

## Citrix CloudPlatform Version 3.0.4 on the Cisco UCS B200 M3 Blade Server



---

# Contents

<b>Overview.....</b>	<b>3</b>
What You Will Learn.....	3
Audience.....	3
Objectives.....	3
The Business Challenge.....	3
<b>Solution Components .....</b>	<b>3</b>
Cisco Components .....	3
Citrix CloudPlatform.....	5
<b>Architectural Design.....</b>	<b>5</b>
Design Fundamentals.....	5
CloudPlatform Design Fundamentals.....	5
Types of Cloud Workloads.....	6
CloudPlatform Supports Both Workload Types .....	8
Cisco Unified Computing System Design Fundamentals .....	10
<b>Solution Validation .....</b>	<b>12</b>
Configuration Topology for Scalability .....	12
Cabling Information.....	14
Cisco Unified Computing System Configuration.....	17
Citrix CloudPlatform Configuration.....	17
Running the Management Server as an Application.....	17
Management Server Cluster Backup and Replication .....	18
LAN Configuration .....	19
SAN Configuration .....	20
Boot from SAN.....	20
Configuring CloudPlatform for Use with Cisco UCS .....	21
<b>Scalability References.....</b>	<b>26</b>
VDI Workload Scaling Example.....	26
Three-tier Application (Open-Source) Scaling Example .....	32
<b>Conclusion .....</b>	<b>39</b>
<b>For More Information.....</b>	<b>39</b>

---

## Overview

### What You Will Learn

This reference architecture guide suggests best practices for configuring, deploying, and scaling Citrix CloudPlatform™ Version 3.0.4 on the Cisco Unified Computing System™ (Cisco UCS®) B200 M3 Blade Server connected to a storage array with Cisco Nexus® 5000 Series Switches. Second-generation Cisco® UCS hardware and software are used. Best practice recommendations and sizing guidelines are provided for large-scale customer deployments of Citrix CloudPlatform running two workload scenarios on Cisco UCS.

### Audience

This document is designed to assist solution architects, sales engineers, field engineers, and consultants with evaluation, planning, design, and deployment of Citrix CloudPlatform Version 3.0.4 on Cisco UCS. The reader should have an architectural understanding of Cisco UCS, Cisco Nexus 5000 Series Switches, Citrix CloudPlatform, shared storage, and related software.

### Objectives

The main objectives of this guide include articulation of architectural design considerations required to successfully deploy Citrix CloudPlatform Version 3.0.4 on the included reference Cisco UCS architecture with shared storage.

The reference architecture is not intended to be a definitive configuration but rather to suggest best practices and provide a starting point for CloudPlatform deployments. Additionally, we've included scaling examples, which can help you determine required resources for a given workload.

### The Business Challenge

Companies everywhere are looking for effective strategies to harness the benefits of cloud computing without disrupting current business models. As the industry ventures into the cloud computing era, service providers and consumer oriented companies are seeking more efficient and differentiated cloud solutions to reduce the total cost of ownership of IT, attract and retain customers, and increase market share. They need open and flexible cloud solutions that free them from vendor lock-in so they can take full advantage of existing investments and choose the best possible components for their clouds. They need access to source code and open APIs to innovate and build value-added services, all while still having enterprise-class support and services. Their customers want to choose the architecture and hypervisor that's right for them. At the same time, enterprises are looking to the cloud to enable more agile, elastic, on-demand IT services. In both cases, they need the right solutions to build, scale, and manage cloud services.

This joint solution comprised of Citrix CloudPlatform running on Cisco UCS provides an agile infrastructure with a proven enterprise cloud infrastructure manager. Citrix CloudPlatform is a broad solution that includes a commercially certified and packaged Apache CloudStack product. Cisco UCS is an extremely scalable, automated, and programmable infrastructure for cloud deployment. The CloudPlatform user interface is integrated with Cisco UCS Manager to provide seamless cloud infrastructure management.

## Solution Components

### Cisco Components

The following components are required to deploy the Unified Access design:

- Cisco UCS

- Cisco Nexus 5500 Series Switch

#### Cisco Unified Computing System

The Cisco UCS is a next-generation approach to blade and rack server computing. It is an innovative data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. Managed as a single system, whether it has one server or 160 servers with thousands of virtual machines, the Cisco UCS decouples scale from complexity. It accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and nonvirtualized systems.

The Cisco UCS consists of the following components:

- [Cisco UCS 6200 Series Fabric Interconnects](#) are a family of line-rate, low-latency, lossless, 10-Gigabit Ethernet and Fibre Channel over Ethernet interconnect switches providing the management and communications backbone for the Cisco UCS. Cisco UCS supports Cisco Data Center Virtual Machine Fabric Extender ([VM-FEX](#)) technology.
- [Cisco UCS 5100 Series Blade Server Chassis](#) supports up to eight blade servers and up to two fabric extenders in a six rack unit (RU) enclosure.
- [Cisco UCS B-Series Blade Servers](#) are Intel based and increase performance, efficiency, versatility, and productivity.
- [Cisco UCS Adapters](#) are available for the Cisco UCS wire-once architecture, offering a range of options to converge the fabric, optimize virtualization, and simplify management. Cisco adapters support VM-FEX technology.
- [Cisco UCS Manager](#) provides unified, embedded management of all software and hardware components in the Cisco UCS.

For more information, see: <http://www.cisco.com/en/US/products/ps10265/index.html>

#### Cisco Nexus 5500 Series Switch

The Cisco Nexus 5000 Series Switch is designed for data center environments with cut-through technology that enables consistent, low-latency Ethernet solutions with front-to-back or back-to-front cooling and data ports in the rear, bringing switching into close proximity with servers and making cable runs short and simple. The switch series is highly serviceable, with redundant, hot-pluggable power supplies and fan modules. It uses data-center-class Cisco NX-OS Software for high reliability and ease of management.

The switch extends the industry-leading versatility of the Cisco Nexus 5000 Series purpose-built 10 Gigabit Ethernet data-center-class switches and provides innovative advances toward higher density, lower latency, and multilayer services. The Cisco Nexus 5500 Series switch is well suited for enterprise-class data center server access layer deployments across a diverse set of physical, virtual, storage-access, and high-performance computing (HPC) data center environments.

The Cisco Nexus 5548UP is a 1RU 10 Gigabit Ethernet (10 GE), Fibre Channel (FC), and Fibre Channel over Ethernet (FCoE) switch offering up to 960 Gbps of throughput and up to 48 ports. The switch has 32 unified ports

---

and one expansion slot supporting modules with 10 Gigabit Ethernet and FCoE ports or connectivity to Fibre Channel SANs with 8/4/2/1 Gbps Fibre Channel switch ports, or both.

For more information, see: <http://www.cisco.com/en/US/products/ps9670/index.html>

### Citrix CloudPlatform

Citrix CloudPlatform, powered by Apache CloudStack, is the only open and flexible software platform that pools compute, storage, and networking resources to build highly scalable and reliable public, private, and hybrid infrastructure-as-a-service (IaaS) clouds. These cloud solutions support legacy enterprise and cloud-era workloads designed specifically for the needs of both enterprises and service providers.

This reference architecture featuring CloudPlatform on the Cisco UCS includes:

- CloudPlatform Management Server 3.0.4 running on CentOS 6.2
- MySQL database Version 5.1
- vCenter 5.0 (optional)
- ESXi 5.0 hypervisor installed onto the Cisco UCS blades
- Network File System (NFS) server for secondary storage (Image and Backup Library)
- NFS server for primary storage (Fibre Channel attached storage)
- Networking preconfigured on the Cisco Fabric Interconnects.

## Architectural Design

### Design Fundamentals

"On-demand" is a key requirement for cloud services, and Citrix CloudPlatform provides a methodology for delivering cloud services on demand.

The challenge of meeting cloud workload demand is typically related to scaling the physical infrastructure supporting the cloud. Our reference architecture aims to provide a design for CloudPlatform that offers maximum scalability with minimal effort.

Using Cisco's Unified Fabric for the networking component of the cloud infrastructure helps to ensure that once cables are run, they do not need to be rerouted as cloud workload requirements change. The Cisco UCS helps to ensure that servers can be delivered quickly and in an automated fashion.

