

Improve Network Throughput by 38 Percent with the Cisco UCS M81KR Virtual Interface Card

What You Will Learn

The revolutionary Cisco[®] UCS M81KR Virtual Interface Card (VIC) helps increase application performance and consolidation ratios, with 38 percent greater network throughput, complementing the latest increases in Cisco Unified Computing System[™] CPU performance and memory capacity. The virtual interface card and the Cisco Unified Computing System together set a new standard for balanced performance and efficiency.

Superior Network Performance for Virtualized Environments


Server virtualization is being widely deployed to increase utilization and reduce power, cooling, and floor space requirements. Companies implementing server virtualization are gaining significant return on their investments as well as reducing the total cost of ownership (TCO) for the infrastructure. In addition, the move to a virtualized infrastructure is giving companies exceptional flexibility and agility, creating pools of resources that are available for immediate deployment, decreasing time to production, and increasing competitive advantage.

However, with server virtualization, each physical system now has multiple workloads sharing compute, storage, and network resources, pushing existing networking infrastructures and architectures well beyond their current capabilities. For example, when multiple virtual servers are moved into a single system, network visibility and monitoring of traffic between the virtual servers becomes obstructed. While the traditional security architectures that use physical switches to protect each physical machine are still viable, new methods for protecting the network connections within a physical system between virtual servers are needed. Management of quality of service (QoS) on a per-virtual machine basis is nearly impossible without the capability to access and see the network traffic all the way to the virtual machine. These challenges are keeping IT departments from migrating their most important business-critical applications to the virtual server environment. To help meet the challenges of this new, shared environment within a single system, software switch solutions that reside on the hypervisor virtualization layer have been developed and deployed, but always with a cost in performance and management overhead.

The Cisco UCS M81KR VIC is the first hardware-based network adapter solution that removes the need for a software virtual switch (vSwitch), freeing server CPU cycles for improved application performance and support for higher virtual machine densities. The Cisco UCS M81KR VIC enables network management and security on a per-virtual machine basis, just as for physical network interfaces, allowing network administrators to use current security and network policy, reducing complexity, and enabling easier management and lower TCO.

The Cisco UCS M81KR VIC is a low-latency, high-performance mezzanine card for the Cisco Unified Computing System that can saturate two 20 Gbps (bidirectional) unified fabric connections, achieving 38 percent better throughput and more than 40 percent better runtime performance than a hypervisor-based software switch solution (Figure 1). This document describes the Cisco UCS M81KR VIC and the test environment used to achieve the results shown in Figure 1.


Figure 1. Throughput Performance Comparison Between a Hypervisor vSwitch (Software), a Cisco UCS M81KR VIC Through the Hypervisor, and a Cisco UCS M81KR VIC Bypassing the Hypervisor



Cisco USC M81KR Virtual Interface Card

The Cisco UCS M81KR VIC is a mezzanine card for the Cisco UCS B-Series Blade Servers with connectivity to two 10-Gbps backplane networks for up to 40-Gbps bidirectional bandwidth. Each card can support up to 128 virtual interfaces configured as network interface cards (NICs) or Fibre Channel host bus adapters (HBAs), managed through Cisco UCS Manager (Figure 2).

Figure 2. Cisco UCS M81KR Virtual Interface Card Supporting 128 Virtual Interfaces



The Cisco UCS B200 M1 and M2 Blade Servers are half-width blade servers with support for one Cisco UCS M81KR VIC; and the Cisco UCS B250 M2 Extended-Memory Blade Server and Cisco UCS B440 High-Performance Blade Server are full-width blade servers with support for two virtual interface cards connecting to a high-performance PCI Express (PCIe) x16 bus for superior throughput. Each card has two full-duplex connections to each of the lossless 10 Gbps unified fabric connections on the Cisco Unified Computing System chassis backplane.

The virtual interfaces defined in the Cisco UCS M81KR VIC are available for connection to virtual machines or non-virtualized applications running within each blade. For these tests, all connections were to virtual machines. Cisco VN-Link technology, implemented in hardware within the card, enables connectivity from each virtual machine to a virtual interface within the parent fabric interconnect (Figure 3). Virtual interfaces are associated with a physical port, but the relationship is not fixed and can change as virtual machines move between servers. Cisco VN-Link technology provides network visibility from each virtual machine to the parent fabric, enabling consistent network and security operations for both physical and virtual servers.

