdge has between 11 and 45 percent greater latency than Cisco UCS.

Real-world latency is more than theoretical assumptions. Cisco remains the leader in ASIC and network design optimization for the end-to-end network stack, enabling business applications and virtual environments to perform better.

Virtual Machine Migration Timing

Cisco tested the amount of time required to migrate a virtual machine from one blade to another under a heavy workload. The VMware ESXi hypervisor was used, and virtual machine allocated memory did not

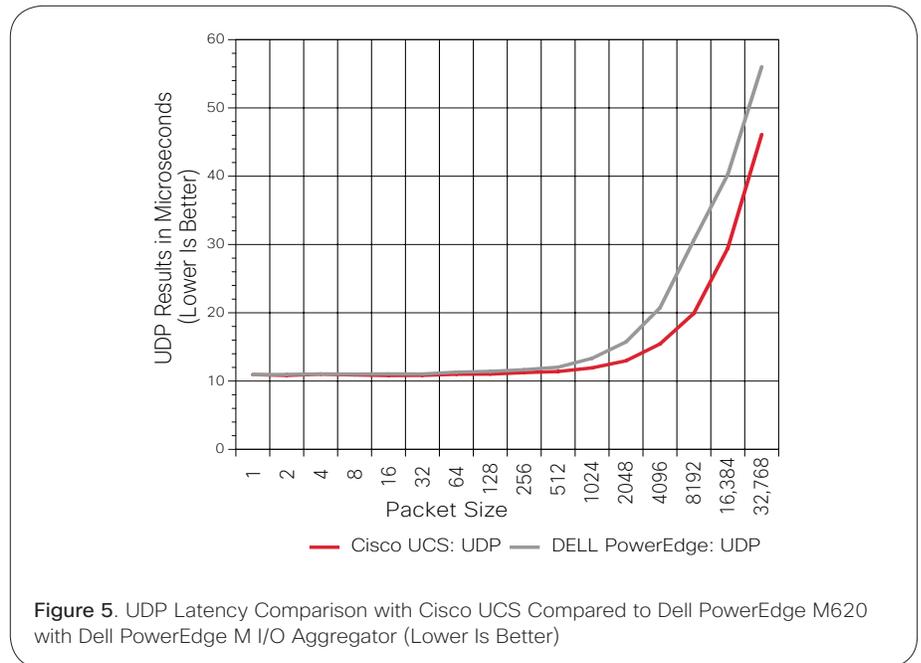


Figure 5. UDP Latency Comparison with Cisco UCS Compared to Dell PowerEdge M620 with Dell PowerEdge M I/O Aggregator (Lower Is Better)

exceed the total host memory. The Microsoft Windows Server 2008 SP2 guest operating system was loaded on each virtual machine, with Prime95 testing software run to push both the memory and processors to their limits. A single network adapter configured with a single network fabric topology was used. The performance profiles configured in the BIOS for both the Cisco and Dell systems were optimized identically.

Cisco ran a minimum of 1000 trials of each of the two separate virtual machine migration tests to capture variances of vCenter completion times. This reflects the way that actual Cisco UCS customers use virtual machines, with more than one

virtual machine running on a host and sharing resources. These configurations really stressed the network during the migrations because all 64 GB of host memory was consumed before the virtual machines were transferred over the Ethernet connection. The first test was configured using four virtual machines, each with 16 GB of memory, and the second test used eight virtual machines, each with 8 GB of memory. Variations in the amount of time needed to complete the migrations occur because the memory in the system is constantly changing.

VMware ESXi views all 160 servers in a Cisco UCS instance as part of a single management domain by default. Therefore, all Cisco UCS tests were

Cisco UCS Outperforms Dell on East-West Network Latency

conducted between multiple chassis (path A in Figure 1). Conversely, the best-case scenario for the Dell PowerEdge was used with all traffic remaining in one chassis (path X in Figure 1).

The results of the virtual machine migration tests revealed the following:

- Cisco UCS demonstrated better performance (faster migration times) than Dell PowerEdge for every virtual machine migration group size tested: 8 and 16 GB (Figure 6).
- As the virtual machine size and network load increases, the Cisco UCS performance advantage also increases.

Conclusion

When well-designed, customer-inspired, and rigorous performance tests are used, the highly efficient architecture of Cisco UCS with unified fabric consistently delivers lower network latency and higher network performance to business applications. The radically simplified and cost-effective design removes complexity and network hops from within its large virtual chassis by using Cisco UCS fabric extenders and fabric interconnects to quickly move data from blade to blade within a single chassis and between multiple chassis—all with just one hop. Contrary to Dell's claims,

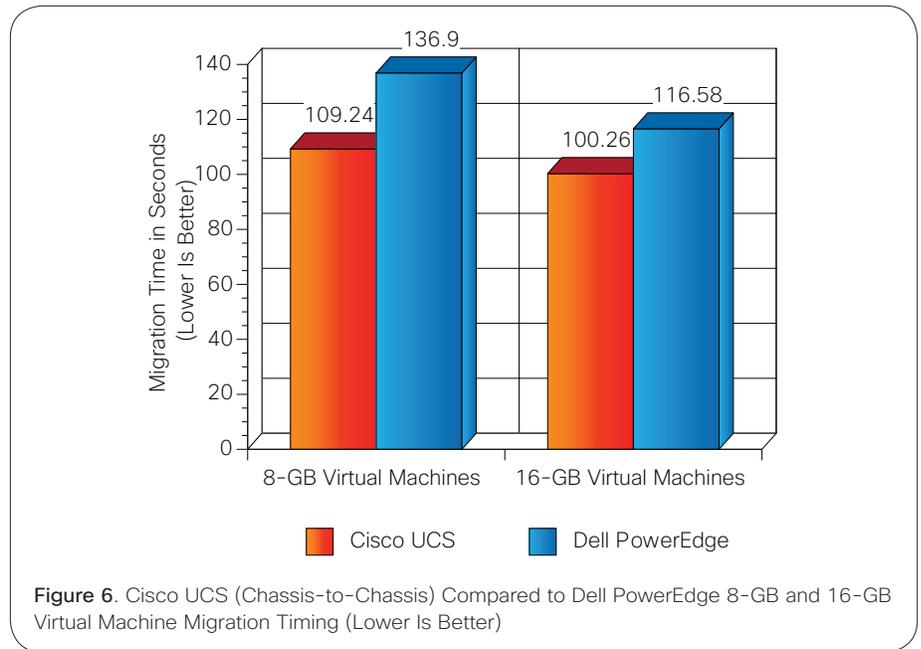


Figure 6. Cisco UCS (Chassis-to-Chassis) Compared to Dell PowerEdge 8-GB and 16-GB Virtual Machine Migration Timing (Lower Is Better)

Cisco UCS, even chassis to chassis, delivers lower latency and greater virtual machine migration performance than Dell PowerEdge M620 blade servers using the Dell PowerEdge M I/O Aggregator.

For More Information

- For more information about the performance tests that produced these results, please contact your Cisco sales representative.
- For more information about Cisco UCS, please visit <http://www.cisco.com/go/ucs>.

- For more information about Cisco UCS fabric extenders, please visit <http://www.cisco.com/en/US/products/ps10265/products.html>.
- For more information about Netperf, please visit <http://www.netperf.org>.
- For more information about Cisco UCS east-west performance compared with HP and IBM performance, please visit http://www.cisco.com/c/en/us/products/servers-unified-computing/industry_benchmarks.html#bm



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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