### CloudPlatform Design Fundamentals

Citrix CloudPlatform is an open source software platform that pools data center resources to build public, private, and hybrid infrastructure-as-a-service (IaaS) clouds. CloudPlatform abstracts the network, storage, and compute nodes that make up a data center and enables them to be delivered as a simple-to-manage, scalable cloud infrastructure. These nodes or components of a cloud can vary greatly from data center to data center and cloud to cloud because they are defined by the unique workloads or applications that they support. With so many options for servers, hypervisors, storage, and networking, it is imperative that cloud operators design with a specific application in mind to help ensure that the infrastructure meets the scalability and reliability requirements of the application.

## Workload-Driven Deployment Process

Figure 1 illustrates the steps a cloud operator typically follows to determine the appropriate deployment architecture for CloudPlatform.

**Figure 1.** Determining Appropriate Deployment Architecture for Citrix CloudPlatform



## Types of Cloud Workloads

Two distinct types of application workloads have emerged in cloud operator's data centers:

- The first type is a traditional enterprise workload. The majority of existing enterprise applications falls into this category. They include applications developed by leading enterprise vendors and have traditional functions such as databases, web servers, and application servers for example. These applications are typically built to run on a single server or on a cluster of front-end and application server nodes backed by a database. Traditional workloads typically rely on technologies such as enterprise middleware clusters and vertically scaled databases.
- Citrix commonly refers to the second type as a "cloud-era workload." Many internet focused companies such have determined that traditional enterprise infrastructure was too inflexible to serve the load generated by millions of users. These Internet companies pioneered a new style of application architecture that does not rely on traditional server clusters but instead on a large number of loosely coupled computing and storage nodes. Applications developed this way often utilize technologies such as database sharding, no-SQL, and geographic load balancing.

There are three fundamental differences between traditional workloads and cloud-era workloads:

- **Scale:** The first difference is scale. Traditional enterprise applications serve tens of thousands of users and hundreds of sessions. Driven by the growth of Internet and mobile devices, Internet applications serve tens of millions of users. The orders of magnitude difference in scale translates to significant difference in

demand for computing infrastructure. As a result, the need to reduce cost and improve efficiency becomes paramount.

- **Reliability:** The difference in scale has an important side effect. Enterprise applications can be designed to run on reliable hardware. Application developers do not expect the underlying enterprise-grade server or storage cluster to fail during their normal course of operation. Sophisticated backup and disaster recovery procedures can be set up to handle the unlikely scenario of hardware failure.. “Cloud-era” applications on the other hand take into account that as hardware scales beyond traditional enterprise and into web-scale, the points of failure become higher. Thusly these applications assume the possibility of failure is higher, and are designed to withstand failures.
- **Agility:** Cloud-era workloads inherently make efficient use of infrastructure resources and can easily add in new resources or lose resources without much impact. Traditional workloads also require agile infrastructure in which resources can be added or removed, or move between workloads to counterbalance the demand.

Table 1 summarizes the differences between traditional and cloud-era requirements.

**Table 1.** Traditional Versus Cloud-Era Workload Requirements

Requirement	Traditional Workload	Cloud-Era Workload
Scale	Tens of thousands of users	Millions of users
Reliability	99.999 % uptime	Assumes failure
Infrastructure	Scale Up	Scale Out
Applications	Database, Back-office, Traditional LAMP/WAMP	Web content, Web apps, social media

Cloud-era workloads assume that the underlying infrastructure might fail or that administrators will induce human error. Instead of implementing disaster recovery as an afterthought, multisite geographic failover must be designed into the application. The application is expected to treat servers and storage as “ephemeral resources,” a term that means resources can be used while they are available but they may become unavailable after a short period of use.

Table 2 lists common application types or use cases for traditional and cloud-era workloads.

**Table 2.** Common Cloud Workloads

Traditional Workload Candidates	
<b>Communications and Productivity</b>	Email servers, Web-based collaborations□
<b>CRM / ERP / Database</b>	Enterprise databases/business operations□
<b>Desktop</b>	Desktop-based computing, desktop service and support applications, and desktop management applications
<b>Development and Test</b>	Software development, test processes, and image management
<b>Disaster Recovery</b>	Onsite/offsite backup and recovery, live failover, cloud bursting for scale
<b>Web Service</b>	Static and dynamic Web content, streaming media, RSS, mashups and SMS
<b>Web Applications</b>	Web service-enabled applications, e-commerce, e-business, Java application servers
<b>IaaS, PaaS</b>	Infrastructure as a service, Platform as a service
Cloud-Era Workload Candidates	
<b>Web Service</b>	Static and dynamic Web content, streaming media, RSS, mashups and SMS
<b>Web Applications</b>	Web service-enabled applications, e-commerce, e-business, Java application servers
<b>Rich Internet Applications</b>	Videos, online gaming, and mobile apps
<b>Disaster Recovery</b>	Onsite/offsite backup and recovery, live failover, cloud bursting for scale
<b>HPC</b>	Engineering design and analysis, scientific applications, high-performance computing
<b>Collaboration and Social Media</b>	Web 2.0 applications for online sharing and collaboration (blog, CMS, File Share, Wiki, IM)
<b>Batch Processing</b>	Predictive usage for processing large workloads, such as data mining, warehousing, analytics, business intelligence
<b>Development and Test</b>	Software development, test processes, and image management

## CloudPlatform Supports Both Workload Types

Citrix CloudPlatform natively supports both traditional enterprise and cloud-era workloads. While cloud-era workloads represent an application architecture that will likely become more common in the future, the majority of applications are written as enterprise-style workloads. With CloudPlatform, a cloud operator can design for one style of workload and add support for the other style later. Or a cloud operator can design for supporting both styles of workload from the beginning.

The ability to support both styles of workload lies in CloudPlatform's architectural flexibility. Cloud operators can, for example, configure multiple availability zones using different hypervisor, storage, and networking capabilities required to support different types of workloads to meet security, compliance, and scalability needs of multiple cloud initiatives.

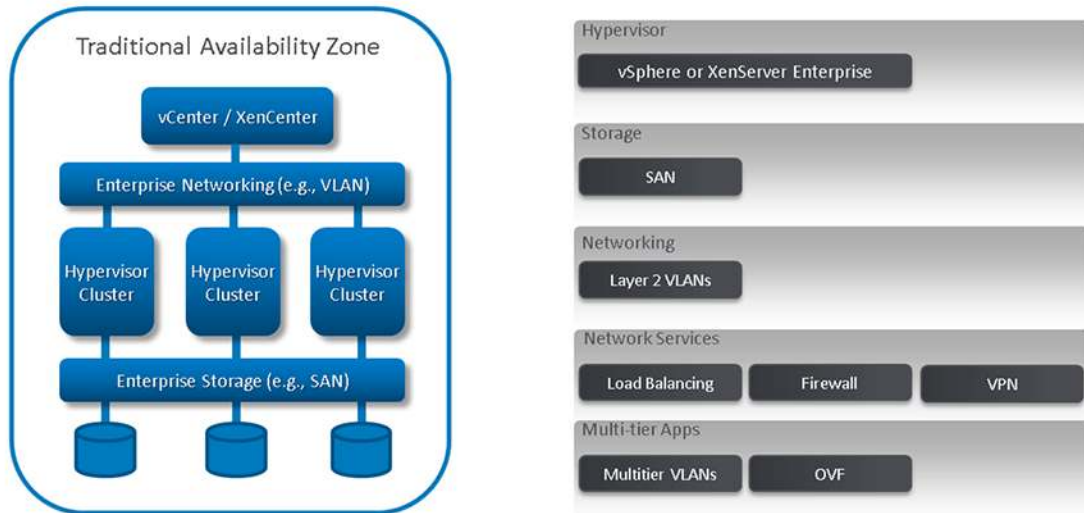
### Traditional Workload

Figure 2 illustrates how a CloudPlatform traditional Availability Zone can be constructed to support a traditional enterprise-style workload.

Traditional workloads in the cloud are typically designed with a requirement for high availability and fault tolerance and use common components of an enterprise data center to meet those needs. This starts with an enterprise-grade hypervisor, such as VMware vSphere or Citrix XenServer, that supports live migration of virtual machines and storage and has built-in high availability. Storage of virtual machine images uses high-performance storage area network (SAN) devices. Traditional physical network infrastructure like firewalls and Layer 2 switching are used and virtual LANs (VLANs) are designed to isolate traffic between servers and tenants. Virtual private network (VPN) tunneling provides secure remote access and site-to-site access through existing network edge devices. Applications are packaged using industry-standard Open Virtualization Format (OVF) files.



