

Private Cloud Computing for Enterprises: Meet the Demands of High Utilization and Rapid Change

What You Will Learn

Enterprise data centers are faced with a critical challenge: The number of applications and amount of data in the data center continue their rapid growth, while IT struggles to provide the resources necessary to make services available to users and meet today's demands using existing infrastructure and organizational silos.

For too long, this siloed approach has hindered IT from adjusting dynamically to new business requests. In existing silos, application workloads are tightly coupled to physical assets, with software linked to operating systems to manage availability, enforce security, and help ensure performance. This tightly coupled model has resulted in proliferation of server and storage devices, with attendant costs and maintenance overhead, to meet user demand.

Unfortunately, only a small portion of each dollar spent on IT today creates a direct business benefit. Customers are spending approximately 70 percent of their budgets on operations, and only 30 percent on differentiating the business. Since data center IT assets become obsolete approximately every 5 years, the vast majority of IT investment is spent on upgrading various pieces of infrastructure and providing redundancy and recoverability: activities that consume approximately 60 to 80 percent of IT expenditures without necessarily providing optimal business value or innovation. In addition, the return on investment (ROI) for new IT projects can take anywhere from 9 to 18 months.

As a result, IT has been forced to focus on simply keeping the data center running rather than on delivering the kind of innovation that meets user needs for faster, better services while also meeting requirements and ensuring business agility.

What is needed is a solution with the scale, flexibility, and transparency to enable IT to provision new services quickly and cost effectively by using service-level agreements (SLAs) to address IT requirements and policies, meet the demands of high utilization, and dynamically respond to change, in addition to providing security and high performance.

Cloud computing provides a solution for meeting this challenge, and using the Cisco® portfolio of products, enterprises can begin moving toward a cloud computing solution today.

The Enterprise Cloud Computing Model

Cloud computing is being proposed as one answer to the challenges of IT silos—inefficiencies, high costs, and ongoing support and maintenance concerns—and increasing user demand for services. While transactional data and high-performance file share applications are best handled now within enterprise data centers, cloud computing is demonstrating its ability to handle the increasing Internet data from rich web applications, services from online service providers, large data processing jobs, and digital media creation with follow-on global distribution.

In addition, many industry analysts and businesses agree on the cloud computing opportunity. In 2008, total spending on cloud computing services was estimated to have a 27 percent compound annual growth rate (CAGR) between 2008 and 2012, growing from US\$16 billion to US\$42 billion, according to IDC; cloud services are projected to grow at 5 times the rate of current enterprise IT spending.

Barriers to Adoption of Cloud Computing Model

While customers agree about the benefits of cloud computing, they have concerns about cost and flexibility, in particular about security, compatibility with existing applications, lack of a migration path from existing applications to clouds, freedom of choice, federation of internal and external resources, lack of SLAs for policy-based management, and interoperability.

The good news is that these challenges can be resolved using a cloud computing approach that includes computing, network, virtualization, and storage resources. Cisco is facilitating the first step for the enterprise toward a multi-tenant cloud computing solution with a unique approach toward enterprise cloud computing: private clouds.

Why Cloud Computing?

According to the Gartner Group, cloud computing is best defined as an infrastructure utility that is open, flexible, predesigned and standardized, virtualized, and highly automated, as well as secure and reliable.

Cisco refines this definition, defining cloud computing as a model that provides resources and services that are abstracted from the underlying infrastructure and provided on demand and at scale in a multi-tenant environment. In addition to its on-demand quality and its scalability, cloud computing can provide the enterprise with:

- Global distribution capabilities, with policy-based control by geography and other factors
- Operational efficiency from consistent infrastructure, policy-based automation, and trusted multi-tenant operations
- Integrated service management interfaces native to the cloud
- Regulatory compliance through automated data policies
- Lower-cost computing

Many companies are rapidly evolving toward cloud computing, though from different starting points and not without debate as to the best direction or computing model. For example, advocates of public cloud computing sometimes advise not owning any software or hardware or employing any IT administrators and instead relying on professional providers of IT applications, platforms, infrastructure, and services. The original Wikipedia definition of cloud computing emphasized that position:

“Cloud computing is Internet (‘cloud’) based development and use of computer technology (‘computing’). It is a style of computing in which dynamically scalable and often virtualized resources are provided as a service over the Internet. Users need not have knowledge of, expertise in, or control over the technology infrastructure ‘in the cloud’ that supports them.”

