

A Platform Built for Server Virtualization: Cisco Unified Computing System

What You Will Learn

This document discusses how the core features of the Cisco Unified Computing System contribute to the ease of deployment, management, performance, and security of virtualized environments. It describes the system's unified fabric, unified and embedded management, service profiles, and virtualization density and the way in which Cisco® VN-Link technology automates and simplifies virtual machine networking and migration.

Introduction

In good economic times, “do more with less” is a motto that helps business organizations sharpen their competitive edge. During challenging economic times, accomplishing more with fewer resources can help organizations with their very survival.

Nearly everyone knows the benefits of virtualization:

- Consolidate workloads; raise utilization levels; and reduce operating, capital, space, power, and cooling expenses.
- Move workloads dynamically within a virtualization pool to increase the flexibility to take servers offline or bring new ones online.
- Manage the relationship of virtual machines to physical machines to optimize performance and maintain service levels.
- Scale existing applications or deploy new ones by creating more virtual machines using an existing pool of resources.
- Use the high-availability and disaster-recovery features of virtualization software to overcome localized and geographic failures.

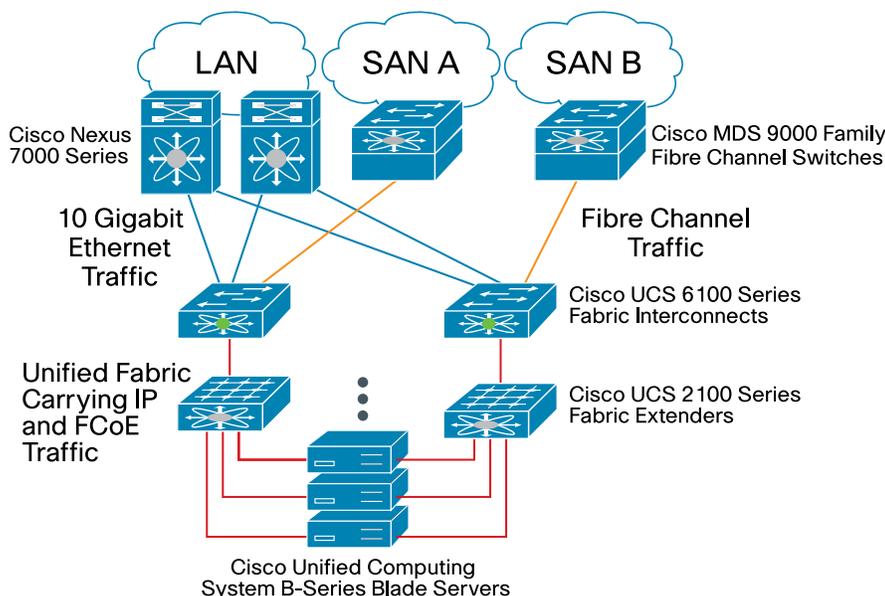
Virtualization decouples application deployment from server purchases, but this and the other benefits of virtualization are best achieved when applications run on one or more uniform pools of server resources. The Cisco Unified Computing System™ is designed to provide just such an environment. Optimized for virtualized environments, the Cisco Unified Computing System is a next-generation data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain.

Unified Fabric

Virtualized environments need consistent I/O configurations that provide uniform support for hypervisors across all the servers that are part of a resource pool. They also need I/O configurations that support the movement of virtual machines (VMs) across servers in a resource pool while maintaining individual VM bandwidth and security requirements. The Cisco Unified Computing System delivers on this need by basing the system on a low-latency, lossless, 10-Gbps unified network fabric. Blade servers in the Cisco UCS 5108 Blade Server Chassis have access to the fabric through mezzanine adapters that provide up to 40 Gbps of throughput per blade server.

The unified fabric enables a “wire once” deployment model in which chassis are cabled to the fabric interconnects just one time, and I/O configurations changes are performed through the management system, not by installing host adapters and recabling racks and switches (Figure 1). The unified fabric dramatically simplifies rack cabling by eliminating the need for multiple redundant Ethernet and Fibre Channel adapters in each server, separate cabling to access-layer switches, and separate switches for each network medium. Instead, all traffic is routed to the central server interconnects, where Ethernet and Fibre Channel then can be separated onto native, nonconsolidated networks.

