Cisco UCS Outperforms HP Blade Servers on East-West Latency

Performance Brief
July 2013

Highlights

Better performance

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

Faster transaction times

- Cisco UCS demonstrated lower latency than the HP BladeSystem c7000 with Virtual Connect 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 HP Virtual Connect FlexFabric and Flex-10/10D for every virtual machine migration group size tested.

Architecture matters. Businesses get faster applications and virtualized environments with more dependable latency with Cisco Unified Computing System™ (Cisco UCS®).

Cisco UCS and the HP BladeSystem 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 a Cisco UCS domain with consistent network performance. Important benefits of the Cisco UCS architecture include reduced latency across all I/O operations and the consistency of I/O latency between physical servers.

Cisco has proven HP's claims about Cisco UCS latency to be incorrect: Cisco UCS, not HP, has the lowest latency and the fastest virtual machine migration times for every use-case test Cisco performed against the HP BladeSystem c7000 with Virtual Connect. This document explains how the Cisco UCS unified fabric works and presents the results of comprehensive latency tests of Cisco UCS and HP BladeSystem c7000 with Virtual Connect.

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 domain unifies computing, networking, management, virtualization, and storage access into a radically simplified, integrated architecture. Unlike in other x86-architecture blade servers, all Cisco UCS connectivity (including Ethernet, Fibre Channel, and management networking), is handled centrally by the Cisco UCS fabric interconnects, reducing cost and simplifying configuration and management. In the Cisco UCS architecture, up to two Cisco UCS fabric extenders bring the unified fabric into each blade server chassis. Each fabric extender communicates with a Cisco UCS fabric interconnect.

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 a Cisco UCS domain into one cohesive system. 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 domain to scale to many chassis, keeping the same one-hop count from any one chassis to any other chassis. This radically simplified architecture delivers consistency and dependability in network performance and has made Cisco UCS a leader in the blade server market.

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 are typically 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

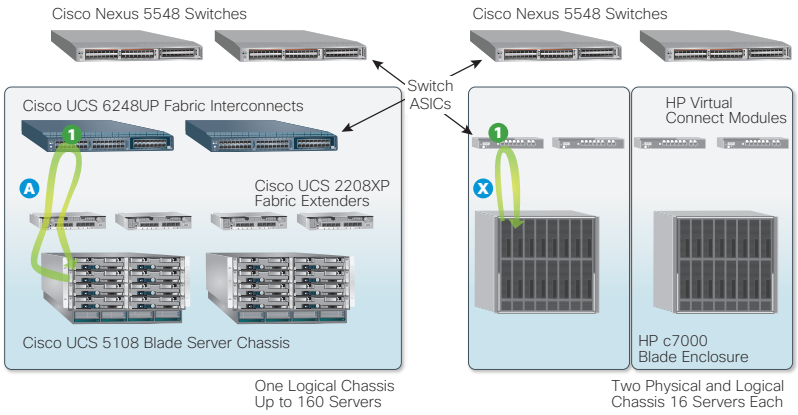

Figure 1. Cisco UCS and HP BladeSystem Traffic Flow To and From a Blade in the Same Chassis

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 Interconnect

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 HP BladeSystem, 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 performance is a concern. For Cisco UCS, traffic everywhere in the domain is optimized and uses a path similar to path A in Figure 1 regardless of physical chassis location. As explained before, the Cisco UCS fabric extender is not a switch, but rather it is an aggregator that passes traffic to the fabric interconnect.

When communications must travel from one chassis to another, the architectural differences become clear. The Cisco UCS fabric interconnect centrally manages the traffic in Cisco UCS as well as the traffic coming into and going out of the system. For traffic that is moving from one chassis to another in the same Cisco UCS, there is no need to exit Cisco UCS and have packets routed through another external switch

(see path B in Figure 2). This feature demonstrates one of the ways in which Cisco UCS functions a single virtual chassis and provides the flexibility to place workloads anywhere in Cisco UCS with the assurance of consistent network performance.

Conversely, following HP best practices for configuring HP BladeSystem chassis, networking, and servers, traffic between HP BladeSystem chassis must first travel through the HP Virtual Connect modules (see path Y in Figure 2). Then the traffic is routed through an external switch and back through the second HP Virtual Connect module, forcing three network hops and greater latency. Although HP calls Virtual Connect a module, it contains a switch application-specific integrated circuit (ASIC) that performs switching functions; therefore, it is a switch.

HP supports stacking connections between HP Virtual Connect modules in up to four chassis. Because Cisco UCS multichassis results were superior to HP's best-case single-chassis results, stacking measurements are not included in this document. Stacking connections are typically not used by customers because they consume ports that could otherwise be used as uplinks, they do not also support Fibre Channel connections, and they cause loss of visibility into internal spanning-tree traffic paths. Please contact your Cisco sales representative for additional test results and detailed information.