On the other side of the debate are those who have spent years building IT infrastructures, whose concerns must be addressed since reliability, security, SLAs, and interoperability will determine the success of the various cloud computing models within the enterprise.

Internal Clouds

Enterprises that want to take full advantage of existing IT infrastructure investments often start with an internal cloud. An internal cloud applies the concepts of cloud computing (on-demand resources, pay-as-you-go pricing, and the appearance of infinite scalability) to resources wholly owned by the enterprise consuming the service. By enhancing existing infrastructure with cloud computing capabilities, the internal cloud can provide economies of scale and the ability to handle both existing and new web-based applications, while providing security and regulatory compliance.

External Clouds

Today, well-known clouds are typically associated with an off-premises, hosted model. These external, or public, clouds involve IT resources and services sold with cloud computing qualities, such as self-service, pay-as-you-go billing, on-demand provisioning, and the appearance of infinite scalability. They are accessed through web browsers or through APIs and offer nearly unlimited capacity on demand at pay-per-use pricing, but with limited customer control. Some external cloud types include:

- Software as a service, in which application services are delivered over the network on a subscription basis: for example, Salesforce.com
- Platform as a service, in which software development frameworks and components are delivered over the network on a pay-as-you-go basis: for example, Google Apps
- Infrastructure as a service, in which computing, network, and storage services are delivered over the network on a pay-as-you-go basis: for example, Amazon EC2

Although many third-party vendors of external clouds are experiencing increasing demand for their offerings, businesses report some concern regarding infrastructures hosted by third parties for mission-critical, highly confidential applications. Businesses and enterprises require security, service-level guarantees, and compliance control, but with public clouds, the vendors are in control of these requisite capabilities.

Evolution Toward Private Clouds

In contrast to external clouds, private clouds—enterprise IT infrastructure services, managed by the business, with cloud computing qualities such as self-service, pay-as-you-go charge-back, on-demand provisioning, and the appearance of infinite scalability—provide a critical benefit: trust.

Initially, most private clouds will be made up almost entirely of internal resources. A private cloud can combine both external and internal cloud resources to meet the needs of an application system, and that combination, which is totally under enterprise control using unified management, can change moment by moment.

With a private cloud, enterprises can run processes internally and externally, having established the private cloud as the control point for workloads. With control through a unified management tool and a user-centric view, the private cloud thus enables IT to make the best decisions about whether to use internal or external resources, or both, and allows that decision to be made on a real-time basis to meet user service needs.

Private clouds should not be confused with hybrid clouds. A hybrid cloud uses both external (under the control of a vendor) and internal (under the control of the enterprise) cloud capabilities to meet the needs of an application system. A private cloud lets the enterprise choose, and control the use of, both types of resources

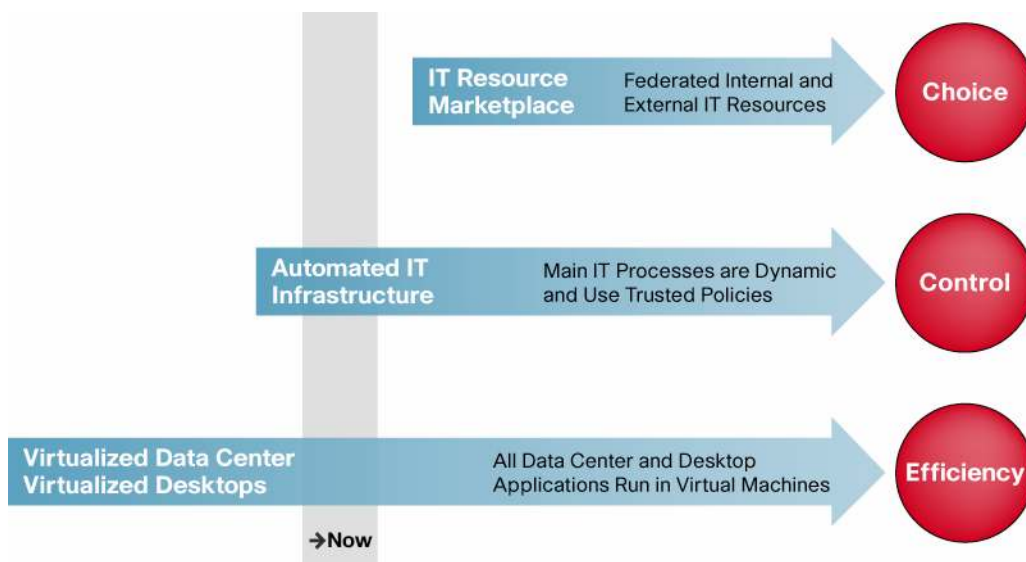
Thus, private clouds eliminate the “rewrite everything” effect of “own nothing” public cloud computing. They also provide the trusted control that enterprises are seeking from internal clouds, including the capability to change the mix of cloud services to be consumed, as needed.

