

Cisco UCS Outperforms IBM Flex System Blades on East-West Network Latency



Performance Brief
July 2013

Highlights

Better performance

- The Cisco Unified Computing System™ (Cisco UCS®) delivered consistent lossless, low-latency, high-workload throughput.

Faster virtual environments

- Cisco UCS delivered better performance than IBM (faster virtual machine migration times) for every group size tested.

Faster transaction times

- Cisco UCS demonstrated lower latency in every transactional latency test with a payload size larger than 256 bytes run against IBM Flex System.

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

Cisco UCS and IBM Flex System have significantly different architectures. These differences help Cisco UCS deliver consistency, visibility, and portability across servers, regardless of whether the servers are physical or virtual. Cisco's innovation is building the end-to-end fabric at the core of Cisco UCS which enables business applications to run better and faster.

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

Contrary to IBM's claims, Cisco UCS performed better than the IBM Flex System in every use-case test. Cisco UCS demonstrated lower latency in every transactional test performed with packets greater than 256 bytes, and faster virtual machine migration timing in every test.

This document explains how Cisco UCS unified fabric works and presents results from comprehensive latency tests of Cisco UCS and the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch.

How Cisco UCS Unified Fabric Works

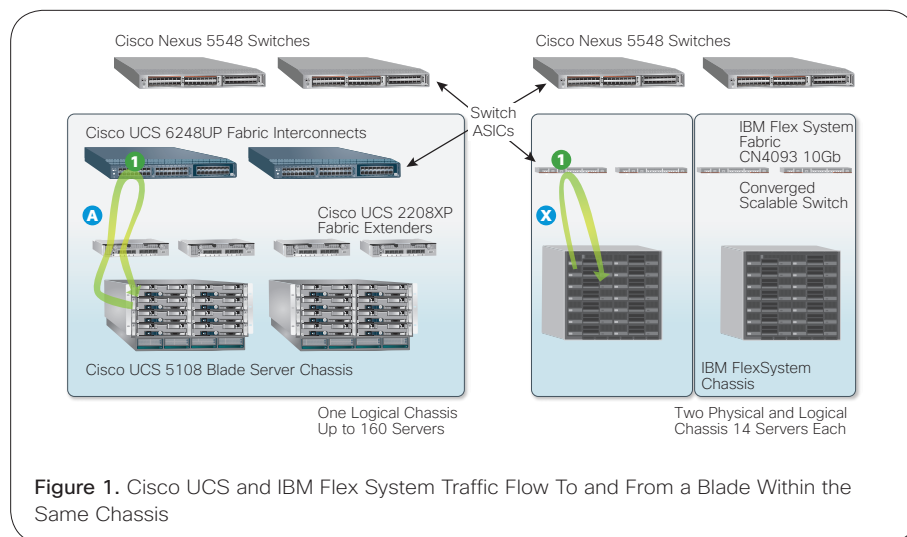
Unlike other x86 blade designs, all Cisco UCS switching (Ethernet, Fibre Channel, and management) has been removed from each blade chassis, delivering cost savings and ease of configuration and management. Cisco UCS uses a unified fabric to create a single virtual chassis that can contain both blade and rack servers. Each Cisco UCS domain unifies computing, networking, management, virtualization, and storage access into a simplified, integrated architecture. Cisco® Unified Fabric is built on industry standards, with Cisco innovations that accelerate end-to-end performance. The focus of the Cisco UCS architecture is on helping business applications run better, faster, and more dependably than with any other design.

Cisco UCS Fabric Interconnect

The Cisco UCS fabric interconnects provide the management and communication backbone for the Cisco UCS B-Series Blade Servers and Cisco UCS blade server chassis. Cisco UCS C-Series Rack Servers can also be integrated into and managed from the Cisco UCS fabric interconnects. All servers and chassis that are attached to fabric interconnects are part of a single, highly available management domain. Cisco UCS fabric interconnects provide the flexibility to support LAN and SAN connectivity for all servers within the domain at configuration time.

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 high-speed, lossless, low-latency 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 as well as the total number of switches required. This architecture allows a single Cisco UCS domain to scale across many chassis without increasing the number of hops when moving from chassis to chassis. The Cisco UCS architecture delivers consistency and dependability in network performance and has made Cisco UCS the leader in blade server architecture.



Cisco UCS fabric extenders also eliminate the need for the separate (and required) management modules with which other blade chassis designs are burdened. 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, providing up to 160 Gbps of I/O bandwidth to each 8-slot chassis along with high availability and resource utilization through active-active uplinks.

