

White Paper

# Delivering on the Promise of Virtualization: Cisco Unified Computing System



In Collaboration With



## Contents

<b>What You Will Learn.....</b>	<b>3</b>
<b>Introducing the Cisco Unified Computing System .....</b>	<b>5</b>
<b>Unified Fabric .....</b>	<b>7</b>
Problem: Multiple Parallel Networks .....	7
Solution: Cisco Unified Fabric .....	8
Problem: Multiple Network Layers.....	8
Solution: Cisco Fabric Extender Technology .....	10
Dramatic Reduction in TCO .....	11
Problem: Traditional Environments Are Difficult to Scale .....	11
Solution: Low-Cost Scaling with Cisco UCS .....	12
<b>Unified Management .....</b>	<b>14</b>
Problem: Server Deployment Challenges Slow Scaling .....	14
Solution: Automated Configuration with Cisco UCS.....	14
Problem: Managing Growth.....	15
Solution: Manage Multiple Cisco UCS Domains with Cisco UCS Central Software.....	16
Problem: Managing and Securing Virtual Networks .....	17
Solution: Cisco Data Center Virtual Machine Fabric Extender .....	17
Problem: Managing Spares and Licenses .....	18
Solution: Fewer Spares Through Dynamic Configuration .....	19
<b>Unified Computing.....</b>	<b>20</b>
Problem: Virtualized Environments Underperform.....	20
Solution: Cisco UCS Is Optimized for Virtualization .....	20
Optimized Performance for Virtualized SAP Environments.....	21
Increased Performance for Virtualized Microsoft SQL Server .....	21
Problem: Memory Constraints.....	21
Solution: Cisco Leads in Memory Capacity.....	21
Problem: Insufficient I/O Bandwidth .....	22
Solution: The Right Kind of Bandwidth .....	22
<b>Easy Deployment .....</b>	<b>24</b>
Cisco Validated Designs .....	24
Cisco SmartPlay Solutions .....	24
Virtual Computing Environment Coalition Vblock Systems .....	24
Health Benefit Administrator Cures Growing Pains with VCE Vblock System-Based Private Cloud .....	24
NetApp FlexPod.....	24
FlexPod Transforms Data Center .....	25
Cisco Services.....	25
<b>Conclusion .....</b>	<b>26</b>



# Delivering on the Promise of Virtualization: Cisco Unified Computing System

White Paper  
March 2013



## Massive Consolidation

NetApp used Cisco UCS to virtualize its testing lab. Using Cisco UCS, it consolidated 51 traditional blade server chassis with 178 servers into a system with 15 blade server chassis and 120 servers. "It took just an hour to deploy the first 112-server Cisco UCS with NetApp storage and VMware vSphere," said Brandon Agee, technical engineering support systems lead.

Read the case study at [http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/case\\_study\\_NetApp\\_Cisco\\_Kilo\\_Lab.pdf](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/case_study_NetApp_Cisco_Kilo_Lab.pdf).

## What You Will Learn

Virtualization has transformed the data center over the past decade. IT departments use virtualization to consolidate multiple server workloads onto a smaller number of more powerful servers. They use virtualization to scale existing applications by adding more virtual machines to support them, and they deploy new applications without having to purchase additional servers to do so. They achieve greater resource utilization by balancing workloads across a large pool of servers in real time—and they respond more quickly to changes in workload or server availability by moving virtual machines between physical servers. Virtualized environments support private clouds on which application engineers can now provision their own virtual servers and networks in environments that expand and contract on demand.

While virtualization has been successful in propelling a major transformation, it also has created many problems, including:

- A proliferation of interfaces, cables, and upstream switch ports to support each server, adding cost and complexity
- Multiple layers of hardware and software switching that makes management difficult
- Too many management points, making it difficult to manage quality of service (QoS) and maintain security
- Scalability made difficult by the amount of time needed to configure servers and integrate them into the network infrastructure
- Performance hampered by the overhead of virtualized environments and constraints on resources

The Cisco Unified Computing System™ (Cisco UCS®) delivers on the promise of virtualization with a single unified system that integrates a unified fabric, embedded management, and powerful servers with intelligent Intel® Xeon® processors. Cisco UCS delivers world-record-setting virtualization performance through a better architecture and better balance of resources, delivering on the promise of virtualization through a system that:

- Simplifies physical and virtual networks, reducing cost while increasing manageability
- Scales better and more rapidly and with lower infrastructure cost per server
- Delivers greater performance for virtualized environments through a better balance of resources
- Increases an organization's responsiveness to changing workloads and business conditions through increased flexibility

This document describes Cisco UCS by looking at how its three main components—unified fabric, unified management, and powerful servers with intelligent Intel Xeon processors—solve the problems that virtualization creates.

## Introducing the Cisco Unified Computing System

Cisco UCS (Figure 1) is the first truly unified data center platform that combines industry-standard, Intel Xeon processor-based servers with unified management, networking, and storage access into a system built for virtualized environments, with optimizations specifically designed for Microsoft Windows 2008 R2 Hyper-V, Red Hat Enterprise Linux, and VMware vSphere. The system is a smart infrastructure that is automatically configured through integrated, model-based management to simplify and accelerate deployment of enterprise-class applications and services running in bare-metal, virtualized, and cloud-computing environments.

The system's x86-architecture rack and blade servers are powered by Intel Xeon processors. These industry-standard servers deliver world-record performance to