Private cloud computing has the potential to eliminate enterprise silos while reducing costs through economies of scale and allows businesses to exploit the new economics, with faster deployment of new business applications and capacity for greater agility. Most important, private cloud computing offers a clear and logical pathway that preserves existing investment in applications and information, while each step delivers immediate value and builds a foundation for the next.

The Movement Toward Private Cloud Computing

Today the movement toward private clouds has started, with the virtualized data center and virtualized desktops. It involves three stages, shown in Figure 1.

Figure 1. The Enterprise Movement Toward a Multi-Tenant Cloud Solution



Stage 1: Consolidation and Virtualization



The movement toward cloud computing began for the enterprise with data center virtualization and consolidation of server, storage, and network resources to reduce redundancy and wasted space and equipment with measured planning of both architecture (including facilities allocation and design) and process. According to John Chambers, Cisco CEO, “Platform-based virtualization with a unified fabric will drive the globalization of IT resources—clients, servers, data centers, networks, and storage.”

Virtualization technologies enable the abstraction and aggregation of all data center resources in order to turn them into a unified logical resource that can be shared by all application loads. Virtualization decouples the physical IT infrastructure from the applications and services being hosted, allowing greater efficiency and flexibility, with any effect on system administration productivity handled by tools and processes.

Consolidation is a critical application of virtualization, enabling IT departments to regain control of distributed resources by creating shared pools of standardized resources that can be rationalized and centrally managed. Many IT departments already are consolidating underutilized computing resources by running multiple applications on a single physical server with virtualization technology from VMware. The percentage of new servers running virtualization as the primary boot option will approach 90 percent by 2012, according to analysts.

Many customers are using server virtualization not only for server consolidation, but also to improve flexibility, speed up service provisioning, and reduce planned and unplanned downtime. Enterprise IT also has undertaken projects to regain control of expanding storage resources by consolidating into scalable SANs.

And server virtualization is just the start; virtualization really needs to extend to network and storage. Virtualization of servers, storage, and networks will enable the mobility of applications and data not only across servers and storage

arrays in the same data center—which customers are already implementing in production—but also across data centers and networks.

After customers start virtualizing and aggregating resources into a single pool, they can operate more efficiently, flexibly, and reliably. Although the most immediate motivations for server, storage, and network virtualization are improved resource utilization and lower capital and operating costs, the ultimate goal is the use of the abstraction between applications and infrastructure to manage IT as a service.

Enterprises that want to begin moving toward cloud computing can start with internal, self-built clouds, utilizing virtualization for consolidation and automation.

Stage 2: Automation and Optimized Virtualization



Control → Hardware and Provisioning Freedom

In the second stage, enterprises need to move from managing underlying infrastructure to managing service levels based on what makes sense for the user of applications. For example, the customer may want to manage factors such as the minimum tolerable application latency or the availability level of an application. Enterprises also must implement automation for central IT and self-service for end users, thus extricating IT from the business of repetitive management procedures and enabling end users to get what they need quickly.

In this stage, virtualization optimizes IT resources and increases IT agility, thus speeding time-to-market for services. The IT infrastructure undergoes a transformation in which it becomes automated and critical IT processes are dynamic and controlled by trusted policies. Through automation, data centers systematically remove manual labor requirements for the run-time operation of the data center.

Provisioning automation is the best-known and most-often implemented form of automation. Sophisticated automation can help reduce operating expenses through:

- On-demand reallocation of computing resources
- Run-time response to capacity demands
- Trouble-ticket response automation (or elimination of trouble tickets for most automated response scenarios)
- Integrated system management and measurement

The ultimate goal, however, is alignment of operations with business needs through creation of a utility computing model.