Figure 1. A Unified Fabric Carries Multiple Traffic Streams to Cisco UCS 6100 Series Fabric Interconnects, Where Ethernet and Fibre Channel Traffic Splits onto Separate Networks



The unified fabric is based on 10 Gigabit Ethernet, with standards-based extensions to support more types of traffic with improved management. The unified fabric supports Ethernet as well as Fibre Channel over Ethernet (FCoE), with management features that allow multiple lanes of traffic, such as Ethernet and FCoE, to be managed independently with bandwidth management and no interference between traffic classes.

What the unified fabric offers virtualized environments is the capability to create large pools of server resources with uniform I/O connectivity that can be programmed to operate in a way consistent with a data center’s current best practices. For environments in which Fibre Channel–based shared storage is used by virtualization software, it eliminates a redundant pair of host-bus adapters (HBAs), transceivers, cables, and upstream switch ports whose costs can approach that of a small server itself. As discussed later, Cisco VN-Link technology supports virtual network links between each virtual machine and the fabric interconnects, easing the management of virtual machines and their networks, including facilitating virtual machine movement while automatically maintaining network characteristics including security.

Embedded, Unified Management

Virtualization helps data centers the most when servers are configured into resource pools that can be harnessed on demand to meet fluctuating workload requirements; deploy new applications without tying them to specific hardware; and move virtual machines between servers to balance workloads, meet service-level agreements (SLAs), or prepare a server for scheduled downtime. Cisco UCS Manager transforms the resources of the Cisco Unified Computing System into a single cohesive system that is ideal for establishing resource pools for virtualized environments.

Cisco UCS Manager acts as the central nervous system of the Cisco Unified Computing System. It integrates the system components from end to end so the system can be managed as a single logical entity. Cisco UCS Manager provides an intuitive GUI, a command-line interface (CLI), and a robust API so that it can be used alone or integrated with other third-party tools. Everything about server configuration—system identity, firmware revisions, network interface card (NIC) settings, HBA settings, and network profiles—can be managed from a single console, eliminating the need for separate element managers for every system component. Cisco UCS Manager is embedded in the system's two fabric interconnects as a highly available clustered pair.

Centralized, comprehensive management combined with the unified fabric means that there is no longer a need to manually configure and integrate separate components to create an effective virtualization pool. Cisco UCS Manager automates and simplifies the procedure of incorporating a new server into the system to the point where a new server can be installed, configured, and put to use in a matter of minutes, rather than the hours or days that traditional configuration approaches require. This capability does more than just help increase IT staff productivity. The capability to quickly scale a virtualization pool in need of more resources can give an enterprise using such a system a strategic advantage.

The role- and policy-based management model of Cisco UCS Manager allows organizations to maintain the current separation of disciplines they have already established in their organizations. For example, the default configuration defines separate server, storage, and network administrator roles in which users in each role define policies that dictate how systems are to be configured. The use of policies allows different organizational teams to work more efficiently together. For example, network administrators define all networking policies, which can then be later incorporated and referenced by server administrators without involving their network administrator colleagues. Meanwhile, high-level administrators are freed to focus more on business-critical issues rather than on the details of individual system configuration.

Service Profiles and Service Profile Templates

Virtualization decouples application deployment from server purchases. When servers are configured into virtualization pools, a data center becomes a dynamic entity in which resources are used efficiently, and the allocation of virtual machines to physical servers can be adjusted dynamically to best balance efficiency and performance. A data center acts like a fine-tuned engine, and adding servers is like fueling that engine: as the overall workload on a virtualization pool increases, more servers are needed to maintain required service levels. Likewise, when workload requirements diminish, it should be straightforward to power off servers and thus closely match power consumption and carbon footprint with actual workload requirements.