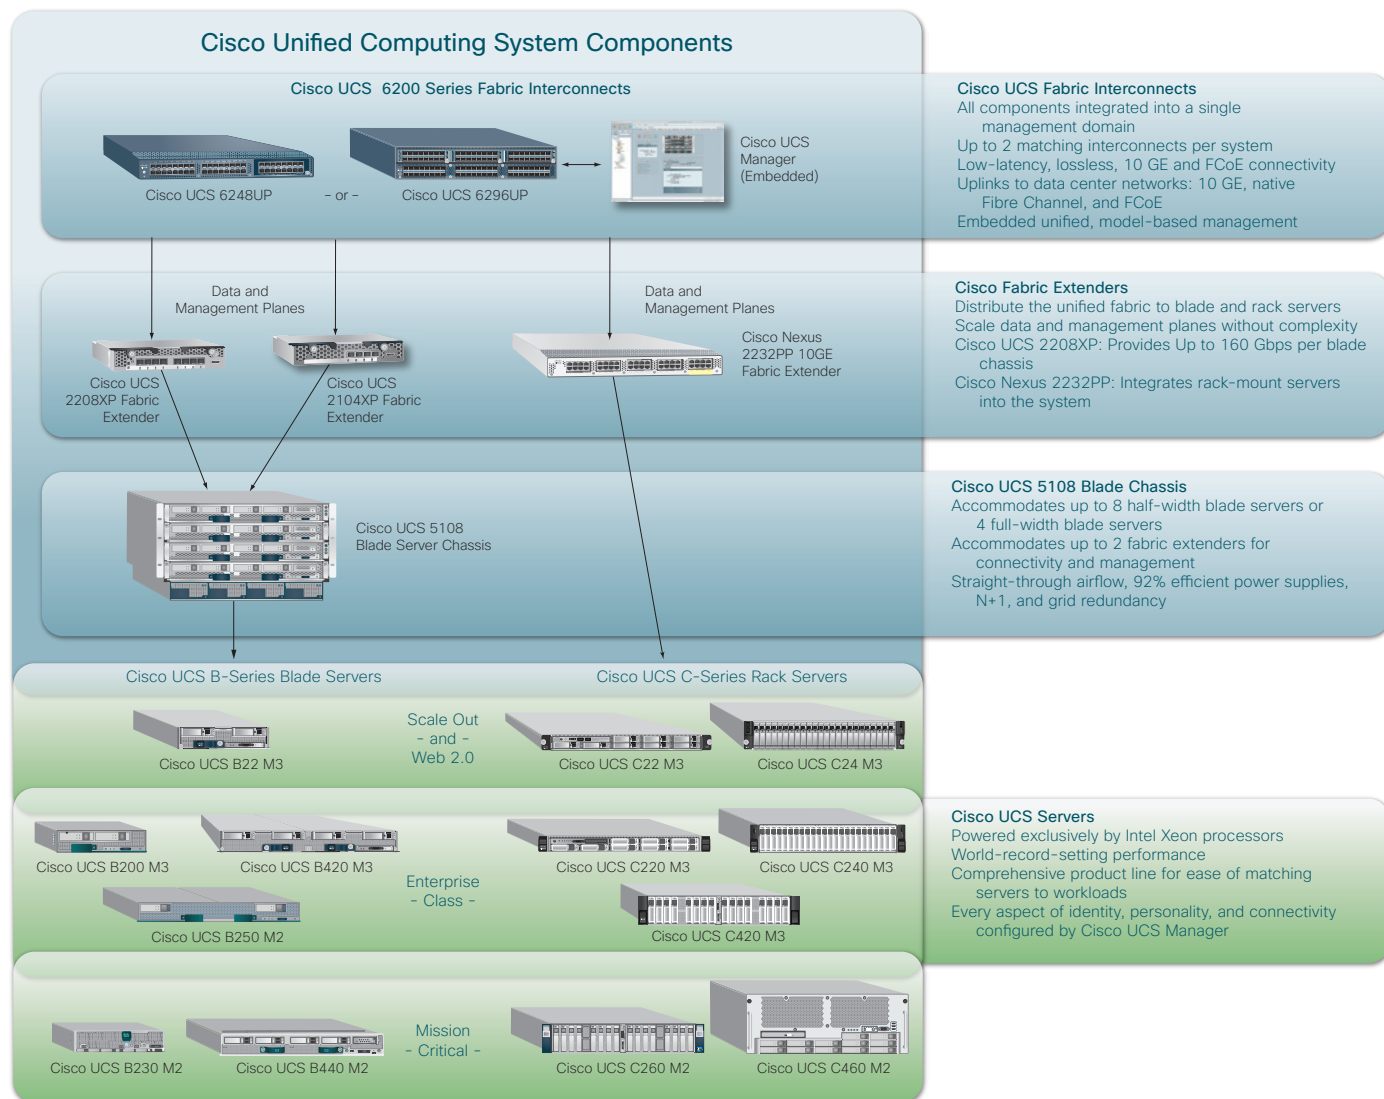


Figure 1. Cisco UCS Component Hierarchy

power virtualized environments, with enhancements at the chip level to accelerate virtual I/O. Cisco® servers, in combination with a simplified, unified architecture, help increase IT productivity and provide a superior price-to-performance ratio for lower total cost of ownership (TCO). Only Cisco servers integrate with Cisco UCS, and only Cisco integrates rack and blade servers into a single unified system.

Building on Cisco's strength in enterprise networking, Cisco UCS is integrated with a standards-based, high-bandwidth, low-latency, virtualization-aware unified fabric. The system is wired once to support the desired bandwidth and carries all Internet protocol, storage, interprocess communication, and virtual machine traffic with security isolation, visibility, and control equivalent to that for physical networks. The system's 10 Gigabit Ethernet network meets the bandwidth demands of today's multicore processors, eliminates costly redundancy, and increases workload agility, reliability, and performance.

The Cisco Fabric Extender Technology (FEX Technology) reduces the number of system components that need to be purchased, configured, managed, and maintained by condensing three network layers into one. It eliminates both blade server and hypervisor-based switches by connecting fabric interconnect ports directly to individual blade servers and virtual machines. Virtual networks are now managed exactly like physical networks, but have massive scalability. This approach represents a radical simplification compared to traditional systems, reducing capital expenditures (CapEx) and operating expenses (OpEx) while increasing business agility, simplifying and accelerating deployment, and improving performance.

Cisco UCS helps organizations go beyond efficiency; it helps them become more effective through technologies that foster simplicity rather than complexity. The result is flexible, agile, high-performance, self-integrating information technology that reduces staff costs, increases uptime through automation, and offers faster return on investment (ROI).

## Unified Fabric

Cisco UCS uses a unified fabric that dramatically simplifies the architecture needed to support virtualized environments, resulting in lower cost and complexity and simpler, easier management.

### Cabling for a Traditional Blade Server Environment



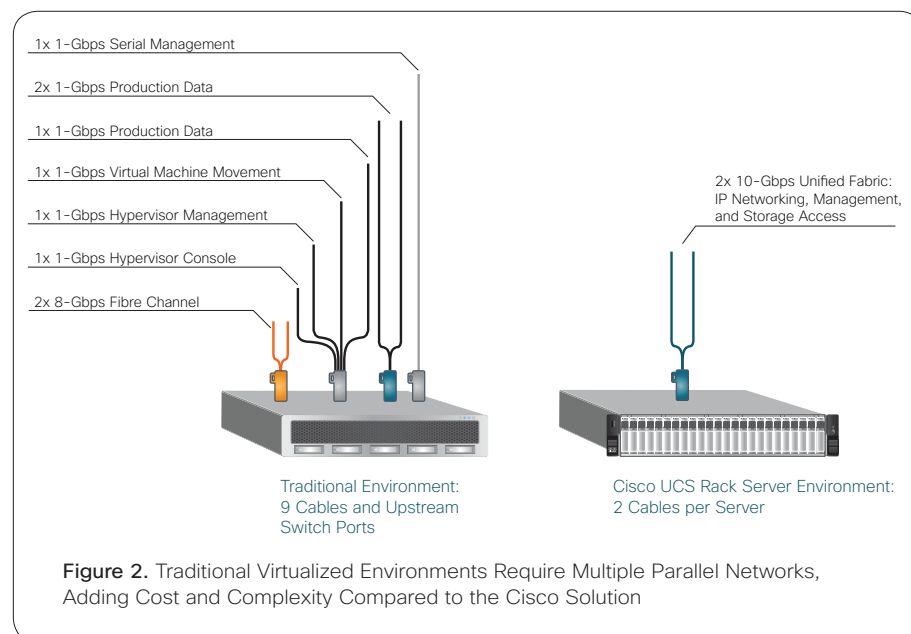
### Cabling for a Cisco UCS Blade Server Environment



#### Problem: Multiple Parallel Networks

Traditional virtualized environments must implement multiple parallel networks to implement the best practices recommended by virtualization software vendors. As in many environments, servers must be connected to IP networks, management infrastructure, and shared storage. The demands of virtualized environments are even more challenging, with best practices requiring separate physical networks for virtual machine production traffic, management of the virtualization software, and movement of virtual machines from server to server.

In 1 Gigabit Ethernet environments, the number of connections can be seven or more per server (Figure 2). Moving to 10 Gigabit Ethernet can reduce the number of connections needed for virtual machine production traffic, but infrastructure functions are often relegated to 1 Gigabit Ethernet links, slowing functions such as virtual machine movement.



Each physical network requires network interface cards (NICs) or host bus adapters (HBAs) in each server, cables, and upstream switch ports, resulting in higher CapEx and OpEx. Each separate network must be sized to handle workload bursts without the capability to share bandwidth between the multiple networks. The complexity of maintaining so much physical infrastructure can lead to cabling errors that can cause downtime or security problems. Server airflow can be obstructed by the massive number of cables, reducing airflow, increasing server temperature, and reducing performance that otherwise would be enhanced by Intel Turbo Boost Technology.

## Virtualizing SAP with Increased Density

Pacific Coast Building Products is the parent company for many subsidiaries that create building materials. Just about every aspect of the company is run on SAP Enterprise Resource Planning (ERP) software. The previous infrastructure was close to the end of its life, so the company needed to upgrade its technology while also reducing the technology footprint through data center consolidation and solution choice.

According to Matt Okuma, enterprise architect for Pacific Coast Building Products, "The size of the server we needed to take us to the next level would have really put a large footprint in our data center. We looked at reducing that footprint. We looked at other blade centers and stuff, but Cisco UCS brought that wire-once technology that was really attractive to us." The company was able to increase virtual machine density on two chassis from 14 to 130 virtual machines, decrease the per-chassis cable count from 15 to 4 cables, and accelerate server provisioning from one week to 30 minutes.

View the video at  
<http://www.youtube.com/watch?v=6WcNVFXAz98>.

## Solution: Cisco Unified Fabric

Cisco UCS is integrated into a high-bandwidth, low-latency, 10-Gbps unified fabric that carries all I/O traffic to the system's fabric interconnects over a single set of cables. Cisco UCS uses a wire-once model in which the system is wired once to meet bandwidth needs, and all features, such as Fibre Channel over Ethernet (FCoE) access to shared storage, are enabled and managed through software.

Use of a single network means fewer cables, adapters, and upstream switch ports, reducing CapEx and OpEx and simplifying the architecture to reduce the chance of errors that can cause downtime. Instead of the waste inherent in sizing separate physical networks for each traffic class, the shared I/O resources in the Cisco Unified Fabric enables more flexible resource allocation: bursts of traffic in one resource class can borrow bandwidth from other classes subject to QoS limitations. For example, instead of limiting the movement of virtual machines to a single 1 Gigabit Ethernet network, Cisco UCS can use the full 10-Gbps unified fabric for the traffic bursts required to move an entire virtual memory image from one server to another.

Hypervisors require virtual machine images to be stored on shared storage so that they can be accessed from all servers in a cluster. For many organizations, SANs meet that need best; however, the cost of discrete Fibre Channel HBAs, transceivers, and optical cables can add significantly to the capital cost of deploying servers. The unified fabric brings FCoE connectivity to each server at no additional cost, giving organizations the capability to choose the shared storage mechanisms that best suit their needs without affecting costs.

Cisco UCS consolidates all I/O modalities at the server, allowing I/O to be split onto separate physical networks (such as native Fibre Channel) at the system's fabric interconnects or higher in the network stack. Cisco virtual interface cards (VICs) present up to 256 PCIe-compliant devices to the hypervisor, with the number and type (NIC or HBA) programmable on demand. Cisco VICs can support all the best practices for network separation by placing different hypervisor traffic flows on separate physical interfaces visible to the hypervisor (Figure 3).