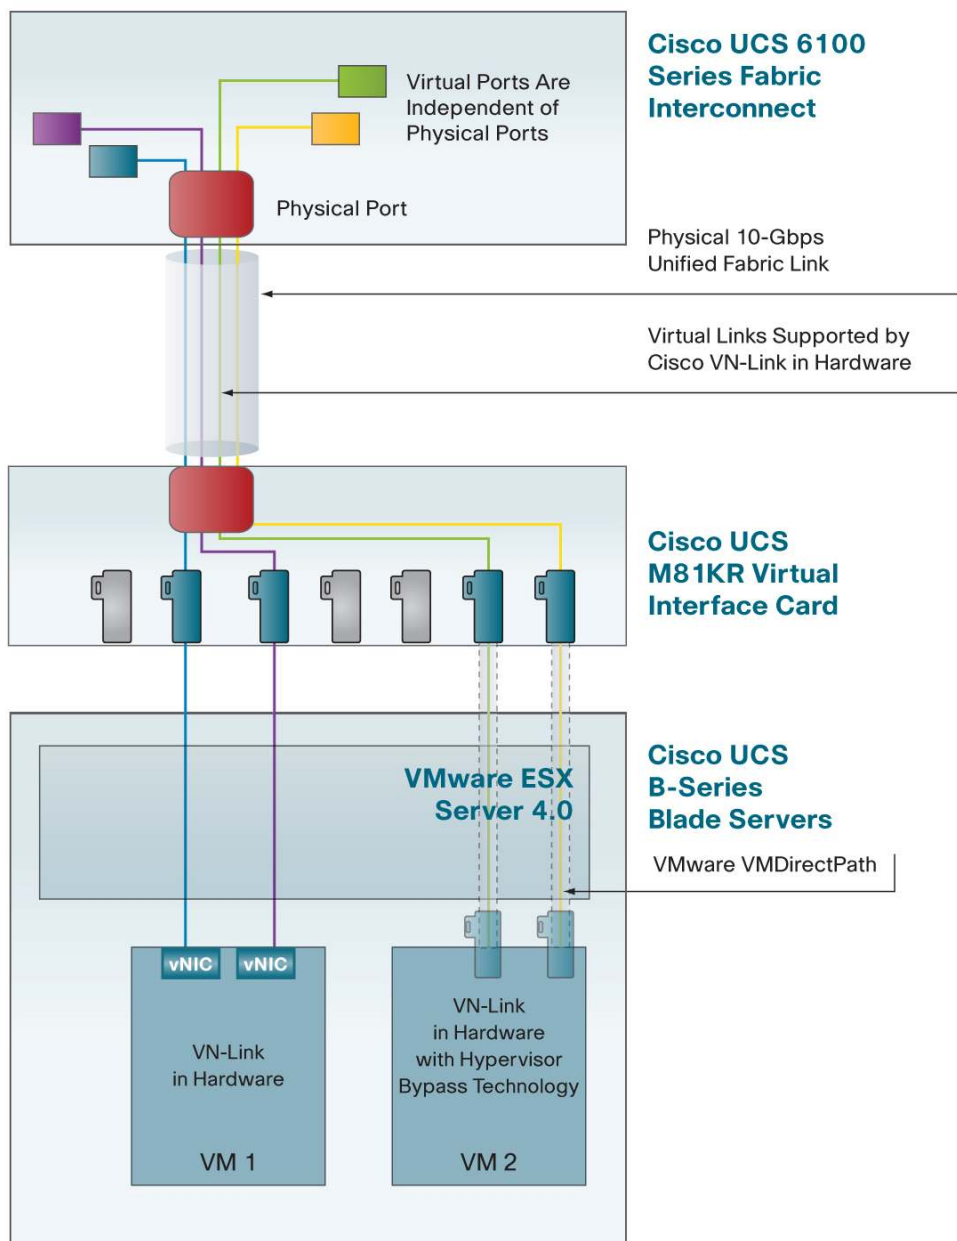
The virtual interfaces in the Cisco UCS M81KR VIC are configured and managed centrally through Cisco UCS Manager using service profiles that specify MAC addresses, QoS, and adapter policies. Cisco UCS Manager collects network policy information such as VLAN configuration directly from the network, enabling consistent management of attributes, such as QoS, VLANs, and access control lists (ACLs) throughout the infrastructure from a single point.

When the system is running VMware vSphere software, the Cisco UCS M81KR VIC can be accessed in two ways: through Cisco VN-Link in hardware and through Cisco VN-Link in hardware with hypervisor bypass technology.

In the first approach, as with any other NIC or HBA, the hypervisor manages all the network connections. As shown on the left in Figure 3, using Cisco VN-Link in hardware the virtual interfaces are configured as described earlier using Cisco UCS Manager. Those interfaces are seen and managed by the VMware hypervisor as if they were physical interfaces.