To create a cloud service, self-service and metering (feedback about the cost of the resources allocated) are offered in addition to automation. With self-service and metering, the computing model resembles a utility. The private cloud then is a technical strategy to turn computing power into utility computing, with the data and costs controlled and managed by the enterprise.

Self-service and metering are breakthrough capabilities for end users and business units, facilitating management and extension of the user experience. Now there is no intermediary between the consumer of a resource and the processes for the acquisition and allocation of resources for core businesses requirements and initiatives. Since the consumer initiates the service requests, now IT is an on-demand service rather than a gatekeeper. With the transition to an on-demand service, the cost structure is dramatically reduced, since the user uses and pays for only what is needed at a specific moment.

The business benefit of this change in cost structure is significant. For example, with the hardware and provisioning freedom that comes with a private cloud, a major pharmaceutical company can perform multiple drug trials that cost

far less in computing power than a single drug trial had cost previously. As a result, the company can now rethink the way that it conducts its research and product development, dramatically improving time-to-market.

While self-service and metering are breakthrough private cloud capabilities for end users and business units, maintaining service delivery in a fully virtualized multi-tenancy environment and providing security, especially for information and services leaving the data center environment, are essential enterprise requirements for IT administrators

With a private cloud utility model enabling these needs and requirements, enterprises can scale and expand by pooling IT resources in a single cloud operating system or management platform. They then can support anywhere from tens to thousands of applications and services and enable new architectures targeting very large-scale computing activities.

Stage 3: Federation



Choice → Open Cloud
Internetworking → Business Freedom

Creation of an open, competitive marketplace, in which IT capabilities in a utility model can be procured, allocated, and provisioned over the Internet on demand by the consumer, with self-service and metering, requires federation.

Federation—linking disparate cloud computing infrastructures with one another by connecting their individual management infrastructures—allows disparate cloud IT resources and capabilities—capacity, monitoring, and management—to be shared, much like power from a power grid. It also enables unified metering and billing, one-stop self-service provisioning, and the movement of application loads between clouds, since federation can occur across data center and organization boundaries, with cloud internetworking.

Cloud internetworking is the network technology enabling the linkage of disparate cloud systems in a way that accommodates the unique nature of cloud computing and the running of IT workloads.

Service Providers and Virtual Private Clouds

To go beyond organizational boundaries through cloud internetworking to reach a third-party, businesses will need service providers and virtual private cloud services. Ultimately, service providers will offer both public and virtual private cloud services on a secure infrastructure, and that will allow enterprises to include and consume those services as part of enterprise private clouds, without exposing content to the general public.

This virtual private cloud service provider model would:

- Enable choice through open intercloud standards and services
- Support federation across internal and external clouds
- Deliver cloud services with security, quality of service (QoS), and manageability
- Use standards for consolidated application and service management and billing

Thus, enterprises will be able to choose freely among the service providers, and service providers will be able to use other providers' infrastructures that allow federation to handle exceptional loads on their own offerings. With virtualization of the private cloud, Cisco will offer solutions that support the packaging of applications and data and the linkages between them into virtual containers. The virtual container, a standards-based representation of the virtual infrastructure and related policies, can be interpreted and implemented as a physical infrastructure.

Virtual private clouds will extend the trust benefit of the private cloud to include third-party components and service provider services.

Why Use Private Clouds?

The private cloud is a new style of computing in which corporate IT infrastructure is available as a ubiquitous, easily accessible, and reliable utility service. Business owners and application owners requesting a new business service can use the infrastructure as a standard service, without the need to understand the complexities of servers, storage, and networks. The private cloud brings the benefits of cloud computing under the control of corporate IT:

- Available on demand
- Faster provisioning of business services
- Reduced cost through economies of scale
- Flexibility and freedom of choice to run workload and applications in the most efficient and effective places
- Use-based, pay-as-you-go model
- Standardized, auditable service levels
- Capability to work with every application, both current and future, without the need to rewrite applications
- Secure and controlled by corporate IT