The programmability of Cisco VICs allows Cisco UCS servers to support different hypervisors—and switch between them—at a moment's notice, increasing flexibility and increasing ROI.

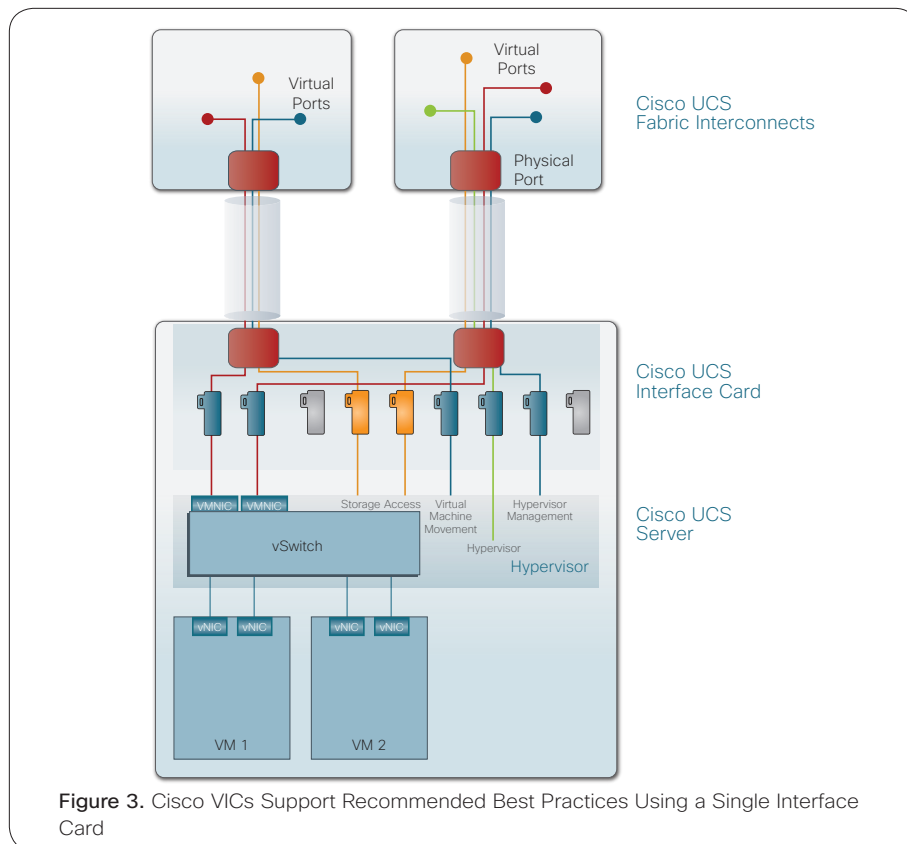
As traffic passes from the VIC and onto the network, traffic from each PCIe device is isolated with air-gap security that is implemented through the IEEE 802.1BR VN-Tag standard. This technology isolates network traffic at the Ethernet layer, terminating each virtual network link at a virtual port within a fabric interconnect. This technology allows each link to be managed as if it were a physical link, giving network administrators complete visibility and control and dramatically simplifying the management of virtualized environments.

## Problem: Multiple Network Layers

The network infrastructure of traditional virtualized environments adds unnecessary cost, complexity, and risk to virtualized environments. In most environments, the network access layer has been fragmented into three layers, making it difficult to maintain visibility and control over network connectivity. To facilitate the movement



Cisco VICs present physical PCIe-compliant devices to hypervisors, supporting them without the need for them to implement Single Root I/O Virtualization (SR-IOV). This approach allows almost any hypervisor or operating system to be supported without additional complexity.

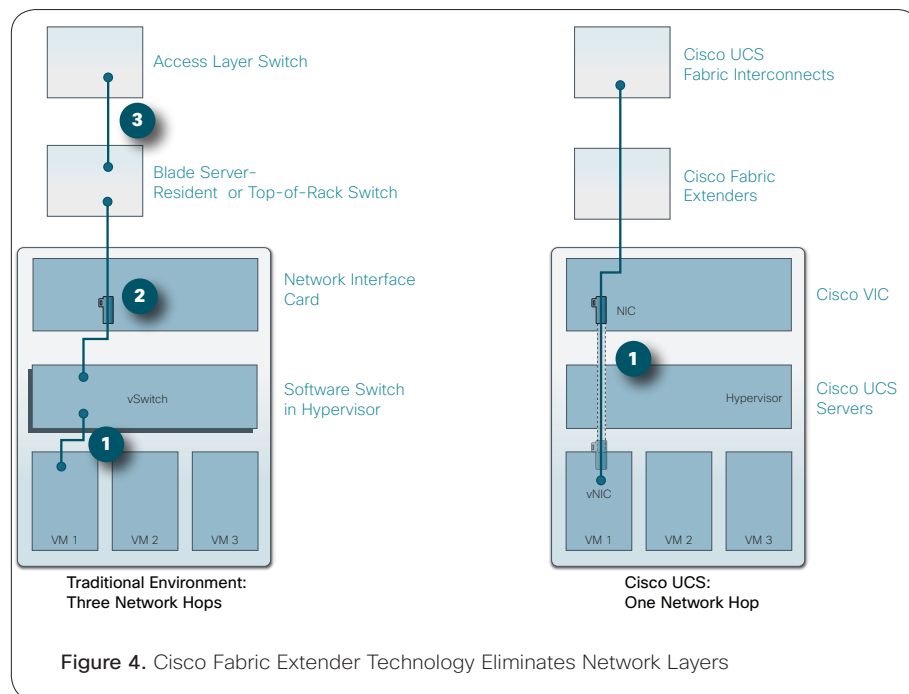


of virtual machines from server to server, security is often reduced so that all three layers do not have to be synchronously changed when a virtual machine moves. These layers add unnecessary and variable latency to virtual networks, and they fragment access-layer management between network and server administrators.

- 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 that reside in blade servers, or additional layers of switches at the top of every rack, add a new layer of switching. Blade-chassis-resident switches are often made by the blade server manufacturer and can have different feature sets than the upstream data center 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 usually configured by server administrators.

Network traffic destined for the access-layer switch requires three network hops, adding significant latency (Figure 4). It also adds inconsistency: traffic between virtual machines on the same server can be switched locally, and traffic between blades in the same blade server can be switched through the blade-server-resident switch. Each path has its own latency characteristics, and it is handled by multiple different switch feature sets. This approach complicates the task of load balancing

across a virtualization cluster because the act of moving a virtual machine from one server to another can affect the latency of the virtual machine's network traffic.



### Solution: Cisco Fabric Extender Technology

Cisco Fabric Extender Technology condenses three network layers into one with an architecture that is physically distributed, but logically centralized. All network traffic passes through the system's fabric interconnects, establishing a single point of management and control regardless of virtual machine location (see Figure 4).

Cisco fabric extenders bring the network fabric to a blade chassis, top-of-rack location, and even to individual virtual machines, passing all traffic in a lossless manner. These low-cost, low-power-consuming devices are physically distributed throughout a Cisco Unified Computing System, but they remain logically part of the fabric interconnects, maintaining a single point of management for the entire system.

In rack server environments, Cisco Nexus® fabric extenders bring the system's unified fabric to the top of every rack, allowing each rack to be self-contained and easily moved within the data center.

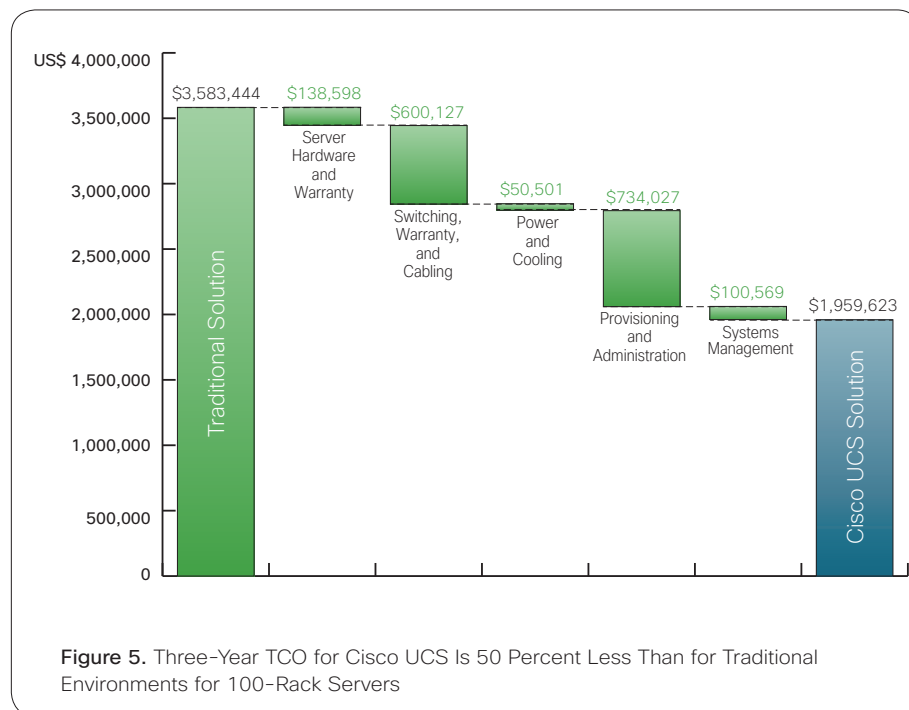
In blade server environments, Cisco UCS fabric extenders bring the system's unified fabric to every blade server chassis, with the current generation of Cisco UCS 6200 Series Fabric Extenders supporting up to 160 Gbps of bandwidth for an eight-blade chassis.