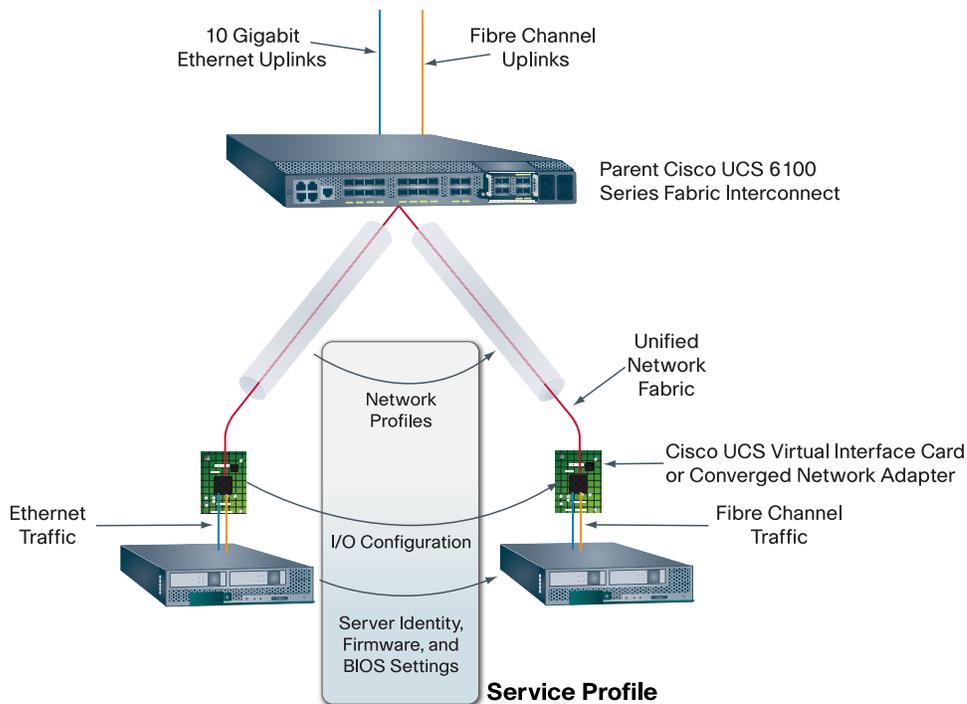
Service profiles and service profile templates are the mechanisms that dramatically simplify the process of incorporating new servers into the system so that they can be added to a resource pool in minutes, rather than in hours or days. Likewise, the same mechanisms can be used to reconfigure servers so that they can be used in a different resource pool than the one in which they were initially members.

Cisco UCS Manager configures and provisions every aspect of a server and its connectivity within the system. It can set or provision unique user ID (UUID), BIOS, and firmware revisions; NIC configurations including MAC address, VLAN, and quality-of-service (QoS) settings; HBA configurations including worldwide names (WWNs), VSANs, bandwidth constraints, and firmware revisions; and uplink port configurations including VLAN, VSAN, QoS, and EtherChannel settings.

Service profiles fully specify a server and all its settings. A service profile can be used to provision a new server or, if necessary, provision a replica of the same server so that an entire software stack can be moved from the original server to a replacement one, all completely transparently to the software and its licensing mechanisms (Figure 2). Whereas a service profile fully describes a server, a service profile template describes how to create a service

profile. Whereas service profiles are like instances of a class, service profile templates define the class; they define the policies for creating a service profile.

Figure 2. Service Profiles Provision Server and Network Resources and Can Move Server Configurations Between Physical Servers