The second approach, shown on the right in Figure 3, uses Cisco VN-Link in hardware with hypervisor bypass technology. In this approach, the Cisco UCS M81KR VIC is accessed directly by the operating system running in a virtual machine using VMware VMDirectPath.

Figure 3. Internal Cisco UCS Network Architecture Utilizing the Cisco UCS M81KR VIC



In the case in which the Cisco UCS M81KR VIC is accessed through VMware ESX Server 4.0, the hypervisor creates a virtual NIC (vNIC) for the virtual machine, assigns the virtual machine vNIC MAC address, and then links to the interface presented by the Cisco UCS M81KR VIC. However, when hypervisor bypass technology with VMware VMDirectPath is used, the virtual machine accesses the Cisco UCS M81KR VIC directly, using the MAC address originally configured with Cisco UCS Manager.

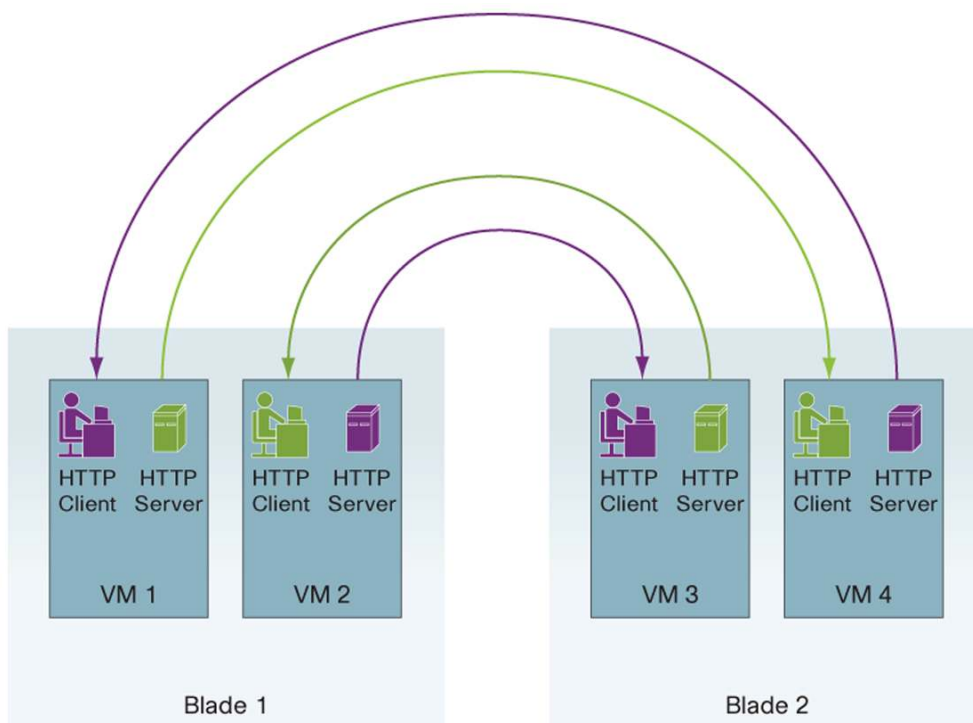
Cisco UCS Manager, together with Cisco VN-Link technology, is tightly integrated with VMware vCenter, enabling the virtual interface configuration and policy to move with the virtual machine when it migrates. This tight integration supports clear demarcation of responsibilities among network, server, and storage administrators – the network and storage administrators set up policies, and the server administrator can apply them for consistent automated deployment of virtual machines – facilitating collaboration and communication among the groups.