Within each server, Cisco VICs act as adapter fabric extenders, bringing the network fabric directly to hypervisors (for best-practices support) and virtual machines. This approach eliminates the need for software switching, delivering improved network performance while freeing CPU cycles for better application performance.

This graph compares the 3-year TCO for 100 HP ProLiant DL380p Gen8 Servers with the 3-year TCO for 100 Cisco UCS C240 M3 Rack Servers. Each server has two Intel Xeon processor E5-2609 CPUs, 64 GB of memory, and two 300-GB 10K 6G SAS HDDs. The networking component compares HP with six Gigabit Ethernet and two 8-Gbps Fibre Channel connections with the Cisco VIC 1225 dual-port 10-Gbps unified fabric and corresponding switches. Pricing is as of February 8, 2013.

### Dramatic Reduction in TCO

The combination of Cisco Unified Fabric and Cisco Fabric Extender Technology is responsible for the radical simplicity of Cisco UCS. This simplicity translates to 50 percent lower TCO for the Cisco environment than for a traditional 100-node rack server environment (Figure 5). Even in a rack server environment, the cost of top-of-rack switching, cabling, and switches to aggregate each rack into a single access layer is the most significant expense, giving the advantage to Cisco UCS.



### Problem: Traditional Environments Are Difficult to Scale

Virtualized environments decouple server and application deployment. The result is that virtualized—and especially cloud-computing—environments need to be elastic, with the capability to expand and contract as the number of applications increases and as periodic workload bursts come and go. Because IT departments often must respond quickly to rapidly changing conditions, scalability is an important consideration in choosing infrastructure.

In traditional virtualized environments, whether blade or rack server based, scalability is hampered by two factors:

- Server deployment challenges: How the server is moved from the loading dock into production, or how existing servers are repurposed (discussed in the next section)
- Infrastructure integration challenges: How new servers are integrated into the network infrastructure of a virtualized environment

In traditional environments, a significant amount of network infrastructure must be deployed to accommodate new servers, and in some cases the infrastructure must

### Network Infrastructure in Each Traditional Rack:

- 2 Production Fibre Channel Switches
- 2 Production Ethernet Switches
- 2 Management Network Switches



### Scalability with Half the Cost and Complexity

Server costs are significant, but so is the cost of the infrastructure to support each server. Cisco Fabric Extender Technology dramatically reduces the number of interfaces, cables, and switches needed to support Cisco UCS blade servers. The result is that the average per-server infrastructure cost is US\$2343 for Cisco UCS compared to US\$3761 for an HP system.\* Every time an HP chassis is added to a rack, it imposes a cost of up to US\$39,739 more than the cost of adding one more chassis to the Cisco Unified Computing System.

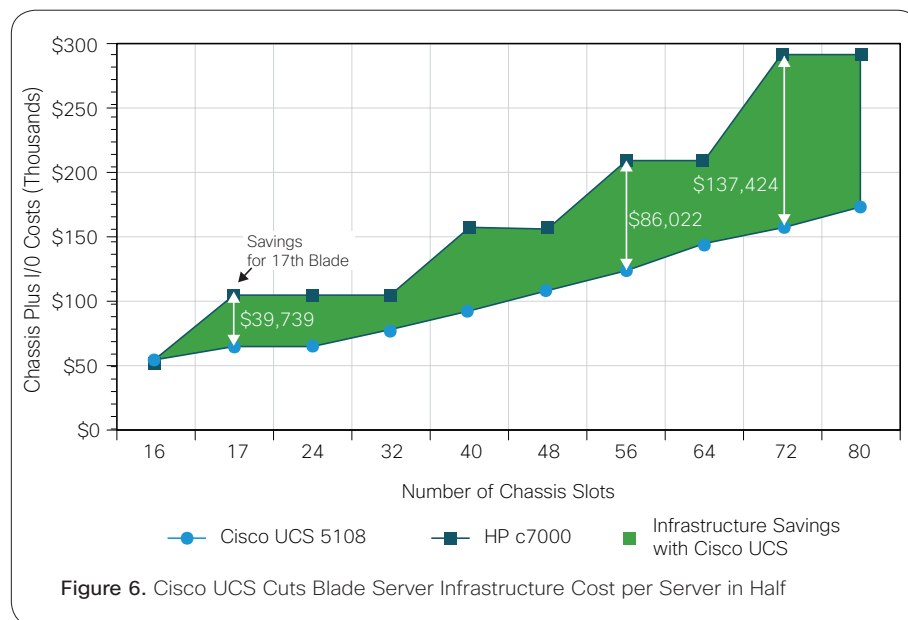
\* Based on the Cisco UCS manufacturer's suggested retail price (MSRP) and HP retail price on January 4, 2012.

be redesigned as the virtualization cluster is scaled up, resulting in not only capital cost, but the opportunity cost of the time lost while the infrastructure is adjusted to accommodate the new servers. In addition, because each new infrastructure component in traditional environments is a new management point, complexity increases as the system scales up.

The capital cost can be seen clearly in the scenario in which one more server than the container size naturally accommodates is added in a traditional environment: what can be called the N+1 problem:

- In traditional rack server environments, an entire set of top-of-rack switches must be purchased and installed to support one more server than the existing racks hold. Then these switches must be integrated into the existing network infrastructure, and sufficient upstream switch ports must be present and ready to support them.
- In traditional blade server environments, a new blade server chassis must be purchased and outfitted with the numerous chassis modules that represent profit centers for traditional manufacturers: redundant Ethernet switches, Fibre Channel switches, and chassis-management modules. The blue line in Figure 6 shows the cost of scaling an HP blade server environment.

In both environments, the infrastructure cost and complexity per server impede scaling.



### Solution: Low-Cost Scaling with Cisco UCS

Cisco UCS has a unified, simplified network infrastructure that scales at lower cost because it is designed from the beginning as a single unified system that is centrally managed yet physically distributed. Cisco UCS supports elasticity in virtual environments with a predefined network architecture that is self-aware and self-integrating, responding quickly to workload spikes and reducing costs.



#### Network Infrastructure in Two Cisco UCS Racks:

- 2 Cisco Nexus Fabric Extenders  
Serving 2 Racks



With Cisco UCS, low-cost, zero-management fabric extenders replace the top-of-rack switches in traditional rack server environments, and the multiple chassis modules in blade-server environments. With Cisco UCS, even the blade server chassis is a comparatively low-cost item consisting of sheet metal, power supplies, and fans, and when integrated into the system it also becomes logically part of the centrally managed system.

Cisco UCS thus mitigates the effects of the N+1 problem because the cost of equipping each new container is less, whether the equipment is a new rack with top-of-rack fabric extenders or a new blade server chassis with fabric extenders inside.

## Consolidate 168 Management Points to Two with Cisco UCS

When NetApp needed to deploy a scalable testing cloud capable of hosting 23,000 virtual machines, the company's engineering support services chose Cisco UCS. The first step consolidated 714 existing servers onto 120 blade servers on a single Cisco UCS platform, reducing 168 management points to just 2: the pair of Cisco UCS 6100 Series Fabric Interconnects.

See [http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/case\\_study\\_NetApp\\_Cisco\\_Kilo\\_Lab.pdf](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/case_study_NetApp_Cisco_Kilo_Lab.pdf).

## Unified Management

In the four years since Cisco UCS was first announced, it has remained the only system built from the beginning so that server personality, configuration, and connectivity can be managed through a single point. Cisco UCS Manager, the system's integrated, embedded management system, automates server configuration, accelerating server deployment and helping make virtualized environments more responsive to workload changes. A single Cisco UCS management domain can support up to 160 servers. Cisco UCS Central Software can aggregate multiple Cisco UCS management domains to support up to 10,000 servers through a single management interface. This scalability can meet the needs of even the largest virtualized and cloud computing environments, and can give IT departments considerable flexibility in mapping virtualization clusters to physical Cisco UCS instances.

### Problem: Server Deployment Challenges Slow Scaling

Virtualized environments, and especially cloud computing environments, must scale rapidly to meet workload demands. Traditional environments, however, depend on days-long server configuration processes to add servers to a virtualization cluster.