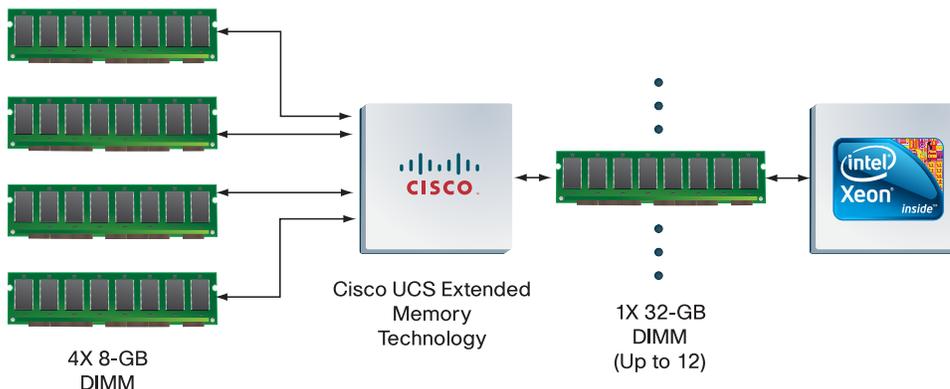
Service profile templates can be invoked to provision servers so that they are all configured exactly the same, yet have their own identities (such as UUIDs, MAC addresses, and WWNs) in instances in which identifiers need to be unique. The result of applying a service profile template is a service profile that fully defines one unique server. Cisco UCS Manager can be configured to discover new servers as they are added to a Cisco Unified Computing System and automatically apply an appropriate service profile template depending on the physical characteristics of the server. This capability makes adding servers to different virtualization pools quick and easy. For example, one pool might consist of servers with 64 GB of main memory, sufficient for one class of applications, while another pool might consist of servers with a minimum of 192 GB of main memory for maximum virtual machine density.

Virtualization Density with Cisco UCS Extended Memory Technology

Virtualization puts the spotlight on the need for servers to have large, cost-effective memory footprints. While high-performance, virtualization-optimized, multicore processors such as the Intel Xeon 5500 series have made it possible to improve virtual machine performance, servers now need an even larger amount of memory to fully utilize a server's processors. The traditional approach to increasing virtual machine density on a virtualized system is to purchase a larger and more expensive four-socket server. But this approach increases both capital and operating costs without actually solving the problem of how to bring more cost-effective memory into a two-socket server.

Cisco Extended Memory Technology offers a potent, cost-effective approach that can help increase virtualization density with lower TCO than other approaches, enabling IT organizations to accomplish more with fewer resources. The technology is available on the Cisco UCS B250 M1 Extended Memory Blade Server and the Cisco UCS C250 M1 Extended Memory Rack-Mount Server. The technology maps four physically distinct DIMMs to a single logical DIMM as seen on the processor's memory channel (Figure 3). This mapping supports extended memory servers with 48 DIMM slots in which traditional servers and blade systems using the same processor can have only up to 12 slots at full performance, or 18 slots at reduced performance.

Figure 3. Cisco Extended Memory Technology Makes Four Physical DIMMs Appear to the CPU as a Single, Large, Logical DIMM



Up to 48 DIMM slots in a Cisco extended memory server can be populated with 4- or 8-GB DIMMs to give organizations exceptional flexibility in balancing processing power, memory capacity, and cost.

- A low-cost option delivers a memory footprint of up to 192 GB using low-cost 4-GB DIMMs rather than the 8-GB DIMMs required for large memory footprints in other servers using the same processors. This option saves 60 percent on memory costs based on publicly available memory pricing obtained in August 2009.
- A large-footprint option can accommodate the most memory-hungry of workloads. With up to 384 GB of memory using 8-GB DIMMs, Cisco extended memory servers offer the largest memory footprint available from any two-socket server using Intel Xeon 5500 series processors.

Whether a large or an extremely large memory footprint is needed to optimize virtualization density, IT organizations can now consolidate more applications and create more virtual machines with better economy.

Networking with Virtual Interfaces and Cisco VN-Link Technology

Traditional blade server implementations add unnecessary cost, complexity, and risk to virtualized environments. In most network implementations, the network access layer has been fragmented into three tiers, making it difficult to maintain control and security over network connectivity, adding unnecessary latency to VM-to-VM networking, and creating obstacles to effective management:

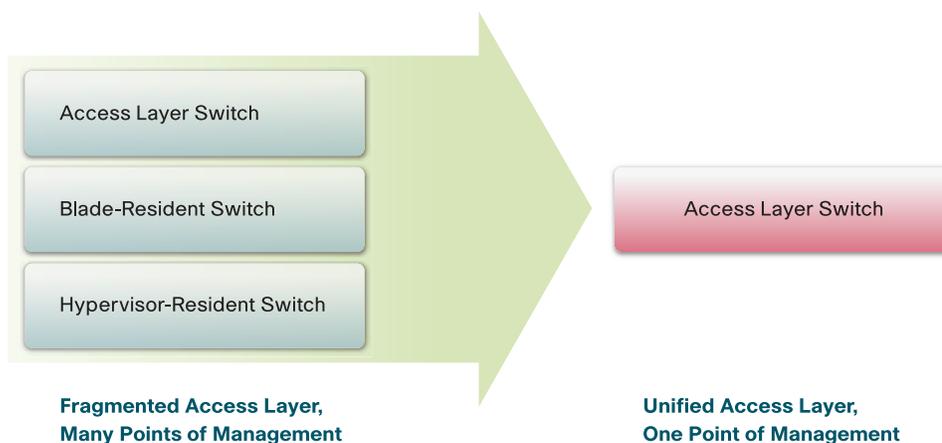
- Access-layer switches are typically part of the data center infrastructure that is managed by network administrators with highly effective control over security and QoS.
- Switches residing in blade servers add a new layer of networking that often uses commodity silicon with feature sets inconsistent with the data center access-layer switches.
- Software switches implemented by virtualization software vendors consume CPU cycles to emulate network hardware at the expense of application performance. These switches are often completely beyond the purview of network administrators and are most often configured by server administrators.

The access-layer fragmentation in these environments results in an arrangement in which different switching devices are used for VM-to-VM communication depending on whether the virtual machines are communicating within the same blade server, within the same chassis, or between chassis. Consistent management between layers is difficult to accomplish, especially when in the face of virtual machines dynamically migrate between servers.

The Cisco Unified Computing System simplifies, speeds up, and secures switching by reducing the switches involved in VM-to-VM communication to a single fabric interconnect. This approach establishes a single point of control and management for all network communication within the system (Figure 4):

- A single access-layer switch, the Cisco UCS 6100 Series Fabric Interconnect, supports all network traffic between virtual machines, regardless of where they reside, establishing a single point of control and management for network traffic.
- Blade-resident switches are eliminated in favor of Cisco UCS 2100 Series Fabric Extenders, devices that are logically part of the fabric interconnects and that pass all traffic from blade servers to the upstream fabric interconnect.
- Hypervisor-resident software switches are eliminated and virtual machines connect directly to physical NICs through pass-through switching or hypervisor bypass technology that is further accelerated by the Intel Virtualization Technology for Directed I/O provided by Intel Xeon 5500 series processors.

Figure 4. Cisco Unified Computing System Establishes a Single Point of Control and Management, Eliminating the Fragmentation of the Network Access Layer