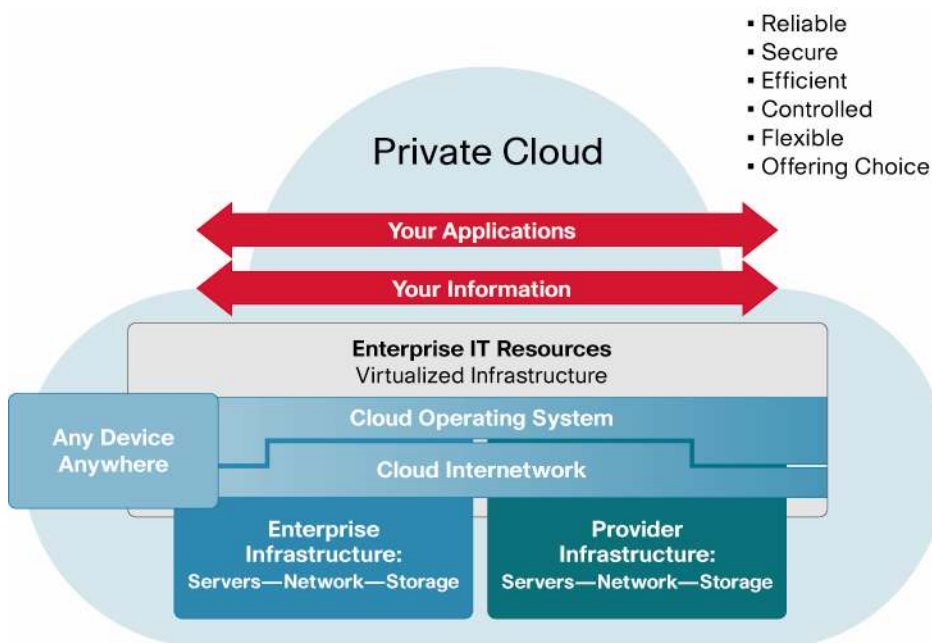
Private clouds serve many types of users in a variety of business settings: corporate and division offices, business partners, raw-material suppliers, resellers, distributors, and production and supply-chain entities.

Building the Private Cloud

While most customers have some understanding of the benefits of IT as a service and the private cloud model, many customers do not know how to build a private cloud.

Today, Cisco is building core elements of a virtualized and cohesive system, integrating desktop systems, mobile devices, data centers, servers, networks, and storage resources through open standards, in secure and trusted, multi-tenant enterprise solutions, starting with private clouds (Figure 2).


Figure 2. Building the Private Cloud



Cisco Virtualized Network Platform

Cloud computing and private clouds operate at the intersection of computing, networking, and virtualization, as shown in the Cisco Unified Computing System (Figure 3).

Figure 3. Cisco Unified Computing System



From Cisco's perspective, the network is the foundation for cloud internetworking, with the capability to deliver:

- Networking technologies to make the cloud enterprise-ready
- Network intelligence for workload portability
- Multimedia (video) delivery in the cloud
- Interoperability across clouds
- User experience with end-to-end service delivery across clouds

As a starting point, existing products in the Cisco Data Center 3.0 and Cisco Unified Computing System portfolios enable the enterprise to transition to the cloud.

The underlying tenets of the Cisco Data Center 3.0 and Cisco Unified Computing System portfolios are consolidation, virtualization, and automation: the same tenets that underpin cloud computing. These technologies are crucial enablers of cloud computing infrastructures.

With these technologies, and in an evolutionary manner, customers can create dynamic virtualized infrastructures that will take advantage of new private cloud solutions from Cisco during the coming year.

Server Consolidation

The first step toward reducing total cost of ownership (TCO) in the enterprise data center is server consolidation. Data centers today, with their siloed environments and high costs, are characterized by fewer than approximately 15 percent of production workloads running on virtual machines. As more customers try to reduce costs using server virtualization technologies such as VMware vSphere—the first cloud operating system to transform IT infrastructures into private clouds—the virtual machine will become the default application platform, with 60 to 80 percent of x86 applications running in a virtualized environment.

Transition to 10 Gigabit Ethernet

From a networking perspective, the increase in virtual machine density encourages a transition to 10 Gigabit Ethernet as the default mechanism for attaching servers. Multiple virtual machines on a single server can quickly overwhelm a single Gigabit Ethernet link, and multiple Gigabit Ethernet links can increase costs. The Cisco Nexus™ Family, including the Cisco Nexus 7000 Series Switches and Cisco Nexus 2000 Series Fabric Extenders, supports existing Gigabit Ethernet attached servers in today's data centers and enables future deployment of 10 Gigabit Ethernet, unified fabric, and virtual machine-aware networking (with Cisco VN-Link).