**Figure 2.** Constructing a CloudPlatform Availability Zone for Traditional Enterprise Workload

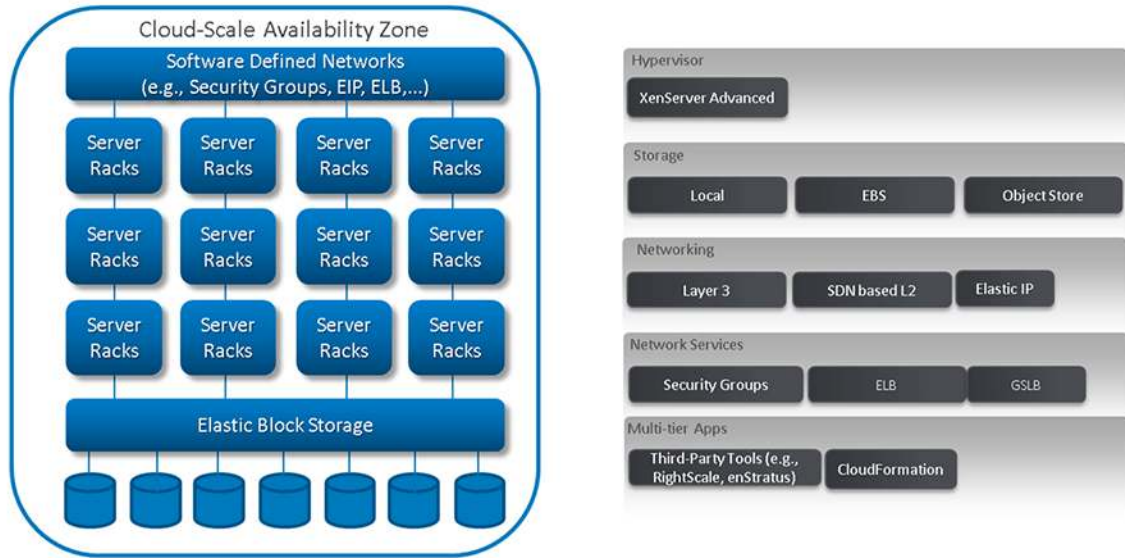


#### Cloud-Era Workload

Figure 3 illustrates how a CloudPlatform cloud-era Availability Zone can be constructed to support cloud-era workloads.

The desire for cost-savings can easily offset the need for features in designing for a Cloud-era workload, making open source and commodity components such as XenServer and kernel-based virtual machine (KVM) a more attractive option. In this workload type, VM images are stored in elastic block store (EBS) volumes, and object stores can be used to store data that must persist through Availability Zone failures. Because of VLAN scalability limitations, software-defined networks are becoming necessary in cloud-era Availability Zones. CloudPlatform meets this need by supporting Security Groups in Layer 3 networking. Elastic Load Balancing (ELB) or Global Server Load Balancing (GSLB) is used to redirect user traffic to servers in multiple Availability Zones. Third-party tools developed for Amazon Web Services to manage applications in this type of environment are readily available and have tested and proven integrations with CloudPlatform.

**Figure 3.** Figure 3. Constructing CloudPlatform Availability Zone for Cloud-Era Workload




## Cisco Unified Computing System Design Fundamentals

### Infrastructure Advantages

In contrast to traditional blade architectures, Cisco's UCS was built after the advent of virtualization and takes virtualization into account within its own design. The management paradigm of Cisco UCS is very similar to that of a virtualized infrastructure. Figure 4 shows the logical similarities between the Cisco UCS and virtual infrastructures.

**Figure 4.** Comparing Cisco UCS Physical Infrastructure with Virtual Infrastructures



Many of the concepts that make virtualization the ideal platform for delivering cloud also apply to UCS, thus making UCS also ideal for cloud deployments. Such common concepts are:

- Single point of infrastructure management, which enables complete control
- Policy-based infrastructure management
- Using templates to capture desired state and subsequently expedite deployment
- Programmatic control of the entire infrastructure

There are three key areas, inherent in the design, where Cisco UCS excels in comparison to traditional architectures, helping to simplify the design and maintenance of cloud solutions.

#### Scalability

Because Cisco UCS management is removed from the blade chassis and centralized in the network, horizontal scalability is simplified. Any blade chassis attached to the Cisco UCS management switches (the Fabric Interconnects) is easily brought under management and blades. Blades can be deployed and managed using a management methodology based on service profiles.

#### Flexibility

Every configurable parameter on a Cisco UCS blade, including BIOS settings, RAID configuration, network interface configuration, networking configuration, etc., is captured within a Cisco UCS service profile. The service profile is the complete definition of a server. Service profiles are decoupled from the blade itself: the service profile can be applied to any blade connected to the Cisco UCS switches. Profile configurations can be enforced through policies. This makes for easy administration of servers and flexible control over the physical infrastructure.

---

## Redundancy

All the components of a Cisco UCS blade server are redundant and designed for high availability. The Cisco UCS switches are designed to fail-over between each other. The blade chassis contains a redundant pair for I/O modules, which connect them to the switches, as well as four power supplies to ensure a constant supply. The ability to move service profiles between blades also helps to ensure that in the unlikely event of a catastrophic failure of a chassis, critical servers can be easily restored in an alternative chassis.

## UCS & Cloud Workload Types

These characteristics of Cisco UCS make it ideal for use in both traditional and cloud-era workloads. In both cases UCS provides the unique ability to rapidly configure compute, network and storage infrastructure while simplifying the management of hardware components. Legacy compute, network, and storage infrastructures add complexity to cloud deployments in that they have discrete management tools and lack a unified fabric to simplify interconnectivity. These fundamental features of UCS are the key in any cloud deployment.


## Solution Validation

This section details the configuration and tuning that was performed on the individual components to produce a complete, validated solution.

## Configuration Topology for Scalability

Figure 5 shows the reference configuration for the Citrix CloudPlatform on Cisco UCS. Table 3 lists the Cisco UCS bill of materials, and Table 4 lists the Cisco Nexus bill of materials.

**Figure 5.** Citrix CloudPlatform on Cisco UCS Configuration



**Table 3.** Cisco UCS Detailed Bill of Materials

Product	SKU	Description	Qty
<b>UCSB200M3 Blade Server</b>	UCSB-B200-M3	Cisco UCS B200 M3 Blade Server without CPU, memory, HDD, mLOM/mezz	8
	UCS-CPU-E5-2690	2.90-GHz E5-2690/135W 8C/20MB Cache/DDR3 1600MHz	16
	UCS-MR-1X162RY-A	16-GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v	128
	UCS-VIC-M82-8P	Cisco UCS VIC 1280 dual 40Gb capable Virtual Interface Card	8
	UCSB-MLOM-40G-01	Cisco UCS VIC 1240 modular LOM for M3 blade servers	8
	N20-BBLKD	Cisco UCS 2.5-inch HDD blanking panel	16
	UCSB-HS-01-EP	CPU Heat Sink for UCS B200 M3 and B420 M3	16
	N20-C6508	Cisco UCS 5108 Blade Svr AC Chassis/0 PSU/8 fans/0 fabric extender	2
	CAB-C19-CBN	Cabinet Jumper Power Cord, 250 VAC 16A, C20-C19 Connectors	8
	UCS-IOM-2208XP	Cisco UCS 2208XP I/O Module (8 External, 32 Internal 10Gb Ports)	4
	N01-UAC1	Single phase AC power module for Cisco UCS 5108	2
	N20-CAK	Access kit for 5108 Blade Chassis including Railkit, KVM dongle	2
	N20-FAN5	Fan module for Cisco UCS 5108	16
	N20-PAC5-2500W	2500W AC power supply unit for Cisco UCS 5108	8
	N20-FW010	Cisco UCS 5108 Blade Server Chassis FW package	2
<b>Chassis</b>			

Product	SKU	Description	Qty
Fabric Interconnects	UCS-FI-6248UP	Cisco UCS 6248UP 1RU Fabric Int/No PSU/32 UP/ 12p LIC	2
	UCS-ACC-6248UP	Cisco UCS 6248UP Chassis Accessory Kit	2
	UCS-PSU-6248UP-AC	Cisco UCS 6248UP Power Supply/100-240VAC	4
	N10-MGT010	Cisco UCS Manager v2.0	2
	UCS-FI-E16UP	Cisco UCS 6200 16-port Expansion module/16 UP/ 8p LIC	2
	CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	2
	UCS-LIC-10GE	Cisco UCS 6200 Series ONLY Fabric Int 1PORT 1/10GE/FC-port license	20
	UCS-FAN-6248UP	UCS 6248UP Fan Module	4
	UCS-FI-DL2	Cisco UCS 6248 Layer 2 Daughter Card	2