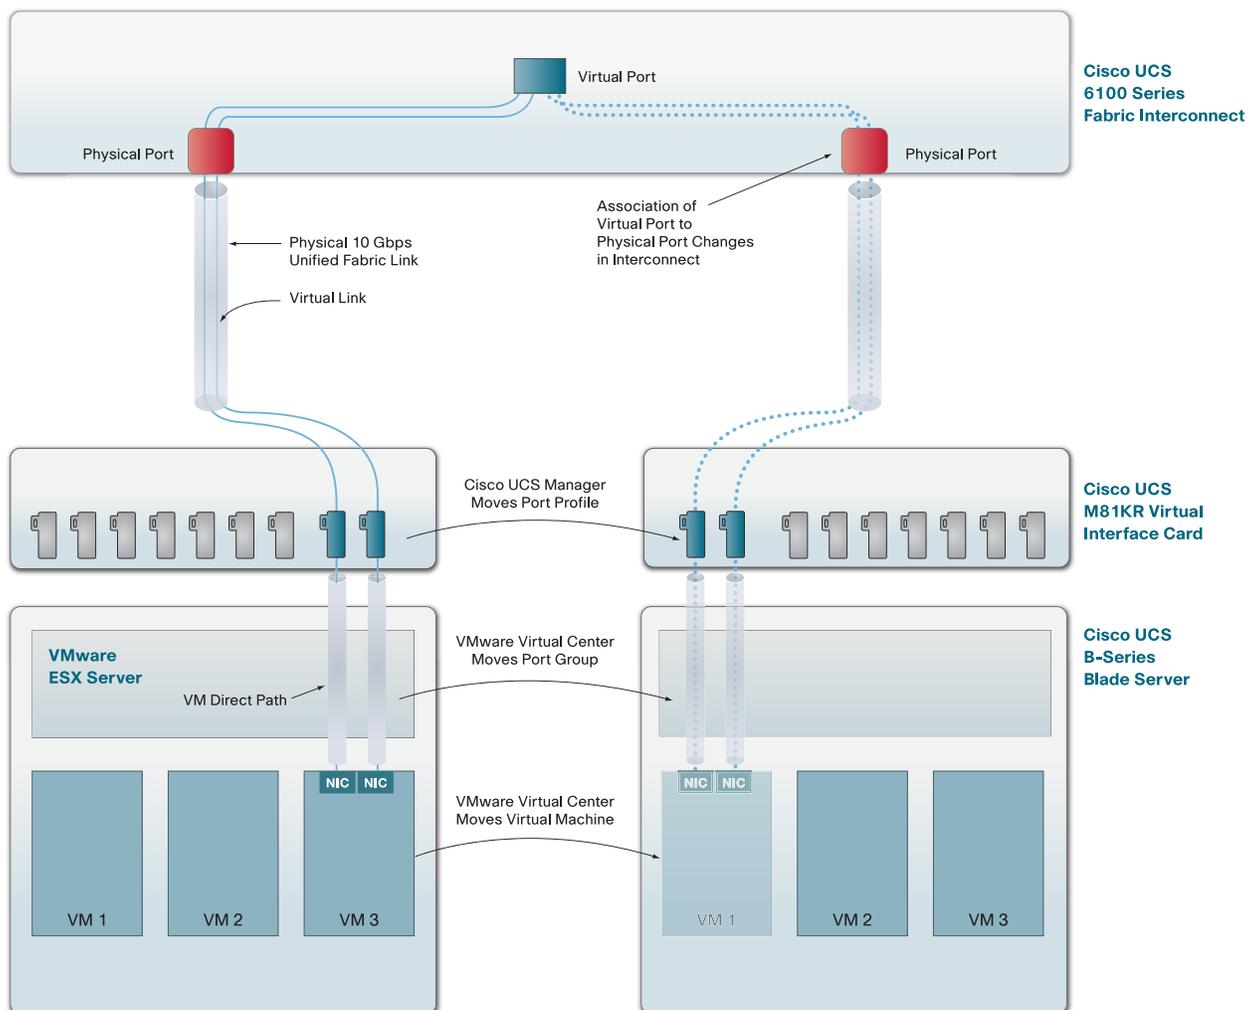
Cisco VN-Link Technology

Each server in a Cisco Unified Computing System is physically connected to the fabric interconnects by one or more physical links. Cisco VN-Link technology allows multiple virtual links to be configured on a single physical link. Virtual links connect a virtual NIC (vNIC) in a virtual machine to a virtual interface within a fabric interconnect. When used with a Cisco UCS M81KR Virtual Interface Card, a virtual link originates with a physical interface connected (through pass-through switching or hypervisor bypass) to a virtual machine (Figure 5). This arrangement allows network connectivity for virtual machines to be managed just like physical links are managed for physical servers. With all traffic to and from virtual machines carried on individual virtual links, attributes including QoS, VLANs, and access control lists (ACLs) can be managed consistently from a single point of management.

Virtual links are terminated at virtual interfaces inside the fabric interconnects. A virtual interface is associated with a physical interface, and that association can change as needed. When a virtual machine moves from one server to another, the virtual interface to which the VM's virtual link is connected is simply associated with a different physical port. Now virtual machines can move from server to server with their network characteristics moved with them, without difficult coordination between multiple layers of switching.

Between the virtual interfaces in the fabric interconnects and the physical interfaces supported by the virtual interface card, the Cisco Unified Computing System provides a hardware-based implementation of Cisco VN-Link technology. When interfaces other than the Cisco UCS M81KR are used to carry virtual machine network traffic, the same management simplicity can be established in software using the Cisco Nexus™ 1000V Series Switches.