Cisco Nexus 1000V Series Switches provide operational consistency down to the individual virtual machine as well as policy portability, so network and security policy follows virtual machines as they move around the data center. The Cisco Nexus 1000V Series can be deployed wherever VMware ESX is currently running, without dependencies based on the server uplink speed or upstream access switch.

Consolidation of Server I/O

To continue to reduce TCO and to deploy virtual machines with technologies like VMware Distributed Resource Scheduling (DRS), all servers must have a consistent and ubiquitous set of network and storage capabilities. One of the simplest and most efficient ways to deliver these capabilities is to deploy a unified fabric. The shift to a unified fabric gives all servers, physical and virtual, access to the SAN, allowing more storage to be consolidated in the customer's SAN for greater efficiency and costs savings.

Unified Fabric

To consolidate server I/O, the server access layer must be adapted to support a unified fabric. Cisco Nexus Family products enable this support, since they support Fibre Channel over Ethernet (FCoE), Small Computer System Interface over IP (iSCSI), and Fibre Channel. On the network side, supporting FCoE is simply a matter of enabling FCoE features in the Cisco Nexus 5000 Series Switches and installing either the Fibre Channel or Fibre Channel and Data Center Bridging (DCB) uplink modules. Any attached Cisco Nexus 2000 Series Fabric Extenders inherit FCoE capability. The Cisco Nexus 7000 Series is FCoE capable; DCB-capable I/O modules (to provide reliable transport) will be available 2010.

Deployment of FCoE on the server side requires either a new converged network adapter (CNA) from a company such as Emulex or QLogic or, in the case of Intel interfaces, a new software driver for the 10 Gigabit Ethernet network adapters. FCoE is supported on VMware vSphere, and the Emulex, Intel, and QLogic interfaces are on the VMware ESX 3.5 hardware compatibility list (HCL), so this step encompasses both physical and virtual servers.

Scalable, Dynamic, Unified Fabric

With a scalable, dynamic, unified fabric, now any data center asset can access any other asset. From a storage perspective, the Cisco data center fabric will support FCoE and iSCSI server access and FCoE, Fibre Channel, and iSCSI attached targets. DCB will expand from the access layer into the aggregation and core layers, simplifying access to the Fibre Channel SAN. FCoE interfaces to the existing Cisco MDS 9000 Family of director-class Fibre Channel switches will simplify access to the existing SAN.

Unified Computing

The unified fabric now enables a fully virtualized data center composed of pools of computing, network, and storage resources. Services such as security and Layer 4 through 7 processing (for example, load balancing) are virtualized and can be implemented whenever they are needed, along with automated management and provisioning capabilities.

In addition, Cisco VN-Link network services are available in VMware vSphere environments with the Cisco Nexus 1000V Series Switches and through the Cisco Unified Computing System. Cisco VN-Link storage services are available as part of the intelligent fabric delivered by the Cisco MDS 9000 Family of multilayer SAN switches.

Cisco VN-Link bridges the server, storage, and network management domains, so changes in one environment are communicated to the others. Cisco VN-Link helps enable new capabilities and features, simplifies management and operations, and allows scaling for server virtualization solutions. Specific benefits include:

- Real-time policy-based configuration
- Mobile security and network policy, moving policy with the virtual machine during VMware VMotion live migration for persistent network, security, and storage compliance
- Nondisruptive management model, aligning management and operations environment for virtual machines and physical server connectivity in the data center

The data center up to this point in its evolution has been confined to the enterprise. The Cisco Unified Computing System facilitates movement outside the enterprise for services.

The Cisco Unified Computing System bridges the silos in the data center, enabling better utilization of IT infrastructure in a fully virtualized environment, and creates a unified architecture using industry-standard technologies that provide interoperability and investment protection.

The architecture unites computing, network, storage access, and virtualization resources into a scalable, modular design that is managed as a single energy-efficient system through Cisco UCS Manager, management software embedded in the Cisco UCS 6100 Series Fabric Interconnects, a family of line-rate, low-latency, lossless, 10-Gbps Cisco DCB and FCoE interconnect switches that consolidate I/O within the system.