Server administrators must install the correct firmware, specify BIOS settings, and configure internal RAID controllers, if used. Network administrators must install NICs and configure them and their upstream switches to manage traffic properly. Storage administrators must install HBAs and configure them and their upstream switches as well. This ponderous, error-prone process can make IT departments ask why the physical infrastructure they support should be any less flexible and scalable than the virtualized environments that run on top of them. Both the virtual and the physical environments must be able to scale in concert for organizations to reap the benefits of virtualization.

### Solution: Automated Configuration with Cisco UCS

Cisco UCS is a self-aware, self-integrating system that was designed from the beginning so that every aspect of its network infrastructure and its server configuration can be automated through software. Cisco UCS Manager resides on the system's fabric interconnects. This model-based manager automates system configuration to accelerate the deployment of new servers in virtualized environments.

Cisco UCS Manager brings a system to life by managing the entire system as a single logical entity. It can be accessed through an intuitive GUI, a command-line interface (CLI), and an XML API that allows virtualization and cloud computing software to manage the physical infrastructure directly. For example, virtualization software can interact directly with Cisco UCS when it manages power use by removing virtual machines from underutilized servers and powering down the servers.

Cisco UCS Manager automates and simplifies the process of incorporating new servers into a virtualization cluster so that new servers can be installed, configured, and put to use in minutes. This capability does more than just increase IT staff productivity. The capability to quickly scale a virtualization cluster offers strategic advantages to organizations that use Cisco UCS.

## Intelligent Infrastructure Means Rapid Deployment

In designing a new center in Scottsdale Arizona, NightHawk Radiology Services looked for technology with the highest level of redundancy and resiliency, as well as technology that would enable rapid deployment

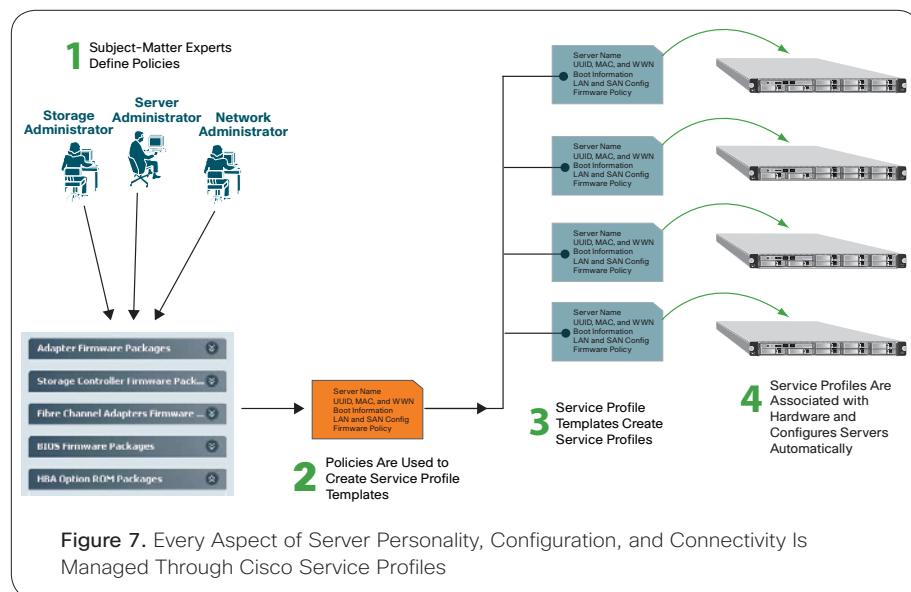
Ken Brande, the company's vice president for IT, reports on the company's success deploying Cisco UCS: "With Cisco UCS Manager service profiles we can very quickly reconfigure any server blade so that it's ready for production in 15–20 minutes. Rapid configuration is critical in our environment, where a server outage is simply unacceptable."

See [http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/case\\_study\\_c36\\_604664\\_ns944\\_Networking\\_Solutions\\_Case\\_Study.html](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/case_study_c36_604664_ns944_Networking_Solutions_Case_Study.html).

Cisco UCS Manager is model based in the sense that it maintains an internal, logical model of the system components and their state, detecting any new devices and integrating them into the system automatically. Cisco service profiles are logical entities that completely define a server's personality, configuration, and connectivity, including firmware revisions, universal user IDs (UUIDs), MAC addresses, worldwide names (WWNs), and the number and type of devices dynamically configured in a Cisco VIC. Associating a service profile with a physical server in Cisco UCS Manager completely configures that server.

In virtualized environments, where complete consistency of server configurations is required, Cisco service profile templates define how to create a service profile, and they can be used to generate a service profile for each server in the system (Figure 7). Configuring new servers or repurposing existing ones takes minutes, and it can even be done automatically by assigning service profiles to specific blade server slots. This means, essentially, that computing infrastructure can keep up with the virtual environment, compliant configurations are guaranteed, and configuration drift is eliminated. IT departments have greater flexibility to redeploy servers to meet temporary workload needs, and they can even switch between hypervisor software by modifying service profiles to reflect a different vendor's best practices.

The role- and policy-based management model of Cisco UCS Manager supports the existing separation of responsibilities in IT departments. It also supports subject-matter experts by freeing them from the tedious, error-prone tasks of server administration, and it allows them to focus on creating policies that any level of administrator can use when creating templates and Cisco service profiles (see Step 1 in Figure 7).



## Problem: Managing Growth

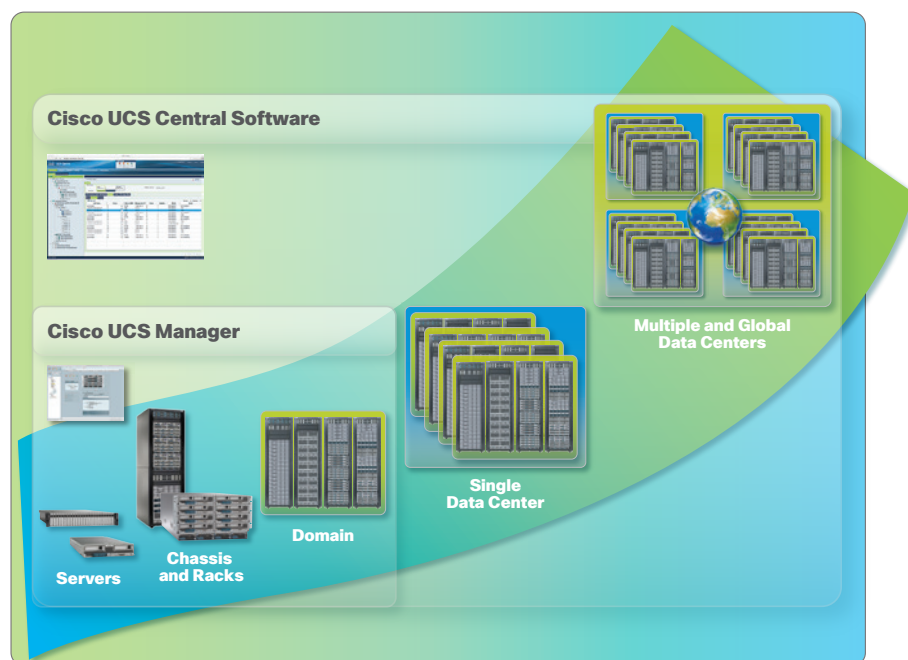
As the use of virtualization increases in an organization, the need to manage multiple virtualization clusters within the same data center and, often, virtualization clusters across the world poses challenges for the traditional environment. IT departments

must be ready to manage massive growth, but they also must not neglect traditional management concerns, including:

- Standards compliance
- Cost containment
- Inventory management
- Business agility
- Compliance with service-level agreements (SLAs)

**Solution: Manage Multiple Cisco UCS Domains with Cisco UCS Central Software**

Cisco UCS Central Software centrally manages multiple Cisco UCS domains using the same concepts that Cisco UCS Manager uses to support a single domain (Figure 8). Cisco UCS Central Software manages global resources (including identifiers and policies) that can be consumed by individual Cisco UCS Manager instances. It can delegate the application of policies (embodied in global service profiles) to individual domains, where Cisco UCS Manager puts the policies into effect. Cisco UCS Central Software can support up to 10,000 servers in a single data center or distributed around the world in as many domains as are used for the servers. The benefits to virtualized environments are plentiful:



**Figure 8.** Cisco UCS Central Software Manages Multiple Cisco UCS Domains in a Single Data Center or Multiple Distributed Data Centers

- Global inventory management: Cisco UCS Central Software maintains a global inventory of all components connected to any Cisco UCS domain that it manages, giving organizations an instant, up-to-date view of the resources they can use to support their virtualized environments.



- Automated standards compliance: Cisco UCS Central Software can maintain global policies that can be used to enforce consistent server identity, configuration, and connectivity policies on a global scale.
- Increased business agility: Server configurations can be adjusted to meet changing workload conditions within minutes, with a global view of the resources that are applied to meet particular application requirements.
- Compliance with SLAs: Greater workload mobility accelerates the alignment of capital resources with the needs of virtualization clusters.
- Cost containment: Operation statistics makes simplify the process of assessing and managing server utilization on a global scale. Software licenses can be used more effectively by shifting server identities to locations where virtualization clusters require more hypervisor instances, reducing the need to size and license each cluster to maximum capacity.