Benchmark Results

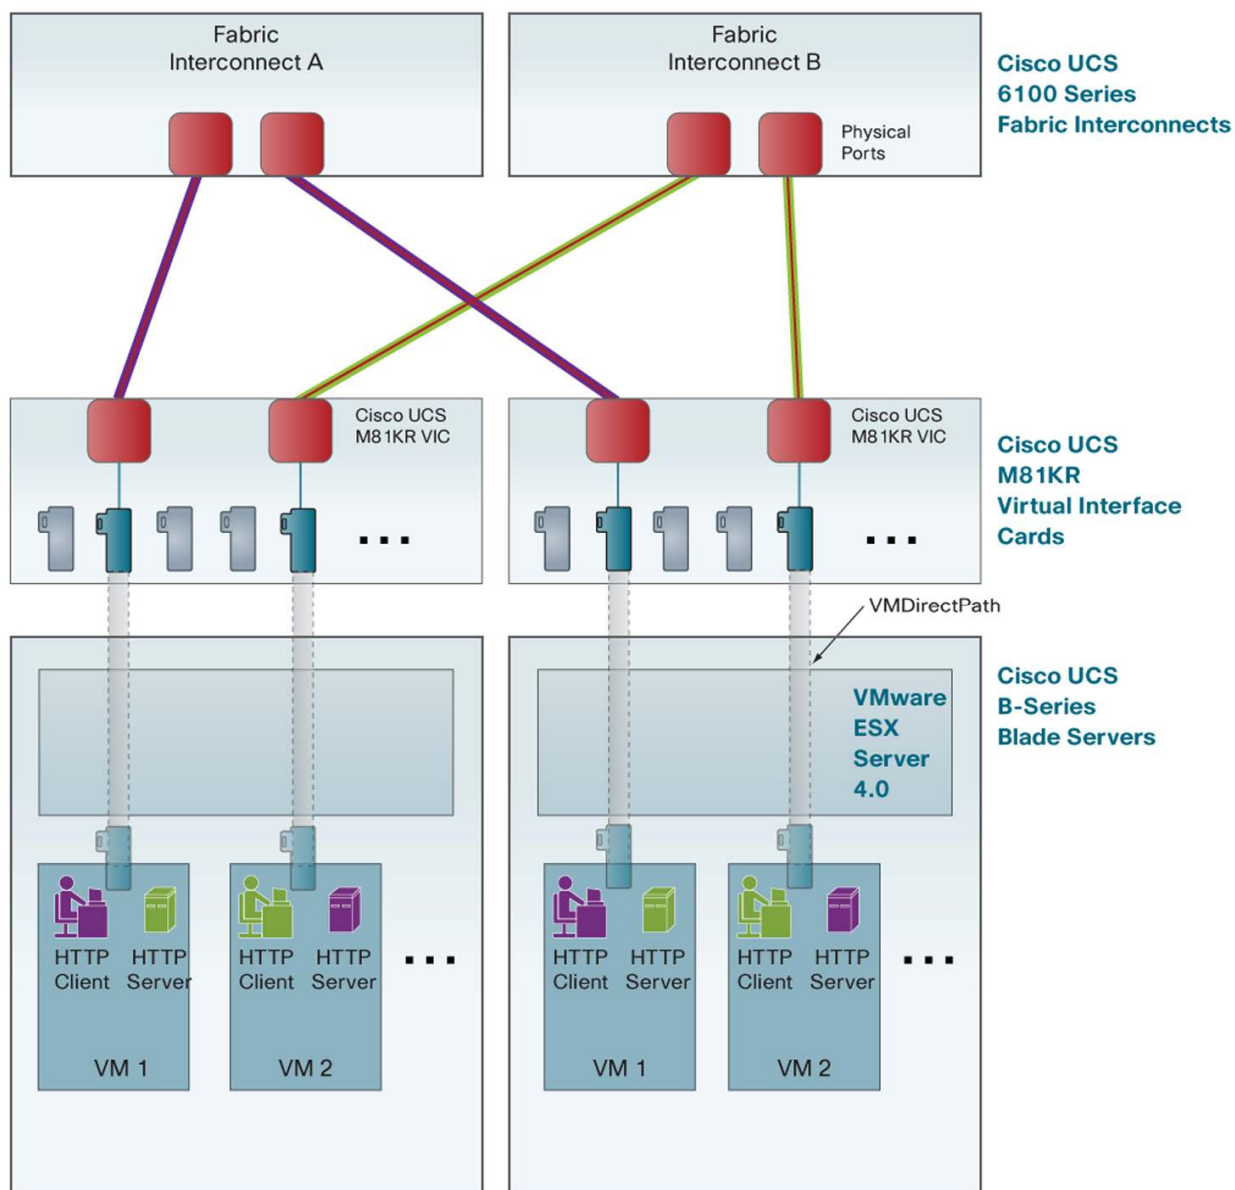
The results of preliminary testing by Cisco show that throughput performance for I/O-intensive tasks between two blades within the Cisco Unified Computing System platform is consistently better using the Cisco UCS M81KR VIC than when using a software switch solution. In the test, pairs of virtual machines were created on each of two blades. Each virtual machine ran one HTTP server application and one HTTP client application, and each HTTP server sent real-world web data to the corresponding client on the other blade. To evaluate performance, the network bandwidth and the runtimes were collected for each benchmark run. This test was run with 16, 24, 32, 40, and 48 pairs of virtual machines to determine how the system responded under both light and heavy I/O conditions.