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 IBM Flex System, that best case can be achieved only by keeping communication within a single 14-blade chassis, as shown by path X in Figure 1. This requirement limits workload placement flexibility when performance is a concern. With Cisco UCS, all traffic everywhere in the domain is optimized and uses a path similar to path A in Figure 1. This path is used regardless of physical chassis location. Cisco UCS fabric extenders are not switches; they are aggregators that pass traffic to the fabric interconnects.

When communications must travel from one chassis to another, the architectural differences become clear. The Cisco UCS fabric interconnects centrally manage the traffic in the Cisco UCS domain as well as the traffic entering and leaving the domain.

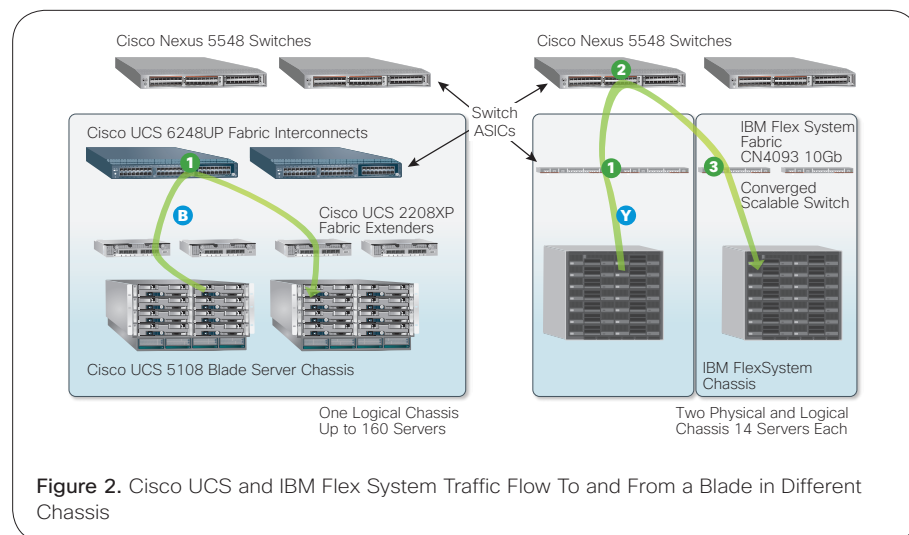
For traffic that is going from one chassis to another within the same Cisco UCS domain, there is no need for the traffic

to exit the Cisco UCS domain. Cisco UCS does not need to have packets routed through a supplementary external switch as IBM Flex System does (Figure 2). This capability illustrates one of the ways in which Cisco UCS functions as a single virtual chassis. This architecture provides the flexibility to place workloads anywhere within the Cisco UCS domain and be assured of consistent, dependable network performance.

Following IBM's best practice for configuring the IBM Flex System chassis, traffic between two or more chassis must pass through multiple switches, which adds hops. First, data travels through the IBM Flex System Fabric CN4093 switch in the first chassis (see path Y in Figure 2), where it is sent to an external switch. The external switch then sends the traffic to a different IBM Flex System Fabric CN4093 switch in the second chassis, forcing three network hops and greater latency.

Test Configurations and Results

Two distinct test cases were measured against both of these scenarios—for traffic within a single chassis and for traffic 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 [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 IBM Flex System blade solution configurations.

The same latency measurement benchmarks 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 IBM systems. The standard network analysis tool, Netperf, was used to measure both UDP and TCP latency, as well as RTT TCP latency. Tests measured actual blade-to-blade latency to obtain real-world results (Figure 3).

The results of the latency tests revealed:

- The Cisco UCS 5108 enclosure configured with Cisco UCS B200 M3 blades demonstrated lower latency in every test with greater than 256 bytes and every chassis-

Cisco UCS Outperforms IBM Flex System
Blades on East-West Latency

Table 1. Cisco and IBM Configurations Used for Latency Testing

Enclosure	Cisco UCS 5108 Blade Server Chassis	IBM Flex System Chassis
Chassis used	2 Cisco UCS 5108 Blade Server Chassis	2 IBM Flex System chassis
Enclosure management modules	2 Cisco UCS 6248UP 48-Port Fabric Interconnects	2 chassis management modules per chassis
Internal I/O modules per chassis	2 Cisco UCS 2204XP Fabric Extenders	2 IBM Flex System Fabric CN4093 10Gb Converged Scalable Switches
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	14
Blade Model	Cisco UCS B200 M3 Blade Server	IBM Flex System x240 Blade
Form factor	Half width	Half width
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)	IBM Virtual Fabric 10-Gbps 2-port adapter
Software	Cisco UCS B200 M3 Blade Server	IBM Flex System x240 Blade
Latency test host operating system	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-IBM-20120919
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