#### Problem: Managing and Securing Virtual Networks

In traditional environments, physical and virtual networks are managed as two separate entities that limit visibility and control from one environment to another. For example, for virtual machines to move between servers, their network profiles (which include characteristics such as VLAN membership and QoS settings) must be accepted on every physical switch on which they might appear. This requirement leads to a least-common-denominator approach to security, making it easier to loosen controls on the physical network than it is to synchronize the physical and the virtual.

Virtual network management and troubleshooting is difficult because virtual networking hides virtual machine traffic from the standard management tools that network administrators routinely use to troubleshoot and manage network problems. For example, if a rogue virtual machine overburdens the network with chatter, administrators must first locate the server on which the problem is occurring and then migrate each virtual machine from that server to a different server until the offending virtual machine is identified. Loss of virtual network visibility and control can make security and QoS management difficult.

#### Solution: Cisco Data Center Virtual Machine Fabric Extender

Cisco VICs can provide static interfaces to hypervisors to comply with recommended best practices. Cisco VICs also can provide dynamic interfaces to virtual machines, allowing the network to connect directly to virtual machines, increasing network performance by up to 38 percent by bypassing hypervisor switches, and freeing the host CPU to deliver better application performance.

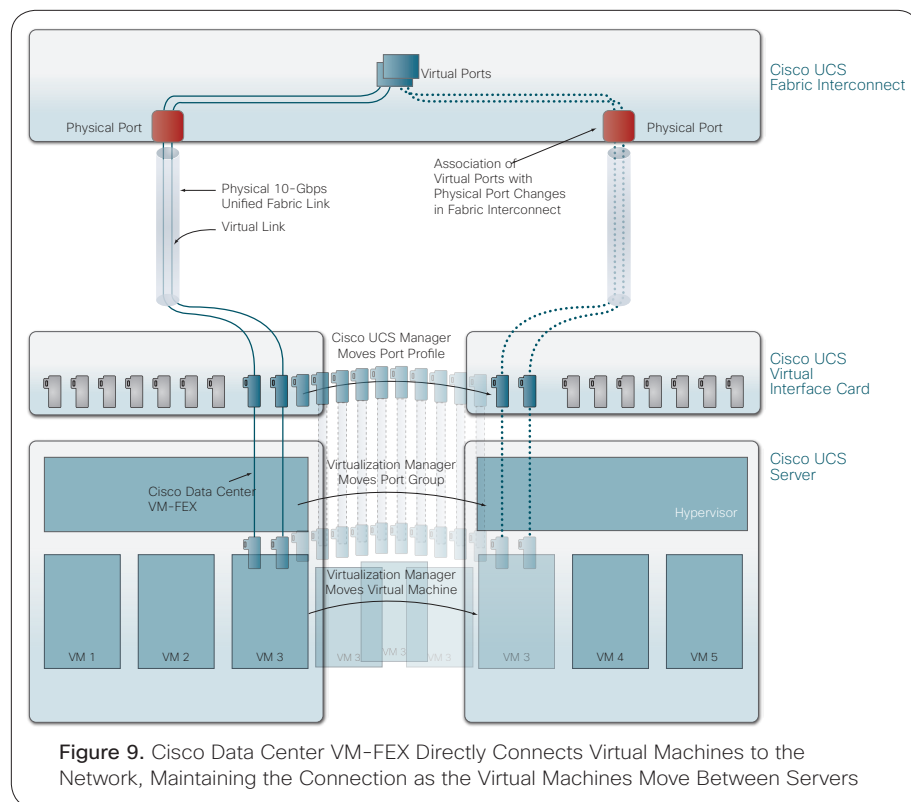
This Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology is a significant innovation that benefits virtualized environments by allowing virtual machines to connect to the network exactly like physical servers do, bringing together the scalability of virtual networks and the manageability of physical networks. Within Cisco UCS, physical and virtual networks are managed in the same way, with physical and virtual network links terminated by physical and virtual ports within the system's fabric interconnects. Now, for example, a rogue virtual machine can be identified by the traffic on its port, and the port can easily be disabled, using essentially the same process as a network administrator would use to handle such an incident on a physical server.

### Rapidly Provision Virtual Machines

Euronet Worldwide, an industry leader and provider of highly secure electronic financial transaction solutions, deployed Cisco UCS and reduced the time needed for virtual server implementation and provisioning by up to 95 percent compared to the time needed for its former infrastructure. The new implementation has resulted in decreased power consumption, cooling needs, and rack space.

See <http://www.marketwire.com/press-release/euronet-deploys-cisco-unified-computing-system-and-cloud-infrastructure-nasdaq-csco-1507480.htm>.

The impact of this innovation on security is equally profound. Cisco UCS Manager coordinates with the major hypervisors to maintain a virtual machine's network profile regardless of the virtual machine's location. Security does not have to be compromised to allow virtual machines to migrate freely from server to server. Now when virtual machines migrate, their network links migrate with them (Figure 9). From the perspective of the fabric interconnects, the virtual port associated with a specific virtual link simply changes the physical port to which it is associated.

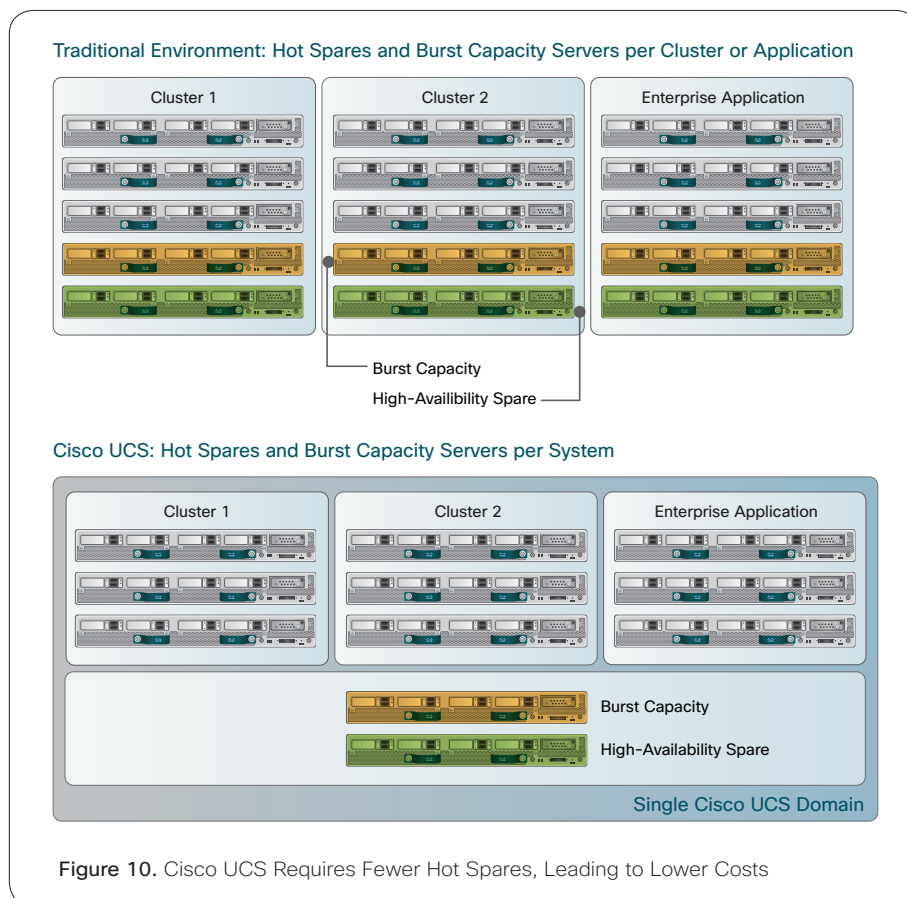


Cisco Data Center VM-FEX helps organizations maintain existing administrator roles across bare-metal applications and virtualized and cloud-computing environments. Because the network remains always in the domain of network administrators, this technology eliminates the overlap of server and network administrator roles that often occurs in virtualized environments.

### Problem: Managing Spares and Licenses

Most organizations using a cluster of servers to support virtualization—or any other mission-critical application—have one or more hot spares ready to replace any server in the event of a failure. Traditional environments usually have one or more hot spares for each virtualization cluster or application, each configured for a specific need and with its own licensed software installed (Figure 10).

This approach leads to higher costs because multiple spares must be maintained and licensed software must be purchased to run on them.



### Solution: Fewer Spares Through Dynamic Configuration

With Cisco UCS, fewer hot spares are needed because server identity, personality, and configuration are all programmable, so a single spare can be made ready to serve a new role within minutes. With the Cisco VIC, the number and type of I/O devices can be configured through Cisco service profiles, and even the server identity can be changed to match that of a failed server. The result is that fewer spares are needed and licensed software can be more effectively utilized with no need to pay for licenses that are not in use.