Cisco UCS Manager provides an intuitive GUI, a command-line interface (CLI), and a robust API for managing all system configuration and operation. Cisco UCS Manager helps increase IT staff productivity, enabling IT managers of storage, networking, computing, and applications to collaborate on defining service profiles for applications. Service profiles help automate provisioning, allowing data center managers to provision applications in minutes instead of days.

The Cisco Unified Computing System provides support for a unified fabric over a low-latency, lossless, 10-Gbps Ethernet foundation. This network foundation consolidates today's separate networks: LANs, SANs, and high-performance computing networks. Network consolidation lowers costs by reducing the number of network adapters, switches, and cables and thus decreasing power and cooling requirements.

The Cisco Unified Computing System also allows consolidated access to both SANs and network attached storage (NAS). With its unified fabric, the Cisco Unified Computing System can access storage over Ethernet, Fibre Channel, FCoE, and iSCSI, providing customers with choices and investment protection. In addition, IT staff can preassign storage access policies for system connectivity to storage resources, simplifying storage connectivity and management.

Cisco UCS 2100 Series Fabric Extenders bring unified fabric to the Cisco UCS 5100 Series Blade Server Chassis, providing up to four 10-Gbps connections each between blade servers and the fabric interconnect, simplifying diagnostics, cabling, and management.

The new Cisco UCS B-Series Blade Servers, based on the future Intel Nehalem processor families (the next-generation Intel Xeon processor), offer patented extended memory technology to support applications with large data sets and allow significantly more virtual machines per server. Cisco UCS B-Series Blade Servers support up to eight blade servers and up to two fabric extenders in a six-rack-unit (6RU) enclosure, without the need for additional management modules. Each blade server uses network adapters for access to the unified fabric.

Cisco UCS network adapters include adapters optimized for virtualization, compatibility with existing driver stacks, and efficient, high-performance Ethernet. With integrated management and "wire-once" unified fabric with the industry-standard computing platform, the Cisco Unified Computing System optimizes virtualization, provides

dynamic resource provisioning for increased business agility, and reduces total overall data center costs, with up to 20 percent reduction in capital expenditures (CapEx) and up to 30 percent reduction in operating expenses (OpEx).

Offering a new style of dynamic IT, Cisco Unified Computing System extends customers' virtualized data centers and creates a foundation for private clouds that federate with compatible service providers. With the virtualized environment defined by a dynamic, scalable data center fabric, a workload really can run anywhere; the resources needed to support a workload can come even from an outside service provider in a cloud computing model.

Business Value of Cisco Unified Computing System

- Simplification through predesigned virtualization architecture
- Reduced design-build-deploy-operate lifecycle
- Extended lifecycle for capital assets through hardware virtualization
- Predictability of investment cycles for IT and facilities
- Scalability (both up and down) as business needs require
- Improved business metrics (return on net assets [RONA], CapEx, and OpEx)

Conclusion

Customers today can take the first step with Cisco toward cloud computing and private clouds, as the enterprise IT infrastructure evolves toward a cost-effective, ubiquitous, easily accessible, reliable and efficient utility service, with unified management for ease of delivery. Private clouds keep the benefits of cloud computing under the control of enterprise IT, while increasing business opportunities for service providers will introduce choice in cloud services. Ultimately, service providers will offer both public and private cloud services in the form of virtual private cloud services that are consumed and controlled within enterprise private clouds.

The introduction of publicly shared core services—much like domain name system (DNS) and peering services—into carrier and service provider networks will enable a more loosely coupled relationship between the customer and the cloud provider. Enterprises will be able to choose among service providers, and federated service providers will be able to share service loads. This looser relationship will increase the elasticity of the cloud market and create a single, public, open cloud internetwork: the intercloud.

With federation and application portability as the cornerstones of the intercloud, businesses will be able to achieve business process freedom and innovate, and users will experience choice and faster, better services.

For More Information

For more information about the strategies and services that Cisco offers to help enterprise IT departments and service providers evolve toward a private cloud model and cloud computing with efficiency, control, and flexibility in next-generation data center architectures, please visit <http://www.cisco.com/go/datacenter>.

To learn more about Cisco Data Center 3.0, Cisco Unified Computing System, and VMware vSphere, go to:

- [Cisco Data Center 3.0](#)
- [Cisco Unified Computing System](#)
- [VMware vSphere](#)



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, CCSI, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCI, CCNA, CCNP, CCS, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, 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. (0903R)