Cisco UCS Outperforms IBM Flex System Blades on East-West Latency

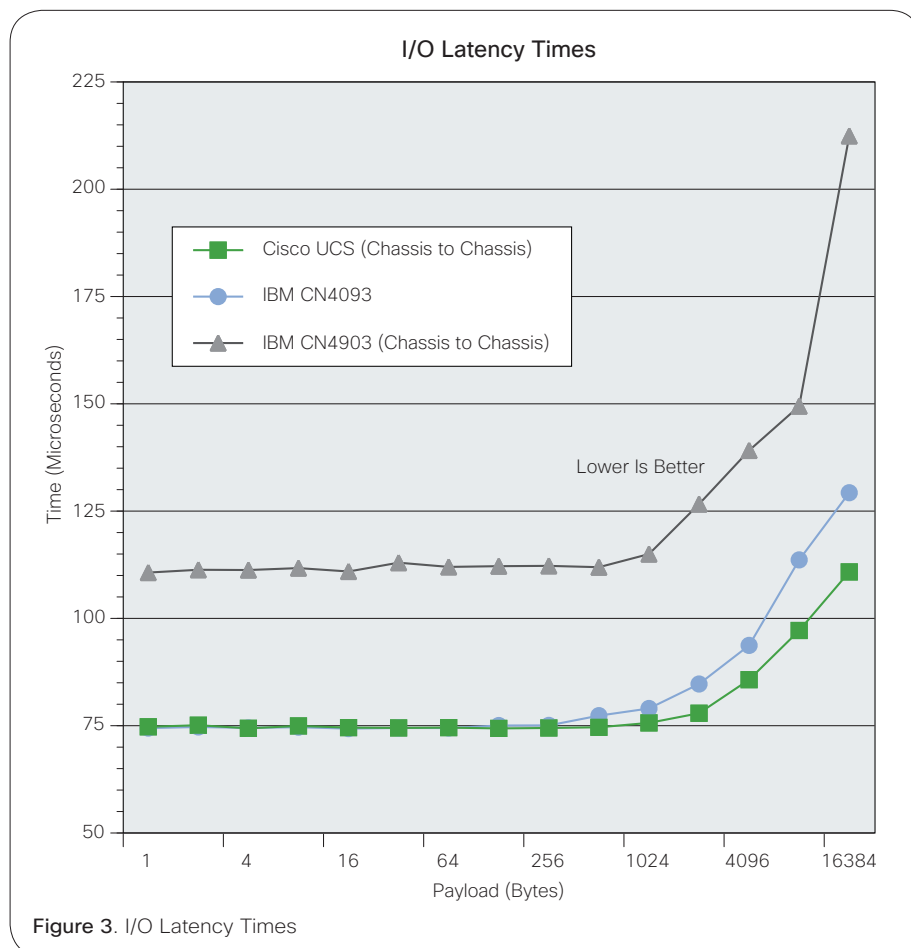
to-chassis test compared to the IBM Flex System x240 blades, using similarly configured hardware, BIOS, and OS settings for each test group and packet size (UDP, TCP, and RTT TCP). For complete test details, please contact your Cisco sales representative.

- As packet sizes increased in each test, the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch disadvantage also increased compared to Cisco UCS.
- As Figure 3 shows, the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch has between 17 and 92 percent more latency, at 16 KB packet sizes, than Cisco UCS.
- With the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch, after traffic leaves the chassis, latency increases dramatically.

Real-world latency is more than theoretical assumptions. Cisco remains the leader in application-specific integrated circuit (ASIC) and network design optimization for the end-to-end network stack, which enables 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. The VMware ESXi hypervisor was used, and the 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, and Prime95 testing software was used 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 IBM systems were tested as they came configured, straight out of the box with no additional optimizations performed.

VMware ESXi views all 160 servers within a Cisco UCS instance as part of a single management domain by default. Therefore, all Cisco UCS migration tests were conducted between multiple chassis. The results of the virtual machine migration tests revealed:

- Cisco UCS demonstrated better performance (faster migration times)

Cisco UCS Outperforms IBM Flex System Blades on East-West Latency

than the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch for every virtual machine migration group size tested: 4 GB, 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.
- Cisco UCS was 22 percent faster than the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch.

Conclusion

The highly efficient architecture of Cisco UCS with unified fabric consistently delivers lower network latency and higher network performance for 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 Cisco UCS 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 IBM's claims, Cisco UCS delivers overall lower and greater performance than the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch.




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.