Performance Test Environment

The test environment used two Cisco Unified Computing System blades each with Intel Xeon 5500 series processors with Intel Virtualization Technology for Directed I/O (VT-d) enabled. Each blade was configured with a single Cisco UCS M81KR VIC. The blades were housed in separate Cisco Unified Computing System chassis but within the same Cisco Unified Computing System domain. VMware vSphere 4.0 was loaded onto the blades, and all virtual machines ran the Red Hat Enterprise Linux U4 64-bit operating system. Every virtual machine ran both an HTTP Apache server and an HTTP client running Gnu Wget. The HTTP server on one blade delivered 128 MB of data (from a mounted RAM disk to reduce latencies and disk I/O rate variability) to a client running on the other blade (Figure 4).

Figure 4. Traffic Flow Topology

For these tests, Port 1 of the Cisco UCS M81KR VIC on Blade 1 communicates with Port 1 of the UCS M81KR VIC on Blade 2 (Figure 5). Likewise, Port 2 on Blade 1 communicates with Port 2 on Blade 2.

Figure 5. Cisco UCS M81KR VIC Port Communications Between Blade 1 and Blade 2

VMware vSphere 4.0 was configured on each blade. For each test run, x virtual machines were created on Blade 1 and y virtual machines were created on Blade 2, where $x=y$. One port group was configured for the Cisco UCS M81KR VICs in each blade, with one vNIC reserved for VMware VMconsole, one vNIC reserved for VMware VMkernel, and all other vNICs associated with virtual machines. No VMware VMotion migration was configured. Each virtual machine had Red Hat Enterprise Linux (RHEL) 5.4 loaded on it.

This configuration was then run with multiple pairs of virtual machines to test network throughput. The test examined throughput and application runtime for 16 pairs of virtual machines (16 virtual machines on Blade 1 communicating with 16 virtual machines on Blade 2), and then again for 24, 32, 40, and 48 pairs of virtual machines. Running 48 pairs of virtual machines, a total of 12 terabytes (TB) of bidirectional traffic was pushed through the virtual interface cards.