## Unified Computing

Cisco UCS is based on industry-standard, x86-architecture servers with Cisco innovations and intelligent Intel Xeon processors. Although many vendors offer servers with the same processors, Cisco integrates them into a system with a better balance of resources. This balance brings processor power to life with at least 70 world-record-setting benchmark results that demonstrate higher virtualization performance and increased consolidation ratios.

### Cisco UCS Overcomes Bottlenecks for Pitt Ohio

Pitt Ohio is a technology leader in the transportation industry. This leadership helps the company attract and retain customers. It is critically important to upper management that the company maintain this reputation. The company uses technology to be agile enough to respond to customer requests and create custom solutions that simply do not exist in other organizations.

According to Jules Thomas, senior systems engineer at Pitt Ohio, "We wanted to go to a full virtual data center. To do so, you need very high resource services for your VMs—along with memory, CPU, storage, and so forth. Prior to our migration to UCS, we had to stop our virtualization because we didn't have the horsepower to accommodate all the VMs that we wanted to get in place."

View the video at [http://www.cisco.com/en/US/prod/collateral/contnetw/ps5680/ps6870/prod\\_case\\_study\\_PittOhio\\_video.html](http://www.cisco.com/en/US/prod/collateral/contnetw/ps5680/ps6870/prod_case_study_PittOhio_video.html).

Cisco offers a range of servers for scale-out and Web 2.0 applications, servers with more power and reliability for enterprise-class applications, and servers with the utmost reliability for mission-critical applications. These servers are equipped with the two most advanced microprocessor families from Intel:

- Intel Xeon processor E7 family: The Intel Xeon processor E7 family is designed to meet the mission-critical IT challenge of managing and keeping business-critical data secure. Powerful, reliable servers such as the Cisco UCS C460 M2 High-Performance Rack Server are equipped with the top-of-the-line Intel Xeon processor E7 family to deliver performance that is excellent for the most data-demanding workloads, with improved scalability and increased memory and I/O capacity. These features help businesses quickly adapt to short-term changes in business needs while addressing requirements for long-term business growth. Advanced reliability and security features help maintain data integrity, accelerate encrypted transactions, and increase the availability of mission-critical applications. The powerful and reliable Intel Xeon processor E7 product family delivers flexibility for business-critical solutions.
- Intel Xeon processor E5 family: The Intel Xeon processor E5 family is at the core of a flexible and efficient data center that meets diverse business needs. This family of processors is designed to deliver versatility, with the best combination of performance, built-in capabilities, and cost effectiveness. The Intel Xeon processor E5 family delivers exceptional performance to a broad range of data center environments and applications: from virtualization and cloud computing to design automation and real-time financial transactions. With these processors, I/O latency is dramatically reduced with Intel Integrated I/O, which helps eliminate data bottlenecks, streamline operations, and increase agility.



### Problem: Virtualized Environments Underperform

Many companies do not achieve the promise of virtualization. Traditional server and networking technologies can limit application performance because of an imbalance in resources.

### Solution: Cisco UCS Is Optimized for Virtualization

Cisco UCS has an architectural advantage that delivers better performance through a better balance of resources. Intelligent Intel Xeon processors are central to the delivery of better virtualization performance, and the rest of the system is designed to enhance that performance. Better airflow means greater potential for engaging Intel Turbo Boost Technology, which boosts CPU clock rates when thermal conditions permit. Intel VT-d Technology accelerates I/O from virtual machines to Cisco VICs, and Cisco VICs deliver network throughput that is limited only by the PCIe buses on which they reside. Cisco has been a leader in offering high memory



For more information about VMware VMmark benchmark results, please visit <http://www.cisco.com/go/ucsatwork>.

See [http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/whitepaper\\_c11\\_703103.pdf](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/whitepaper_c11_703103.pdf) for more information.

See [http://www.cisco.com/en/US/prod/collateral/ps10265/ps10281/whitepaper\\_c11-711915.html](http://www.cisco.com/en/US/prod/collateral/ps10265/ps10281/whitepaper_c11-711915.html) for more information.

For more information, please refer to [http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns377/ucs\\_b200\\_vdi\\_0312.pdf](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns377/ucs_b200_vdi_0312.pdf).

capacities that allow more virtual machines to run on each server. All of these factors contribute to proven virtualized performance.

Cisco has demonstrated long-term leadership by establishing 14 world records on VMware® VMmark™ benchmarks that evaluate not just virtualization performance, but also the effectiveness of the underlying infrastructure in performing common tasks such as virtual machine migration. Cisco's outstanding scalability is due to the combination of technologies, including its server designs, large memory capacity, and unified fabric, with Cisco fabric extender technology.

When deploying virtualization clusters on Cisco UCS, organizations can use both blade and rack servers to meet their specific requirements. Both blade and rack servers are managed through single-wire management, with automation provided by Cisco UCS Manager and optional Cisco UCS Central Software.

#### Optimized Performance for Virtualized SAP Environments

Cisco innovations result in tangible application performance improvements. In SAP environments, Cisco Data Center VM-FEX technology eliminates the burden of virtual machine switching from the hypervisor, giving SAP applications access to more host CPU cycles and I/O throughput.

The combination of Cisco UCS and Cisco Data Center VM-FEX technology delivers an integrated solution that enables performance gains through support for up to 11 percent more users and through reduced database access latency, effectively accelerating user queries by up to 29 percent.

#### Increased Performance for Virtualized Microsoft SQL Server

The capability of Cisco Data Center VM-FEX technology to reduce network latency, increase throughput, and provide more CPU power for application performance is proven by Cisco's study of Microsoft SQL Server 2012 running on Microsoft Windows Server 2008 R2 Hyper-V.

When Cisco Data Center VM-FEX is used, the virtualized environment delivered 12 percent faster query response time than when a traditional hypervisor software switch is used. This I/O-intensive environment further demonstrated a 30 percent gain in I/O throughput and a 43 percent reduction in disk latency when the two environments were compared with and without the software switch in the hypervisor.

#### Problem: Memory Constraints

Many environments run out of memory before they run out of CPU, and upgrading to four-core processors to add more memory capacity is costly, requiring CPU upgrades and additional software licensing.

#### Solution: Cisco Leads in Memory Capacity

Cisco was the first to recognize the memory shortage problem and was the first to respond, delivering patented Cisco Extended Memory Technology to enable more effective resource utilization and lower licensing costs. Today, Cisco still leads the way by delivering one of the highest-density blade servers, with up to 768 GB in a half-width form factor, and rack servers with up to 2 TB of main memory. Choosing

Cisco means that you choose a vendor committed to staying ahead of the industry with its focus on this concern.

The benefit of Cisco's memory capacity is seen in the density that can be achieved in virtual desktop infrastructure workloads running on Cisco UCS, with more than 185 desktops running the knowledge worker profile on a single 2-socket server.

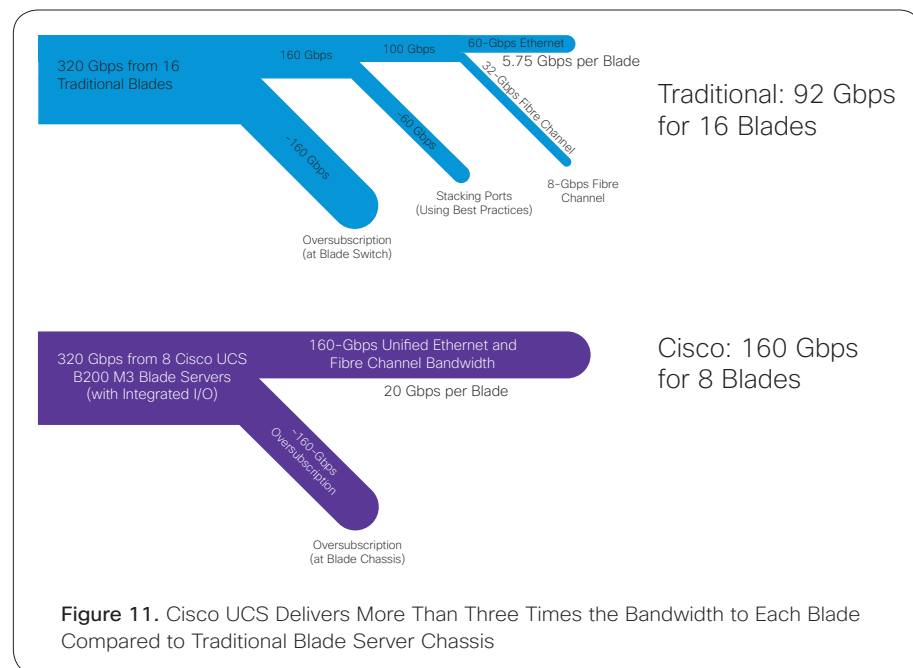
#### Problem: Insufficient I/O Bandwidth

With innovators such as Intel regularly increasing CPU power, organizations with traditional systems are increasingly locked into architecture that limits bandwidth and the flexibility to shift I/O resources to meet short-term bursts in traffic.