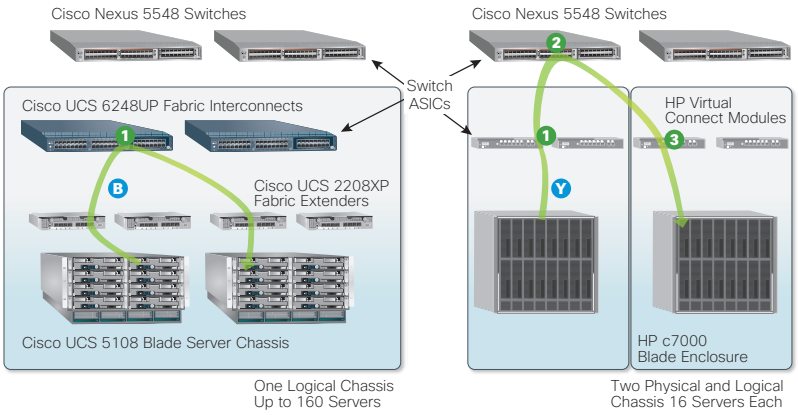

Figure 2. Cisco UCS and HP BladeSystem Traffic Flow To and From a Blade in Different Chassis

Test Configurations and Results

Two distinct test cases were measured—for traffic within a single chassis and across multiple 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. For latency and virtual machine migration tests to be relevant and equitable between manufacturers, each server was configured identically. Table 1 shows the Cisco UCS and the HP BladeSystem c7000 blade solution configurations.

The same latency measurement benchmark and virtual machine workloads were run on each blade server solution operating in the same environmental conditions. Thousands of samples were collected in a variety of fabric configurations, and a subset of these results is reported here. Please contact your Cisco sales representative for additional test results and detailed information.

Application Latency TCP Transaction Round-Trip Time

Cisco loaded the SUSE Linux Enterprise Server 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 HP systems. The standard network analysis tool, Netperf, was used to measure both UDP and TCP latency and RTT

Table 1. Cisco and HP Configurations Used for Latency Testing

Enclosure	Cisco UCS 5108 Blade Server Chassis	HP BladeSystem c7000
Chassis used	2 Cisco UCS 5108 Blade Server Chassis	2 HP c7000 Chassis
Enclosure management modules	2 Cisco UCS 6248UP 48-Port Fabric Interconnects	2 HP Onboard Administrator Modules per chassis
Internal I/O modules per chassis	2 Cisco UCS 2204XP Fabric Extenders	2 HP Virtual Connect FlexFabric 10Gb Modules or 2 HP Virtual Connect Flex-10/D Modules (both were tested)
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**	**HP BladeSystem BL460c Gen8**
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 drive	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)	HP FlexFabric 10Gb Adapter
Software	**Cisco UCS B200 M3 Blade Server**	**HP BladeSystem BL460c Gen8**
Latency test host OS	SUSE Linux Enterprise Server 11 SP2	SUSE Linux Enterprise Server 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-HP-5.34.23
Virtual machine migration test VMware ESXi network driver	Release 2.1.2.22	Release 4.2.327.0
Virtual machine migration test VMware ESXi guest operating system	Microsoft Windows Server 2008 R2 SP1	Microsoft Windows Server 2008 R2 SP1

TCP latency. Tests measured actual blade-to-blade latency to obtain real-world results (Figure 3). The results of the latency tests revealed:

- Cisco UCS demonstrated lower latency than the HP BladeSystem c7000 with Virtual Connect for every test group and every packet size (UDP, TCP, and RTT TCP). For full test details, please contact your Cisco sales representative.
- As packet sizes increased in each test, the HP BladeSystem c7000 with Virtual Connect disadvantage also increased compared to Cisco UCS.
- As Figure 3 shows, the performance is almost identical for both single-chassis tests and multichassis tests for Cisco UCS. In addition, the HP BladeSystem c7000 with Virtual Connect has between 31 and 70 percent more latency than Cisco UCS.
- With the HP BladeSystem c7000 with Virtual Connect, after traffic leaves the chassis, latency increases dramatically.

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.