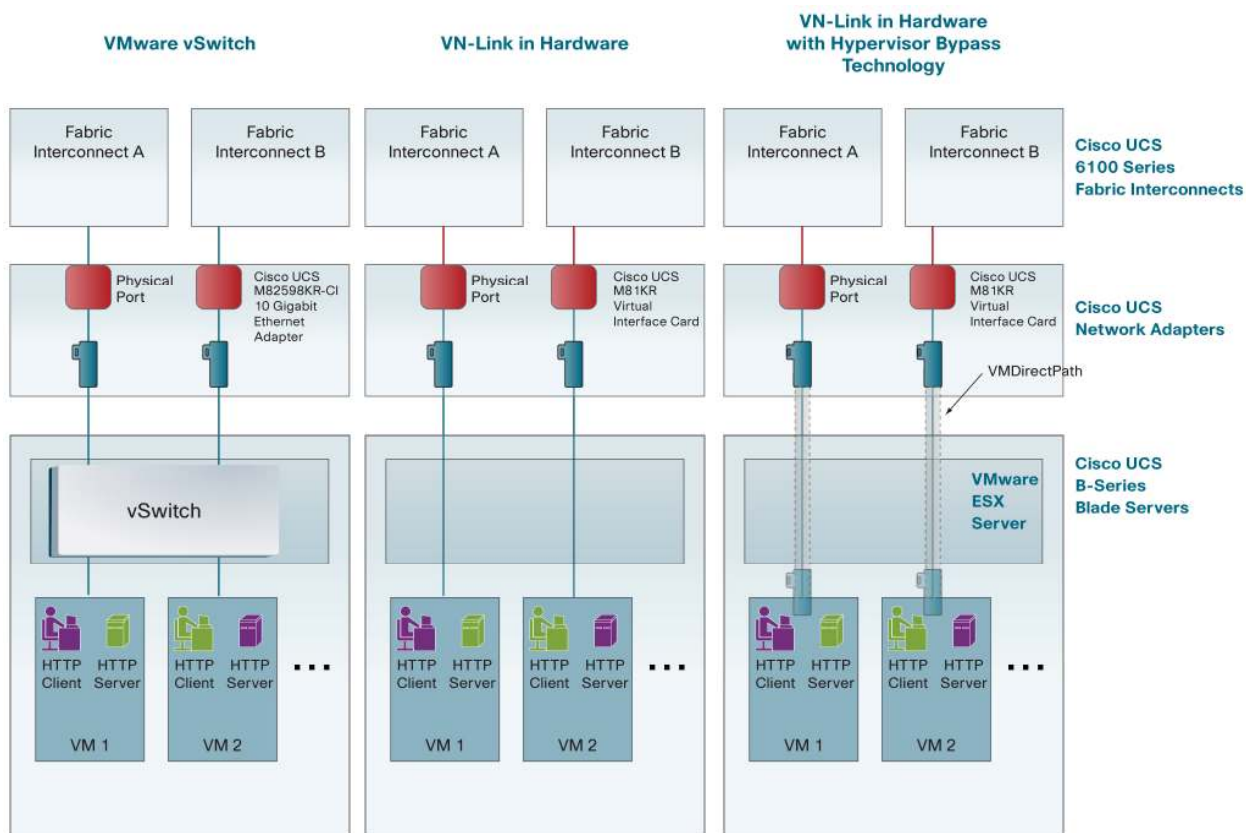
Test Scenarios

Three scenarios were tested using this environment. The first scenario used the VMware vSwitch software switch that resides in each blade's VMware vSphere 4.0 hypervisor. For these tests, the Cisco UCS 82598KR-CI 10 Gigabit Ethernet Adapter, based on the Intel 82598EB 10 Gigabit Ethernet Controller, was used. VMXNET 3 network drivers were used inside the virtual machines to interface with the hypervisor. A port profile was configured on the vSwitch, and each virtual machine was allocated a port from the profile. In this scenario, all network traffic switching takes place within the hypervisor on each blade (consuming CPU).

For the second scenario, the Cisco UCS M81KR VIC with Cisco VN-Link in hardware was used with traffic passing through the VMware hypervisor. Again, VMXNET 3 network drivers were used to interface with the hypervisor. The traffic was directed up to the parent switch, where the virtual machine network traffic switching is handled in hardware, just as if it were physical machine traffic.

The third scenario also used the Cisco UCS M81KR VIC with Cisco VN-Link in hardware; however, this scenario used VMware VMDirectPath, which bypasses the hypervisor for direct access from the operating system in each virtual machine to the virtual interfaces (vNICs) in the Cisco UCS M81KR VIC (Figure 6). For these test runs, the default enic driver from RHEL 5.4 was used to enable direct communication between the operating system and the virtual interface card, bypassing the hypervisor. Again, the traffic was directed to the parent switch, where the virtual machine network traffic switching is handled in hardware, just as if it were physical machine traffic.

Figure 6. Three Test Scenarios



Test Results

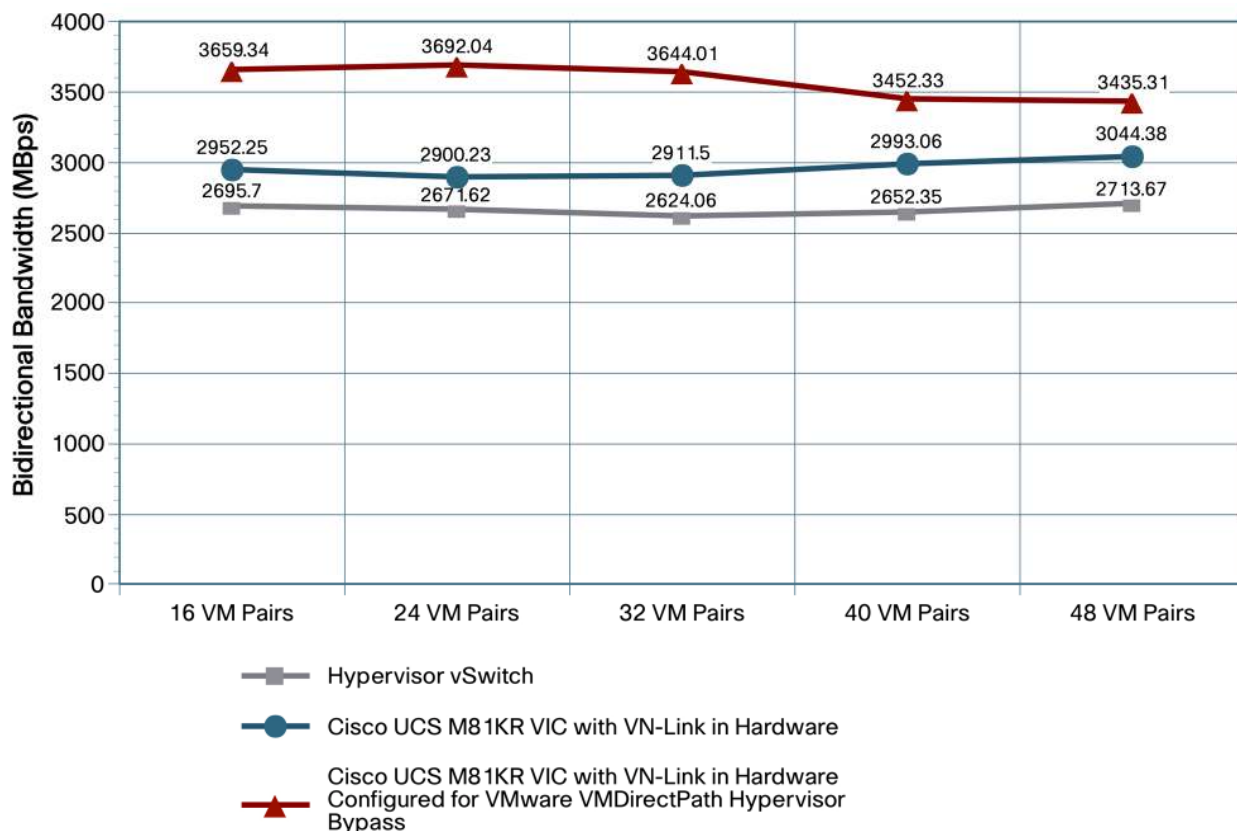
The results of testing the three scenarios showed that using the Cisco UCS M81KR VIC with VN-Link in hardware bypassing the hypervisor using VMware VMDirectPath gave the best throughput results by 38 percent compared to the vSwitch software option (Figure 7). Additionally, the combination of the Cisco UCS M81KR VIC and hypervisor bypass nearly saturated the network, to 92 percent of the total bandwidth available. This scenario achieved the greatest throughput when running 24 pairs and 32 pairs of virtual machines, with only 1.3 percent reduction in throughput with the higher numbers of virtual machine pairs.