**Table 4.** Cisco Nexus Detailed Bill of Materials

Name	Part Number	Description	Qty
Cisco Nexus 5548UP Switch	N5K-C5548UP-OSM.P	Cisco Nexus 5548UP Storage Solutions Bundle, Full Stor Serv Lic, OSM	2
	N55-48PO-SSK9.P	Cisco Nexus 5500 Storage License, 48 Ports, OSM	2
	N55-DL2.P	Cisco Nexus 5548 Layer 2 Daughter Card	2
	N55-M-BLNK.P	Cisco Nexus 5500 Module Blank Cover	2
	N55-PAC-750W.P	Cisco Nexus 5500 PS, 750W, Front to Back Airflow(Port-Side Outlet)	4
	N5548P-FAN.P	Cisco Nexus 5548P and 5548UP Fan Module, Front to Back Airflow	4
	CAB-C13-C14-2M.P	Power Cord Jumper, C13-C14 Connectors, 2 Meter Length	4
	N5548-ACC-KIT.P	Cisco Nexus 5548 Chassis Accessory Kit	2
	N5KUK9-521N1.1.P	Cisco Nexus 5000 Base OS Software Rel 5.2(1)N1(1)	2
	DCNM-SAN-N5K-K9=.P	DCNM for SAN License for Nexus 5000	2

## Cabling Information

This section details the physical cabling of the Unified Computing platform.

Note this guide to cabling does not include integration of the management infrastructure.

### Cisco UCS Connectivity

Tables 5 and 6 list Ethernet cabling requirements for the Cisco Nexus 5548UP Switches. Tables 7 and 8 list Cisco UCS Fabric Interconnect cabling requirements. Tables 9 and 10 give cabling requirements for the Cisco UCS 2208XP I/O modules.

**Table 5.** Cisco Nexus 5548UP Ethernet Cabling for Switch A

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5548 A	Eth 1/9-10	10 GE	Cisco Nexus 5548 B	Eth 1/9 -10
	Eth 1/13-14	10 GE	Cisco UCS Fabric Interconnect A	Eth 1/13-14
	Eth 1/15-16	10 GE	Cisco UCS Fabric Interconnect B	Eth 1/13-14

**Table 6.** Cisco Nexus 5548UP Ethernet Cabling for Switch B

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5548 B	Eth 1/9-10	10 GE	Cisco Nexus 5548 A	Eth 1/9 -10
	Eth 1/13-14	10 GE	Cisco UCS Fabric Interconnect A	Eth 1/13-14
	Eth 1/15-16	10 GE	Cisco UCS Fabric Interconnect B	Eth 1/13-14

**Table 7.** Cisco UCS Fabric Interconnect A Ethernet Cabling

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	Eth1/ 13-14	10 GE	Cisco Nexus 5548 A	Eth1/ 13-14
	Eth1/ 15-16	10 GE	Cisco Nexus 5548 B	Eth1/13-14
	Eth1/1-4	10 GE	Cisco UCS B-Series FEX A	Port 1-4
	Eth 2/1-4	8 Gbps (FC)	Cisco Nexus 5548UP A	Eth 1/29-32

**Table 8.** Cisco UCS Fabric Interconnect B Ethernet Cabling

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect B	Eth1/ 13-14	10 GE	Cisco Nexus 5548 A	Eth1/ 13-14
	Eth1/ 15-16	10 GE	Cisco Nexus 5548 B	Eth1/ 15-16
	Eth1/1-4	10 GE	Cisco UCS B-Series chassis FEX B	Port 1-4
	Eth/ 2/1-4	10 GE	Cisco Nexus 5548UP B	Eth 29-32

**Table 9.** Cisco UCS 2208XP I/O Module for FEX A

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS-IOM FEX A	Eth1/1	10 GE	Cisco UCS Fabric Interconnect A	Eth1/1
	Eth1/2	10 GE	Cisco UCS Fabric Interconnect A	Eth1/2
	Eth1/3	10 GE	Cisco UCS Fabric Interconnect A	Eth1/3
	Eth 1/4	10 GE	Cisco UCS Fabric Interconnect A	Eth1/4

**Table 10.** Cisco UCS 2208XP I/O Module for FEX B

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS-IOM FEX B	Eth1/1	10 GE	Cisco UCS Fabric Interconnect B	Eth1/1
	Eth1/2	10 GE	Cisco UCS Fabric Interconnect B	Eth1/2
	Eth1/3	10 GE	Cisco UCS Fabric Interconnect B	Eth1/3
	Eth 1/4	10 GE	Cisco UCS Fabric interconnect B	Eth1/4

### Unified Access Layer Design Fibre Channel Connectivity

Tables 11 and 12 give requirements for the Cisco Nexus 5548UP Fibre Channeling cables. Tables 13 and 14 give requirements for the Cisco UCS Fabric Interconnect Fibre Channel cables. Figure 6 provides a color-coded wiring diagram.

**Table 11.** Cisco Nexus 5548UP Fibre Channel Cabling for Switch A

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5548 A	FC1/25	8 Gbps (FC)	Storage Array	Differs by Array
	FC1/26	8 Gbps (FC)	Storage Array	Differs by Array
	FC1/27	8 Gbps (FC)	Storage Array	Differs by Array
	FC1/28	8 Gbps (FC)	Storage Array	Differs by Array
	FC1/29-32	8 Gbps (FC)	Cisco UCS Fabric Interconnect A	Eth 2/1-4

**Table 12.** Cisco Nexus 5548UP Fibre Channel Cabling for Switch B

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5548 B	FC1/25	8 Gbps (FC)	Storage Array	Differs by Array
	FC1/26	8 Gbps (FC)	Storage Array	Differs by Array
	FC1/27	8 Gbps (FC)	Storage Array	Differs by Array
	FC1/28	8 Gbps (FC)	Storage Array	Differs by Array
	FC1/29-32	8 Gbps (FC)	Cisco UCS Fabric Interconnect B	Eth2/1-4

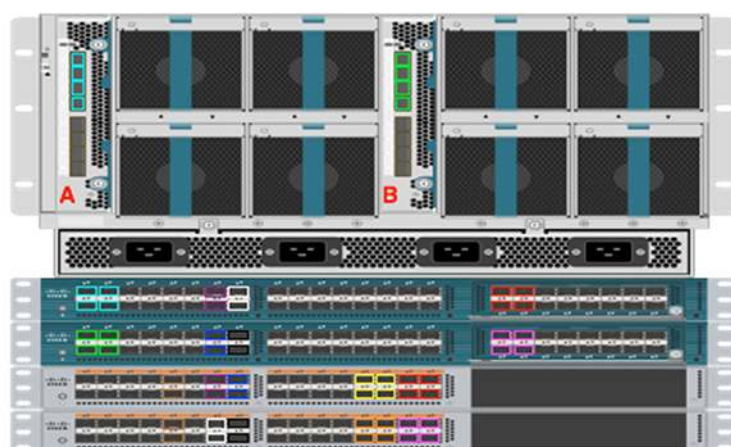
**Table 13.** Cisco UCS Fabric Interconnect A Fibre Channel Cabling

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	FC 2/1-4	8 Gbps (FC)	Cisco Nexus 5548UP Switch A	FC1/ 29-32

**Table 14.** Cisco UCS Fabric Interconnect B Fibre Channel Cabling

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect B	FC 2/1-4	8 Gbps (FC)	Cisco Nexus 5548UP Switch B	FC1/ 29-32

**Figure 6.** Color-Coded Wiring Map



Cisco UCS B-Series 5108 Chassis

Cisco UCS Fabric Interconnect A

Cisco UCS Fabric Interconnect B

Cisco Nexus 5548UP Switch A

Cisco Nexus 5548UP Switch B

Orange—To storage array (Side A)

Yellow—To storage array (Side B)