#### Solution: The Right Kind of Bandwidth

Today's high-performance, multicore Intel Xeon processors demand the utmost in I/O connectivity to perform well, and Cisco is positioned to deliver that bandwidth today and well into the future. Three kinds of bandwidth are needed to deliver excellent virtualization performance:

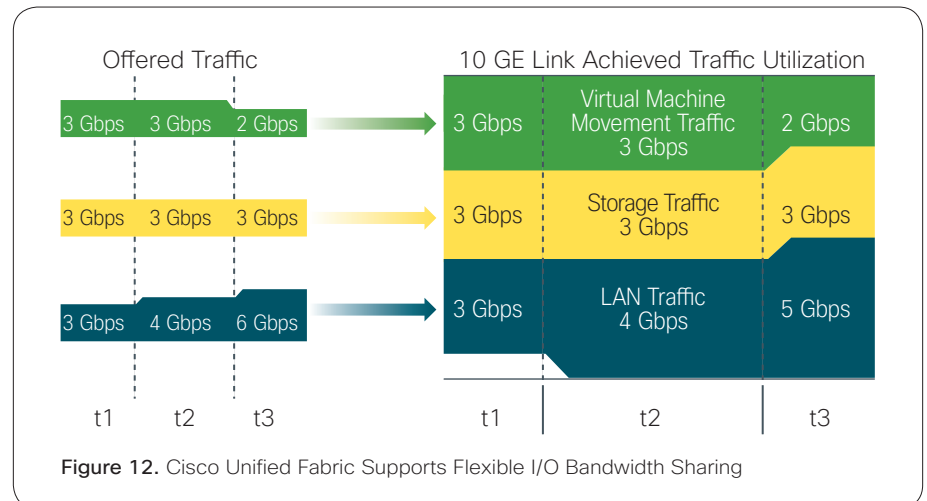
- Average bandwidth: At times, virtualized workloads can put significant, sustained I/O workloads onto the network infrastructure. The Cisco UCS 5108 Blade Server Chassis equipped with Cisco UCS 2208XP Fabric Extenders supports up to 160 Gbps per 8-blade chassis, for an average of 20 Gbps of bandwidth per half-width blade (Figure 11). Traditional environments lose significant amounts of bandwidth when stacking is required, with a traditional 16-blade chassis delivering only 5.75 Gbps average bandwidth per blade split between Ethernet and Fibre Channel connectivity.



- Peak bandwidth: Like any workload, virtualized workloads tend to be bursty. The Cisco blade server chassis connects each half-width blade server with up

to 80 Gbps of bandwidth, and each full-width blade with up to 160 Gbps of bandwidth. Although actual achieved peak bandwidth is less because of server PCIe bus constraints, the bandwidth available through the Cisco blade server chassis testifies to the longevity of the platform and its capacity to support future generations of blade servers with future generations of bus speeds.

- Shared bandwidth: A flexible I/O infrastructure can share network bandwidth between servers and also between I/O modalities such as IP networking, storage access, and management networks. Traditional environments require decisions at installation time about how much bandwidth to partition for Ethernet compared to Fibre Channel connectivity. After these decisions are made, they become fixed because a separate set of cables is used from the server on up. When traffic bursts arrive, there is no way to share some Fibre Channel bandwidth to help an overloaded Ethernet network.
- With the wire-once philosophy of Cisco UCS, the system is wired once for bandwidth, and different I/O modalities can be managed and bandwidth shared through QoS rules that administrators can determine. As Figure 12 illustrates, if a burst of Ethernet traffic pushes the demand past the allocated bandwidth, the traffic can share bandwidth ordinarily used for other purposes, such as virtual machine movement, if the guaranteed minimum bandwidth for this service class is met (see time t3). This kind of sharing cannot be accomplished if each traffic class has its own dedicated physical networks.



## Easy Deployment

IT departments have a wide range of choices for deploying virtualized environments on Cisco UCS. These options include the use of tested and validated designs to configure systems in house, deployment of prepackaged solutions, and engagement of Cisco Services for part or all of the implementation.

### Cisco Validated Designs

Cisco Validated Designs make deployment of virtualized environments easier. Cisco Validated Designs describe solutions using Cisco UCS blade and rack servers that are designed, tested, and documented, much like a deployment recipe, to facilitate, simplify, and improve customer deployments. These designs incorporate a wide range of technologies and products into solutions that have been developed to address the business needs of Cisco customers.

### Cisco SmartPlay Solutions

Cisco has created bundles of the top-selling Cisco UCS blade and rack servers at highly competitive prices. These bundles are created based on application-specific Cisco Validated Designs to accelerate both procurement and deployment of data center application infrastructure. For organizations wanting to deploy virtualized environments, Cisco offers SmartPlay solutions using Microsoft Windows 2012 R2 Hyper-V and VMware vSphere, all configured to help small and medium-sized businesses begin their journey with Cisco UCS.

### Virtual Computing Environment Coalition Vblock Systems

Virtual Computing Environment (VCE) coalition Vblock™ systems accelerate the adoption of converged infrastructure and cloud-based computing models. The VCE Vblock system is optimized to help ensure secure and predictable performance through preengineered and modular infrastructure. The infrastructure integrates VMware virtualization software, Cisco UCS blade servers, and EMC storage.

### Health Benefit Administrator Cures Growing Pains with VCE Vblock System-Based Private Cloud

With 70 million patients under care, CareCore National launched a private cloud with two Vblock systems. The organization reduced the time required to launch new lines of business from six months to two weeks, freed software engineers to increase their development time from 50 to 80 percent, and enabled call-center agents to handle 20 percent more calls daily.

Read the case study at [http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/CaseStudy\\_CareCore.pdf](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/CaseStudy_CareCore.pdf).

### NetApp FlexPod

FlexPod data center solutions from Cisco and NetApp provide validated, shared infrastructure that scales to accommodate many different application workloads and environments. The basis of the solution is Cisco UCS and NetApp storage, creating a unified architecture across computing, networking, and storage layers. Virtualization solutions span VMware vSphere and View, Microsoft Windows Server Hyper-V, Citrix XenDesktop, and Red Hat Enterprise Linux (RHEL).

You can access Cisco Validated Designs at the Cisco website at [http://www.cisco.com/en/US/netsol/ns743/networking\\_solutions\\_program\\_home.html](http://www.cisco.com/en/US/netsol/ns743/networking_solutions_program_home.html).

More information about Cisco SmartPlay bundles can be found at <http://buildprice.cisco.com>.

More information about VCE Vblock systems can be found at <http://www.cisco.com/go/vblock> or on the VCE website: <http://www.vce.com>.

More information about FlexPod can be found at <http://www.cisco.com/go/flexpod>.



#### **FlexPod Transforms Data Center**

Hedrick Automotive Group built two virtualized, co-located data centers with Cisco UCS, VMware, and NetApp architecture and accelerated new application deployment time from weeks to a day. The solution boosted IT staff productivity by more than 30 percent and helped transform the company's data center from a cost center to a revenue center.

Read the case study at [http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/hendrick\\_external\\_case\\_study\\_fnl\\_11\\_28\\_12.pdf](http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/hendrick_external_case_study_fnl_11_28_12.pdf).

#### **Cisco Services**

Cisco offers a broad range of services for data center optimization through server and network consolidation, cloud enablement, and desktop virtualization. Cisco Services works with customers to develop architectural strategies, roadmaps, and designs and validate, implement, and migrate (if necessary) new virtualized solutions. Cisco Services also work to optimize the environment for application performance and simplified service management. Cisco wants you to succeed at every phase of your data center transformation with services from Cisco and our partners.

## Conclusion

Cisco UCS is a single unified system that helps virtualized environments overcome the limitations of traditional environments. Cisco UCS scales better and more rapidly and with lower infrastructure cost. It delivers greater performance through a better balance of resources. It simplifies management for rapid deployment of physical and virtual machines, and for the first time equates physical and virtual networking to provide increased visibility and control. Cisco UCS increases flexibility, letting you move workloads between servers without worrying about changes in network latency, and it allows resources to be shared so that the environment can expand and evolve without the constraints of physical partitioning between networks.

Cisco developed Cisco UCS to meet the demands of virtualized environments, and it has already supported multiple generations of server, network, and virtual interface technologies. For a virtualized environment that needs to excel both today and into the future, Cisco UCS is the choice that delivers on the promise of virtualization.

## For More Information

- For more information about Cisco UCS, please visit <http://www.cisco.com/go/ucs>.
- For more information about Cisco virtualization solutions, please visit <http://www.cisco.com/en/US/netsol/ns1145/index.html#~overview>.
- For more information about Cisco UCS virtualization performance, please visit <http://www.cisco.com/go/ucsatwork>.
- For more information about application performance using Cisco Data Center VM-FEX, please visit <http://www.cisco.com/go/vmfex>.



**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](http://www.cisco.com/go/offices).