With VMware VMDirectPath, after the system gets beyond 24 virtual machine pairs, the hypervisor is polled for CPU resources by the virtual machine. However, the vNIC works outside the control of the hypervisor, and all scheduling is performed directly by the virtual machine. With large numbers of virtual machines, the duplicate effort by the hypervisor and the virtual machine to schedule I/O can cause additional overhead.

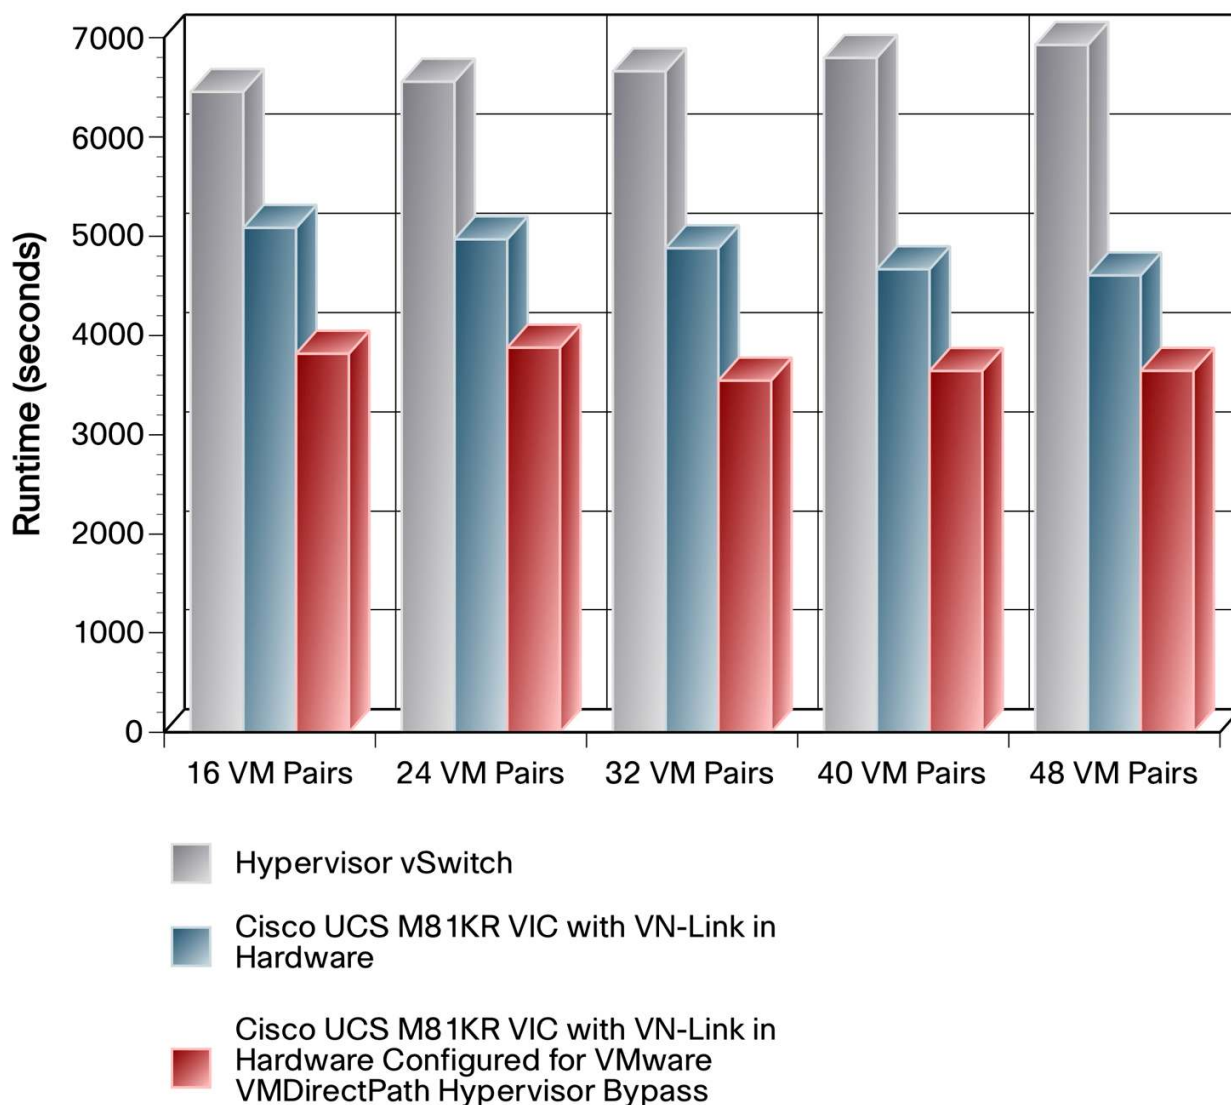
The scenario using the Cisco UCS M81KR VIC with VN-Link in hardware configured with the network links going through the hypervisor pushed network throughput to up to 76 percent of the total bandwidth available. This scenario resulted in up to a 9 percent improvement in network throughput compared to the vSwitch software scenario.