## Cisco Unified Computing System Configuration

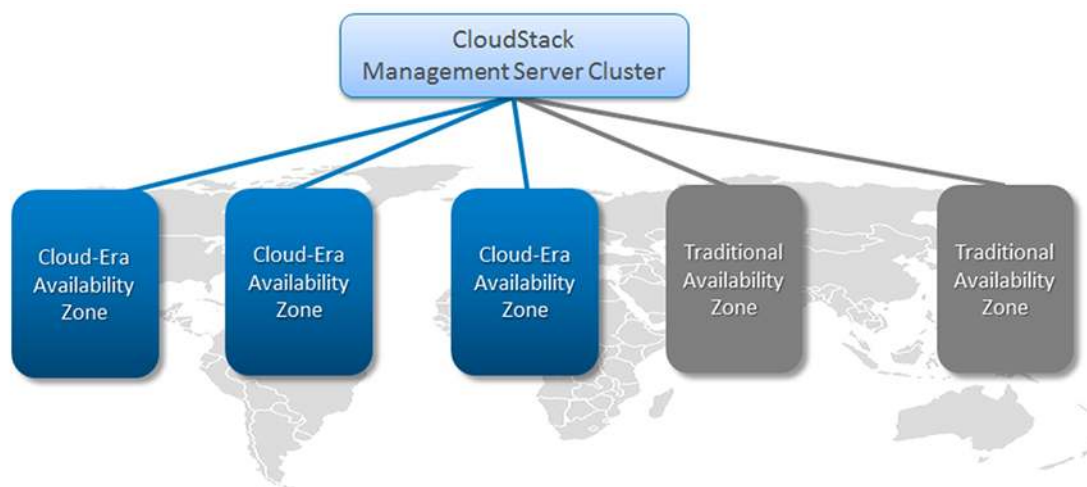
The basic preparation and configuration of the Cisco Unified Computing System are beyond the scope of this document. Detailed instructions can be found at:

- [The racking, power, and installation of the chassis](#)
- [Cisco UCS Command-line Interface \(CLI\) Configuration Guide](#)
- [Cisco UCS Manager GUI Configuration Guide](#)

## Citrix CloudPlatform Configuration

The management server deployment is not dependent on the underlying style of workload. A single management server cluster can manage multiple Availability Zones across multiple data centers, enabling cloud operators to create different Availability Zones to handle different workload types as needed. Figure 7 illustrates how a single cloud can contain both cloud-era Availability Zones and traditional Availability Zones that are local or geographically dispersed.

**Figure 7.** Multiple Types of Availability Zones in a Single Cloud



## Running the Management Server as an Application

CloudPlatform management server is designed to run as a traditional enterprise-grade application or traditional workload. It is designed as a simple, lightweight, and highly efficient application with the majority of work running inside system VMs ([see CloudPlatform Administration Guide—Working with System Virtual Machines](#)) and executed on computing nodes. This design choice was made for two reasons.


First, managing a cloud is not a cloud-scale problem. In CloudPlatform Version 3.0.x, each management server node is certified to manage 10,000 computing nodes. This level of scalability is sufficient for today's production cloud deployments. When CloudPlatform deployments continue to grow, it is expected that Citrix will be able to tune management server code so that each individual management server node can scale to many times the number of computing nodes.

The second reason for designing the management server as an enterprise application is a pragmatic one. Few people who deploy CloudPlatform will have a cloud-era infrastructure already in place. Without an existing IaaS cloud and third party management tools like RightScale or EnStratus in place, deploying a cloud workload is not an easy task. Building CloudPlatform management server for a cloud-era workload would therefore lead to a bootstrap problem.

### Management Server Cluster Backup and Replication

As a traditional-style enterprise application, the management server cluster is front-ended by a load balancer and connects to a shared MySQL database. While the cluster nodes themselves are stateless and can be easily recreated, the MySQL database node should be backed up and replicated to a remote site to ensure continuing operation of the cloud. Figure 8 illustrates how a standby management server cluster is set up in a remote data center.

**Figure 8.** Setup of Standby Management Server Cluster in Remote Data Center



During the normal course of operation, the primary management server cluster serves all user interface (UI) and API requests. Individual server failures in the management server cluster are protected as other servers in the cluster will take over the load.

To help ensure that the management server cluster can recover from a MySQL database failure, an identical database machine is set up to serve as the backup MySQL server. All database transactions are replayed in real time on the backup MySQL server in an active-passive setup. If the primary MySQL server fails, the administrator can reconfigure the management server cluster to point to the backup MySQL server.

To help ensure that the system can recover from the failure of the entire Availability Zone 1 that contains the primary management server cluster, a standby management server cluster can be set up in another Availability Zone. Asynchronous replication is set up between the backup MySQL server in the primary management server cluster and the MySQL server in the standby management server cluster. If Availability Zone 1 fails, a cloud

administrator can bring up the standby management server cluster and then update the Domain Name System (DNS) server to redirect cloud API and UI to the standby management server cluster.

### LAN Configuration


The access layer LAN configuration consists of a pair of Cisco Nexus 5548UP Switches, which are low-latency, line-rate, 10 Gigabit Ethernet and FCoE switches.

### Cisco UCS Connectivity

Four 10 Gigabit Ethernet uplink ports are configured on each of the Cisco UCS 6248UP Fabric Interconnects and they are connected to the Cisco Nexus 5548UP pair in a PortChannel as shown in Figure 9. The Cisco 6248UP is in end-host mode because both Fibre Channel and Ethernet network attached storage (NAS) data access are deployed following the recommended best practices for Cisco UCS. This configuration was built out for scale and more than 40 Gbps per Fabric Interconnect is provisioned.

**Note:** The upstream configuration is beyond the scope of this document. There are some good reference documents that cover best practices for [Cisco Nexus 5000](#) and [7000 Series Switches](#).

**Figure 9.** Network Configuration with Upstream Cisco Nexus 5000 Series Switches from Cisco UCS



### Storage LAN Connectivity

The Cisco Nexus 5548UP is used to connect to the storage system for Fibre Channel access. If the storage system is equipped with dual-port, 8-Gbps Fibre Channel modules, they should be connected to the pair of Cisco Nexus 5548UP Switches to provide block storage access to the environment.

If the storage system supports dual-port 10 Gigabit Ethernet, they should be configured in a PortChannel and connected to the pair of Cisco Nexus 5548UP Switches. It is recommended to ensure jumbo frames are

---

implemented on the ports and have priority flow control (PFC) enabled, with platinum class of service (CoS) and quality of service (QoS) assigned to the virtual network interface cards (vNICs) carrying the storage data access on the Fabric Interconnects.

### SAN Configuration

The same pair of Cisco Nexus 5548UP Switches can be used to connect the Fibre Channel ports on the storage system to the Fibre Channel ports of the Cisco UCS 6248UP Fabric Interconnect expansion modules. The Fibre Channel connection should be exclusively used for configuring boot-from-SAN. For more information, see the following section, "Boot from SAN."

Single virtual SAN (vSAN) zoning was set up on the Cisco Nexus 5548UP Switches to make these storage system logical unit numbers (LUNs) visible to the infrastructure and test servers.

### Boot from SAN

Booting from the SAN is another feature that facilitates the move toward stateless computing, where there is no static binding between a physical server and the OS and the applications that it is tasked to run. This feature is highly recommended for use with CloudPlatform as it enables administrators to easily move cloud workloads. The OS is installed on a SAN LUN and boot-from-SAN policy is applied to the service profile template or the service profile. If the service profile was moved to another server, the worldwide port names (WWPNs) of the host bus adapters (HBAs) and the boot-from-SAN policy would move along with it. The new server now takes the same exact view as the old server, demonstrating the true stateless nature of the blade server.

The main benefits of booting from the network include:

- **Smaller server footprint:** Boot from SAN alleviates the necessity for each server to have its own direct-attached disk, eliminating internal disks as a potential point of failure. Thin diskless servers also take up less facility space, require less power, and are generally less expensive because they have fewer hardware components.
- **Simplified disaster and server failure recovery:** All the boot information and production data stored on a local SAN can be replicated to a SAN at a remote disaster recovery site. If a disaster destroys the functionality of the servers at the primary site, the remote site can take over with little downtime. Recovery from server failure is simplified in a SAN environment. With the help of snapshots, mirrors of a failed server can be recovered quickly by booting from the original copy of the server image. As a result, boot-from-SAN capability can greatly reduce the time required for server recovery.
- **High availability:** A typical data center is highly redundant, with redundant paths, redundant disks, and redundant storage controllers. Storage of operating system images on the SAN supports high availability and eliminates the potential for the mechanical failure of a local disk.
- **Rapid redeployment:** Businesses that experience temporary high production workloads can take advantage of SAN technologies to clone the boot image and distribute the image to multiple servers for rapid deployment. Such servers may need to be in production for only hours or days and can be readily removed when the production need has been met. Highly efficient deployment of boot images makes temporary server use a cost-effective endeavor.
- **Centralized image management:** When operating system images are stored on networked disks, all upgrades and fixes can be managed at a centralized location. Changes made to disks in a storage array are readily accessible by each server.