Figure 5. Cisco VN-Link Technology Enables Per-Virtual-Machine Links That Are Managed and Moved Independently of Physical Links



Cisco VN-Link Technology with the Cisco UCS M81KR Virtual Interface Card

The full power of the Cisco Unified Computing System is achieved when servers are configured with Cisco UCS M81KR Virtual Interface Cards. These cards have a dynamically configured I/O configuration that allows Cisco UCS Manager to create any combination of up to 128 (8 are reserved for system use) Ethernet NICs or Fibre Channel HBAs with identities (MAC addresses and WWNs) programmed dynamically. This virtual interface card eliminates the need for hypervisor-resident switching by providing enough interfaces so that each virtual machine can have one or more dedicated physical interfaces.

Cisco UCS Manager works in conjunction with VMware vCenter software to coordinate the creation and movement of virtual machines along with the interfaces to which they directly communicate. A port profile in Cisco UCS Manager defines the NIC configuration used by a virtual machine and directs the manager to configure the interfaces needed by a virtual machine dynamically, just as the VM is created or moved to a server. The name of a port profile corresponds to the name of a port group in VMware ESX Server. When VMware ESX Server wants to create a new virtual machine or set up the destination for a virtual machine migration, it communicates the port group name to the virtual interface card. The virtual interface card asks Cisco UCS Manager for the port profile having the same name, and the virtual machine is able to connect to the network using devices that it expects to access.

VM Direct Path Technology with the Cisco UCS M81KR Virtual Interface Card

The VM Direct Path technology further optimizes virtual machine I/O by bypassing the hypervisor altogether on all I/O traffic, thereby eliminating one of the remaining bottlenecks for the virtualization of I/O-intensive workloads. Cisco's support for VM Direct Path with the Cisco UCS M81KR is built on the same VN-Link foundation for configuration and management.

End-Host Mode

End-host mode for Ethernet and Fibre Channel connections from the fabric interconnects to aggregation-layer switches completes the Cisco Unified Computing System as a single cohesive system. (End-host mode for Fibre Channel is based on N-Port ID Virtualization (NPIV)) These features cause the fabric interconnects to assume the identities of the physical and virtual machines in the system, making the existence of multiple servers transparent to the aggregation layer, further simplifying management of a data center's Layer 2 network.

Conclusion

Server virtualization offers numerous benefits to data centers everywhere, but it also presents challenges. Resource pools are more easily managed if they have consistent I/O configurations and sufficient bandwidth to support a significant number of VMs per physical server. They can more quickly adapt to rapidly changing business and workload conditions when new server resources can be added to pools in minutes, without requiring hours or days of tedious, time-consuming, and error-prone manual configuration to prepare a server and its interfaces, firmware, and settings for production use. The economic benefits of server virtualization are more easily achieved when virtual machine density can be increased through a balance of CPU power and cost-effective memory footprints. Finally, virtual environments can be managed more efficiently, more securely, and with better QoS when the network access layer is unified, and when links to virtual machines are treated the same way as links to physical servers.

The Cisco Unified Computing System meets these challenges with a next-generation data center platform that unites compute, network, storage access, and virtualization support in a cohesive system that is managed centrally and coordinated with virtualization software such as VMware ESX Server. The system integrates enterprise-class servers in a 10 Gigabit Ethernet unified network fabric that provides the I/O bandwidth and functions that virtual machines and the virtualization software that supports them both need. Cisco Extended Memory Technology offers a highly economical approach for establishing the large memory footprints that high virtualization density requires. Finally, the Cisco Unified Computing System integrates the network access layer into a single, easily managed entity in which links to virtual machines can be configured, managed, and moved as readily as physical links. The Cisco Unified Computing System continues Cisco's long history of innovation and delivers innovation in architecture, technology, partnerships, and services.

For More Information

For more information, please visit <http://www.cisco.com/go/unifiedcomputing>.



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.

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks. Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco.Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0910R)