The scenario using the vSwitch software in the hypervisor configured with a 10 Gigabit Ethernet NIC achieved up to 67 percent use of the total bandwidth available. Additionally, as the number of virtual machine pairs increased, the throughput remained relatively constant (within 2 percentage points).

Figure 7. Test Results Comparison for the Three Scenarios Run



In addition to evaluating network throughput performance, the tests recorded elapsed time for each scenario to determine whether there were any differences in the speed with which the applications were able to complete the work. Figure 8 shows the elapsed time results.

Figure 8. Elapsed Time of Tests for Each of the Three Scenarios with Varying Numbers of Virtual Machines

In addition to handling a greater amount of network throughput, the Cisco UCS M81KR VIC using VMware VMDirectPath hypervisor bypass ran significantly faster than other test scenarios. In fact, the Cisco UCS M81KR VIC using hypervisor bypass ran nearly 50 percent faster than the hypervisor switch in some cases, enabling applications to perform up to twice the work in the same amount of time. Additionally, the scenario with the Cisco UCS M81KR VIC connected through the hypervisor ran up to 30 percent faster than the software hypervisor switch solution.

What This Means for Your Environment

The results of these tests show that the Cisco UCS M81KR VIC easily handles I/O-intensive workloads, similar to enterprise application workloads. The Cisco UCS M81KR VIC excels in performance with various virtual machine densities while maintaining evenly distributed bandwidth availability to all virtual machines under heavy network load.

Additionally, the absence of a software-based switch increases the CPU cycles available to the virtual machines and applications. This increase enables more business-critical work to be accomplished in a shorter amount of time, increasing competitive advantage.

When configuring any virtualized system, one common best practice is to balance the CPU cycles across the workloads. This best practice also applies to network I/O access. The best throughput will be obtained when bidirectional I/O access is balanced across the virtual machine workloads and network infrastructure.

Conclusion

Consolidating servers using virtualization makes good business sense. The Cisco Unified Computing System with the Cisco UCS M81KR VIC allows greater server consolidation densities and increases throughput performance by 38 percent over software-based switch solutions. The hardware-based solution offloads performance-intensive I/O tasks from the CPU, enabling greater (nearly 50 percent) application performance and greater virtual machine densities while supporting network throughput to handle increased densities and application load. With the Cisco Unified Computing System with the Cisco UCS M81KR VIC, the network is no longer the bottleneck to virtual server performance.

A single Cisco UCS M81KR VIC performs the functions of both NICs and HBAs, lowering costs and simplifying the entire networking infrastructure. Integrated role- and policy-based management enables network and storage administrators to manage virtual interfaces the same way they manage physical interfaces, reducing the number of management systems and building on current personnel skills for an overall lower TCO.

For More Information

For more information about the Cisco UCS M81KR Virtual Interface Card, visit http://www.cisco.com/en/US/prod/ps10265/ps10280/cna_models_comparison.html or contact your local account representative.



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.

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 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)