Details for configuring SAN boot can be found in the following configuration guides

- [Cisco UCS Command-line Interface \(CLI\) Configuration Guide](#)
- [Cisco UCS Manager GUI Configuration Guide](#)

### Configuring CloudPlatform for Use with Cisco UCS

CloudPlatform manages the power and flexibility of Cisco UCS hardware with a single, simple web management interface.

Start by defining a zone (see Figure 10). A zone is the largest organizational unit within a CloudStack deployment. A zone typically corresponds to a single data center, although it is permissible to have multiple zones in a data center.

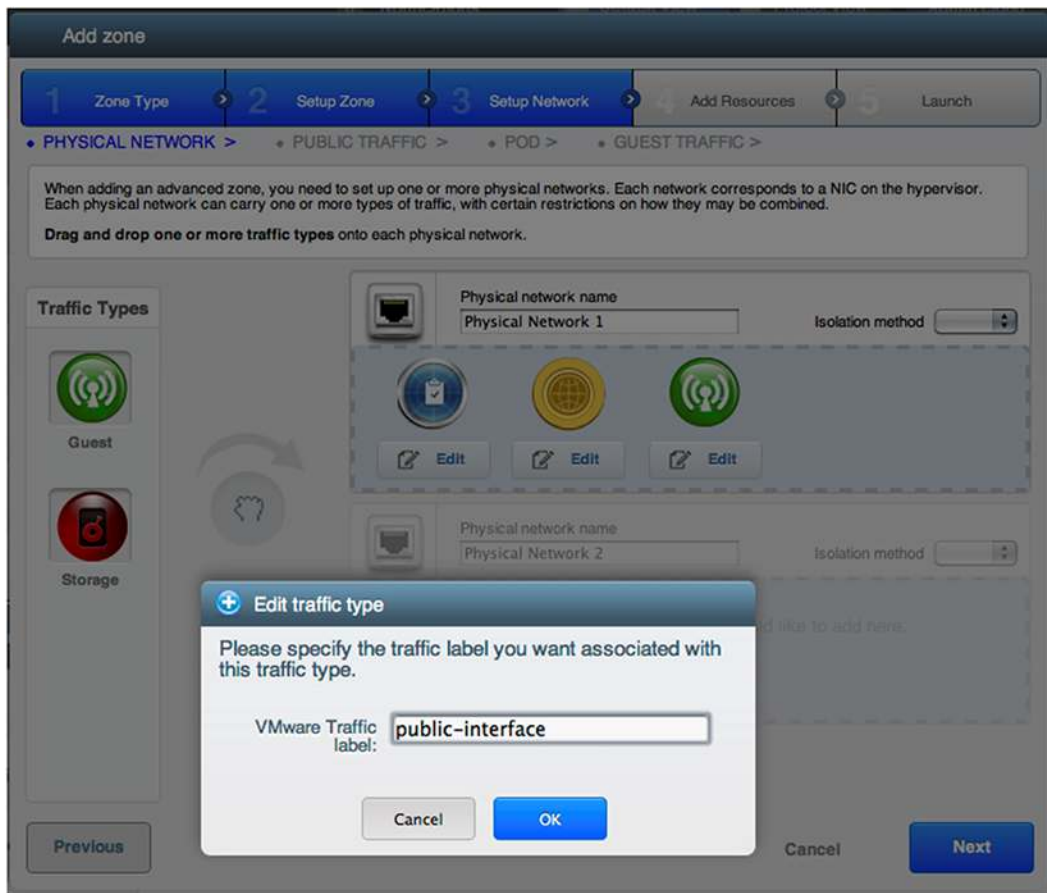
The benefit of organizing infrastructure into zones is to provide physical isolation and redundancy. CloudPlatform was deployed using an advanced networking model, which allows for more sophisticated network topologies. This network model takes advantage of the flexibility that Cisco UCS provides for easily splitting management, storage, public, and guest traffic types.

**Figure 10.** Adding a Zone

The screenshot displays the 'Add zone' web interface. At the top, a progress bar shows five steps: 1. Zone Type (selected), 2. Setup Zone, 3. Setup Network, 4. Add Resources, and 5. Launch. Below the progress bar, a text box explains: 'A zone is the largest organizational unit in CloudStack, and it typically corresponds to a single datacenter. Zones provide physical isolation and redundancy. A zone consists of one or more pods (each of which contains hosts and primary storage servers) and a secondary storage server which is shared by all pods in the zone.' The main form contains several input fields: 'Name' (Cisco-SC), 'DNS 1' (72.52.126.11), 'DNS 2' (72.52.126.12), 'Internal DNS 1' (10.119.110.253), 'Internal DNS 2' (10.119.110.254), 'Hypervisor' (XenServer), and 'Network Domain' (empty). At the bottom, there are 'Previous', 'Cancel', and 'Next' buttons.

Figure 11 shows the ease of creating these mappings using a drag-and-drop interface.

**Figure 11.** Creating Mappings with Drag and Drop Interface



Next, a pool of public IP addresses is added to the cloud zone. Public traffic is generated when VMs in the cloud access the Internet. End users can use the CloudPlatform UI or API to acquire these IP addresses to implement Network Address Translation (NAT) and load balancing services hosted on their VMs (Figure 12). Public network traffic is metered and billing metrics are generated when users acquire more IP addresses or generate traffic.

**Figure 12.** Figure 12. Adding IP Addresses to the Cloud Zone

**Add zone**

1 Zone Type > 2 Setup Zone > 3 Setup Network > 4 Add Resources > 5 Launch

• PUBLIC TRAFFIC > • POD > • GUEST TRAFFIC >

Public traffic is generated when VMs in the cloud access the internet. Publicly-accessible IPs must be allocated for this purpose. End users can use the CloudStack UI to acquire these IPs to implement NAT between their guest network and their public network.

Provide at least one range of IP addresses for internet traffic.

Gateway	Netmask	VLAN	Start IP	End IP	Add	Actions
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Add"/>	
72.52.119.1	255.255.255.0	119	72.52.119.2	72.52.119.254		<input type="button" value="X"/>

After the public network is added, the pods are the next resource to be added (Figure 13). A pod often represents a single rack. Hosts in the same pod are managed through the same subnet.

A pod is the second-largest organizational unit within a CloudStack deployment. Pods are contained within zones. Each zone can contain one or more pods. A pod consists of one or more clusters of hosts and one or more primary storage servers.



**Figure 13.** Adding Pods

**Add zone**

1 Zone Type > 2 Setup Zone > 3 Setup Network > 4 Add Resources > 5 Launch

• PUBLIC TRAFFIC > • **POD** > • GUEST TRAFFIC >

Each zone must contain in one or more pods, and we will add the first pod now. A pod contains hosts and primary storage servers, which you will add in a later step. First, configure a range of reserved IP addresses for CloudStack's internal management traffic. The reserved IP range must be unique for each zone in the cloud.

\* Pod name: testpod

\* Reserved system gateway: 10.223.101.1

\* Reserved system netmask: 255.255.255.0

\* Start Reserved system IP: 10.223.101.210

End Reserved system IP: 10.223.101.230

Previous Cancel Next

A cluster provides a way to group hosts. It is a XenServer server pool, a set of KVM servers, or a VMware cluster preconfigured in vCenter. The hosts in a cluster all have identical hardware, run the same hypervisor, are on the same subnet, and access the same shared primary storage. VM instances can be live-migrated from one host to another within the same cluster without interrupting service to the user. ESXi 5.0 was used in the scope of this document but could have been easily implemented using any of the supported hypervisors. Figure 14 shows how clusters are added.



**Figure 14.** Adding Resources

**Add zone**

1 Zone Type > 2 Setup Zone > 3 Setup Network > 4 Add Resources > 5 Launch

• CLUSTER > • HOST > • PRIMARY STORAGE > • SECONDARY STORAGE >

Each pod must contain one or more clusters, and we will add the first cluster now. A cluster provides a way to group hosts. The hosts in a cluster all have identical hardware, run the same hypervisor, are on the same subnet, and access the same shared storage. Each cluster consists of one or more hosts and one or more primary storage servers.