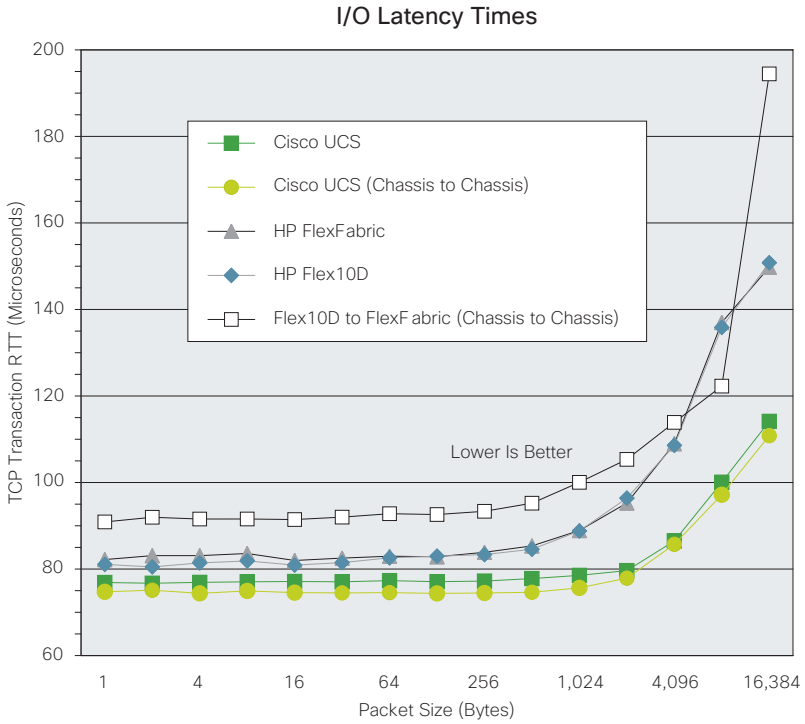

Figure 3. I/O Latency Times

Virtual Machine Migration Timing

Cisco tested the amount of time required to migrate a virtual machine from one blade to another. The VMware ESXi hypervisor was used, and virtual machine allocated memory did not exceed the total host memory. The Microsoft Windows Server 2008 SP1 guest operating system was loaded on each host, 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. Both the Cisco and HP systems were tested as they came configured, straight out of the box with no additional optimizations performed.

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

The results of the virtual machine migration tests revealed:

- Cisco UCS demonstrated better performance (faster migration times) than HP Virtual Connect FlexFabric and Flex-10/10D for every virtual machine migration group size tested: 4 and 8 GB (shown in Figure 4) and 16 GB.

- As the virtual machine size and network load increases, the Cisco UCS performance advantage also increases.
- HP Virtual Connect Flex-10/10D and HP Virtual Connect FlexFabric displayed severe limitations in bandwidth performance. The average bandwidth data captured from HP Virtual Connect modules compared to Cisco UCS was as follows:
  - Cisco UCS averaged 9.5 Gbps
  - HP Virtual Connect Flex-10/10D averaged 6 Gbps
  - HP Virtual Connect FlexFabric averaged 5.5 Gbps

Conclusion

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 HP's claims, Cisco UCS delivers lower latency and greater performance than HP BladeSystem c7000 with Virtual Connect.

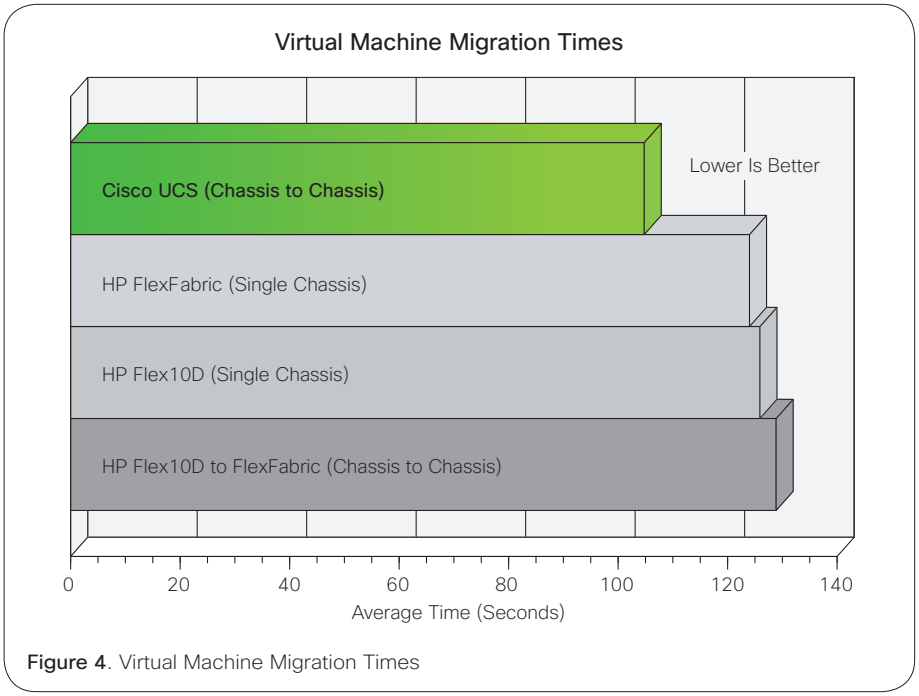

Figure 4. Virtual Machine Migration Times

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 performance, please visit http://www.cisco.com/go/ucsatwork.

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**.

LE-40201-00 07/13