Hypervisor: VMware

\* Cluster Name:

\* vCenter Host:

\* vCenter Username:

\* vCenter Password:

\* vCenter Datacenter:

Previous Cancel Next

Primary storage is associated with a cluster and it stores the disk volumes for all the VMs running on hosts in that cluster. You can add multiple primary storage servers to a cluster. At least one is required. It is typically located close to the hosts for increased performance

The final step is to associate a secondary storage server with the zone. Secondary storage acts as a library or repository for:

- Templates: OS images that can be used to boot VMs and can include additional configuration information, such as installed applications
- ISO images: Disk images containing data or bootable media for operating systems
- Disk volume snapshots: Saved copies of VM data which can be used for data recovery or to create new templates

At this point, VMs, storage, multitiered networks, and run workloads may be provisioned in the cloud, all orchestrated with CloudPlatform on Cisco UCS infrastructure.

---

## Scalability References

The following scaling references are intended to provide examples of how the Cisco UCS B200 M3 Blade Servers can scale given a sample workload. Two common cloud workloads were selected as examples:

- XenServer-based VDI workload
- Apache Geronimo Daytrader workload

The workload examples provided are based on the capabilities of the blade in similar infrastructure configurations using enterprise-class storage arrays.

### VDI Workload Scaling Example

One common type of cloud workload is virtual desktop infrastructure (VDI). This VDI scaling example shows the capabilities of the Cisco UCS B200 M3 Blade Server when used in conjunction with Citrix XenDesktop. This example should be used to get a general idea of the compute requirements to support CloudPlatform-backed VDI environments.

#### VDI Configuration Details

##### About Citrix XenDesktop 5.5

Citrix XenDesktop is a desktop virtualization solution that transforms Windows desktops and applications into an on-demand service available to any user, anywhere, on any device. With XenDesktop, you can securely deliver individual Windows, Web, and software-as-a-service (SaaS) applications or full virtual desktops, to PCs, Macs, tablets, smartphones, laptops, and thin clients—all with a high-definition user experience. XenDesktop 5.5 adds Personal vDisk technology to all editions of XenDesktop, making it far easier for customers to deploy highly personalized virtual desktops at a substantially lower cost than conventional VDI solutions. Citrix XenDesktop 5.5 includes High Definition User Experience™ (HDX™) technology and MediaStream Flash Redirection (MFR), which moves the processing of Adobe Flash content to users' local devices to avoid using server and network resources and provide a better user experience.

To learn more about Citrix XenDesktop 5.5, visit:

[http://www.citrix.com/English/ps2/products/product.asp?contentID=163057&ntref=prod\\_cat](http://www.citrix.com/English/ps2/products/product.asp?contentID=163057&ntref=prod_cat).

##### About Login VSI

Login VSI benchmarks virtual desktop sessions to determine the relative scalability of a particular virtualized system. Specifically, it benchmarks a virtual desktop solution by simulating Windows-based Microsoft Office user workloads. The medium workload of Login VSI, which has been tested as indicative of a typical knowledge worker, opens and closes the following applications and runs their respective tasks:

- Microsoft Outlook®: Browsing a message
- Microsoft Word® (TimerDoc): Initiating response timer to see how the program responds throughout the workload
- Microsoft Internet Explorer® instance one: Maximizing, scrolling, and minimizing
- Microsoft Internet Explorer instance two: Navigating a Web site, maximizing, and scrolling
- Adobe® Flash® KA movie trailer
- Microsoft Word (UserRead): Reading and typing text, and printing to PDF

- 
- Bullzip: Generating a PDF
  - Adobe Reader®: Reading a PDF
  - Microsoft PowerPoint®: Watching a presentation and adding a slide
  - Microsoft Excel®: Reading and minimizing
  - 7-Zip: Saving a ZIP file

Login VSI Version 3.0 (Release 6) benchmarks user experience more effectively than previous versions of Login VSI because its workloads and what the VSI Index measures more accurately reflect the tasks actual users perform on their virtual desktops. Reported response times are higher in Login VSI 3.0 than in Login VSI 2.0 and other previous versions because the benchmark uses this heavier workload. The Login VSI benchmark mandates the minimum acceptable response time for the testing.

The Login VSI 3.0 benchmark uses seven operations to determine the VSImax, the maximum number of users the system can handle before suffering serious degradation in performance. By using seven operations instead of only two, as earlier versions of Login VSI did, Login VSI 3.0 better reflects what a user actually experiences. The seven operations are as follows:

- Copying a new document from the document pool in the home drive
- Starting a Microsoft Word document
- Starting the File Open dialogue
- Starting the Search and Replace dialogue
- Starting the Print dialogue
- Starting Notepad
- Compressing the document into a ZIP file with 7-zip command line

Login VSI records response times, the time taken to execute a given task, in milliseconds. Login VSI then reports minimum, average, and maximum response times, as well as the VSI Index average while performing the workload. The Login VSI Index average is similar to the average response time, as it averages the maximum and minimum response times, but it removes 2 percent from the maximum and minimum response time before calculating the average. VSImax is then calculated in one of two ways: Classic and Dynamic.

When the VSI Index average is higher than the default threshold of 4000 ms, Classic VSImax is achieved. Dynamic VSImax calculates a dynamic threshold based on the average response times of the first 15 sessions and applies the formula  $\text{baseline} \times 125\% + 3000$ , when the VSI Index is higher than the dynamic baseline then dynamic VSImax is achieved. In our testing, Dynamic VSImax was calculated to be 182 sessions.

It is important to note that variations in hypervisor, application, guest OS, and virtual desktop infrastructure (VDI) settings can have a significant impact on expected user density in these tests. Tuning frames per second, image compression, screen resolution, and other user-experience-specific settings can increase or decrease the number of desktops a system can support. Generally, improving user experience will decrease the number of supported desktops. It is therefore important to understand and quantify the specific needs of VDI users and create baseline settings to ensure that the results are representative of your environment.

For more information on Classic and Dynamic VSImax, see:

<http://www.loginvsi.com/en/admin-guide/analyzing-results#h0-1-calculating-vsimax>.

For more information on Login VSI 3.0, see:


[http://www.loginconsultants.com/index.php?option=com\\_content&task=view&id=390](http://www.loginconsultants.com/index.php?option=com_content&task=view&id=390).

### Test Results

After all desktops are idle, Login VSI incrementally logs users into virtual desktop sessions and begins workloads on each. Login VSI measures the total response times of seven typical office operations from each session and calculates the VSI Index by taking the average response times and dropping the highest and lowest 2 percent. The average response time of the first 15 sessions determines a baseline; the Dynamic VSImax is baseline x 125% + 3000ms. As more sessions begin to consume system resources, response times degrade and the VSI index increases until it is above the Dynamic VSImax. When this condition is met, the benchmark records a Login VSImax, which is the maximum number of sessions that the platform can support.

Because the VSI Index drops the highest two percent of response times, we needed to use 191 virtual desktop sessions to reach the Login VSImax of 182 for the Cisco UCS B200 M3 Blade Server. Figure 15 shows the VSI Index average and average response times for all active sessions recorded during the test. The Cisco UCS B200 M3 Blade Server was able to support 182 virtual desktops based on the Login VSImax assigned by the Login VSI benchmark. User response time degraded only when all 16 cores were nearly saturated.

**Figure 15.** Average Virtual Desktop Response Times



To initiate our test, we enabled our Citrix XenDesktop 5.5 pool to start up all Windows 7 VMs and reach an idle state. We monitored our test bed during the startup and perceived no bottlenecks in server CPU, network, or storage I/O at the Citrix XenDesktop default startup rate of 10 actions per minute. When the VMs were idle, we started Login VSI testing. Figure 16 shows the processor utilization throughout the test. With 191 simultaneous users, 182 of which achieved an acceptable response time as determined by the Login VSI 3.0 benchmark, nearly all 16 processor cores were at 100 percent utilization. The graph line represents the average utilization across all 16 cores (32 threads) throughout the boot, idle, and testing phases. When the test was complete, all virtual desktops began to log off.

**Figure 16.** Processor Utilization During Test




Figure 17 shows the memory usage throughout the test. The steep increase at the beginning of the chart reflects the beginning of the test with the 191 virtual desktops powering on.

**Figure 17.** Memory Using During Test

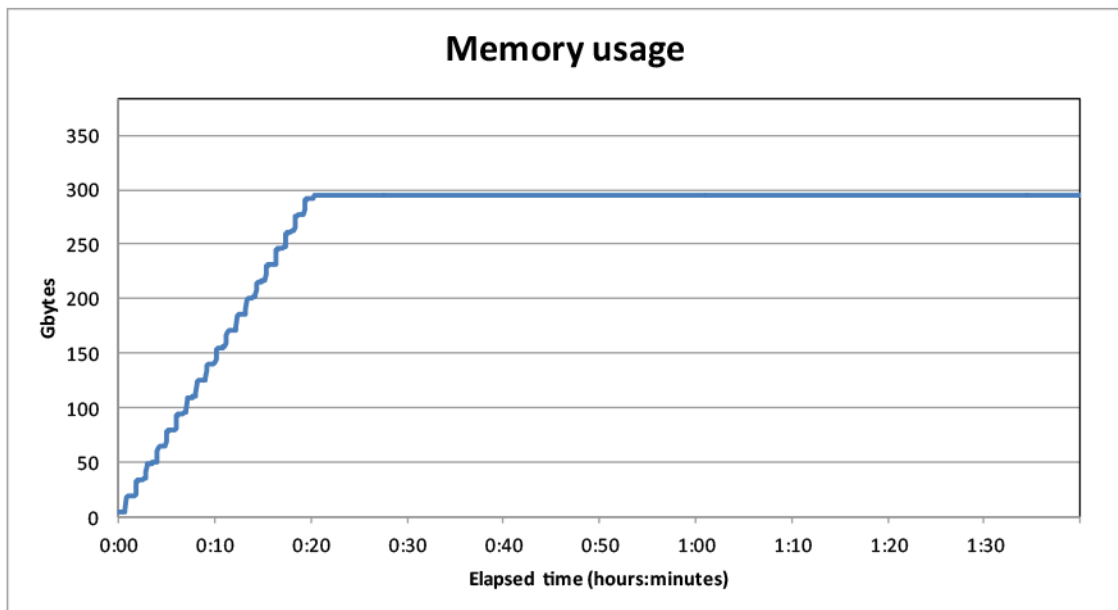


Figure 18 shows the fabric usage throughout the test. The peak usage was due to virtual desktop receiving streaming VHDs via Citrix Provisioning Services at boot up.

**Figure 18.** Fabric Usage During Test

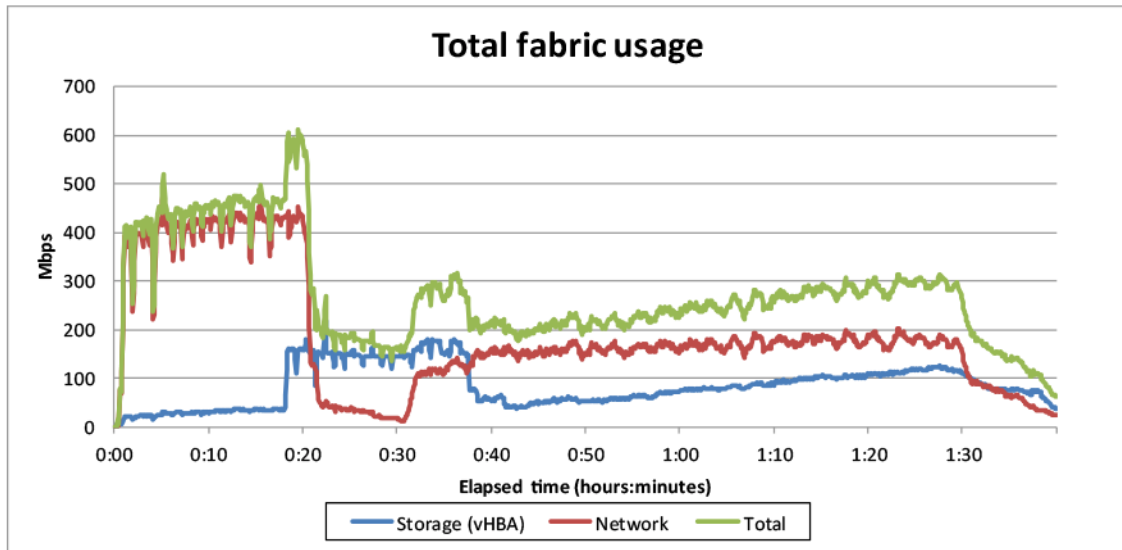
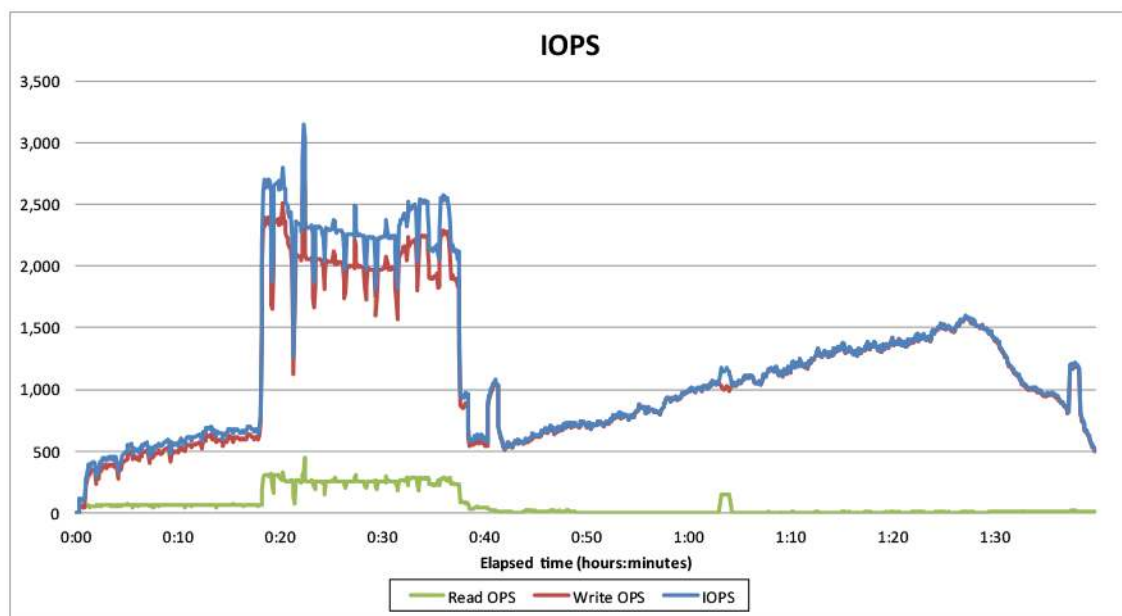



Figure 19 shows the input/output operations per second (IOPS) recorded throughout the test. Citrix Provisioning Services caches the desktop image in memory, resulting in a significant reduction in disk read operations (see Figures 20 and 21).


**Figure 19.** IOPS DuringTest



**Figure 20.** Provisioning Services IOPS Throughout Test



**Figure 21.** Provisioning Services Network Usage During Test



### Three-tier Application (Open-Source) Scaling Example

For any typical three-tier enterprise applications, there are several technologies to implement including some of the widely know technolgoies as such as LAMP stack, ASP .NET framework and JAVA/J2EE platform.

LAMP is a technology stack build using open source software that includes the GNU/Linux operating system, Apache Web server, MySQL database and scripting languages such as PHP, Perl and/or Python. The exact combination of software included in a LAMP package may vary, especially with respect to the web scripting software, as PHP may be replaced or supplemented by Perl and/or Python. Individually, Linux, Apache Web server, MySQL database, Perl, Python or PHP are each a powerful component in its own right. The key to the idea behind LAMP is to provide open source software alternatives which are readily and freely available. This has lead to them often being used together. In the past few years, their compatibility and use together has grown tremendously in small, medium and even in large enterprise web deployments.

### Deployment Architecture

For the purposes of this study, Citrix CloudPlatform, Citrix XenServer, and an open source DayTrader application were installed on a VM. The components used are listed in Table 15. Figure 22 shows the deployment topology.

**Table 15.** List of Components

<b>Citrix CloudPlatform</b>	Citrix CloudPlatform 3.0.5 is deployed on a Cisco UCS B200 M3 Blade Server equipped with two eight-core Intel Xeon E5-2690 processors at 2.90 GHz with a physical memory of 32 GB.
<b>Citrix XenServer</b>	Citrix XenServer version 6.0.2 is deployed on a second Cisco UCS B200M3 blade server equipped with two eight-core Intel Xeon E5-2690 processors at 2.90 GHz with a physical memory of 256 GB.
<b>DayTrader Application</b>	Apache 2.2 (the Web server) / JBOSS 5.1 (the application server) / MySQL-server 5.1.52 (the database server) is deployed on a VM with 2 vCPUs / 8 GB of physical memory.
<b>Storage</b>	EMC Clariion CX4 is configured with RAID1/0 using 100 GB of storage space.
<b>Guest Operating System</b>	Red Hat Enterprise Linux Server Release 6.0.2



**Figure 22.** Deployment Setup




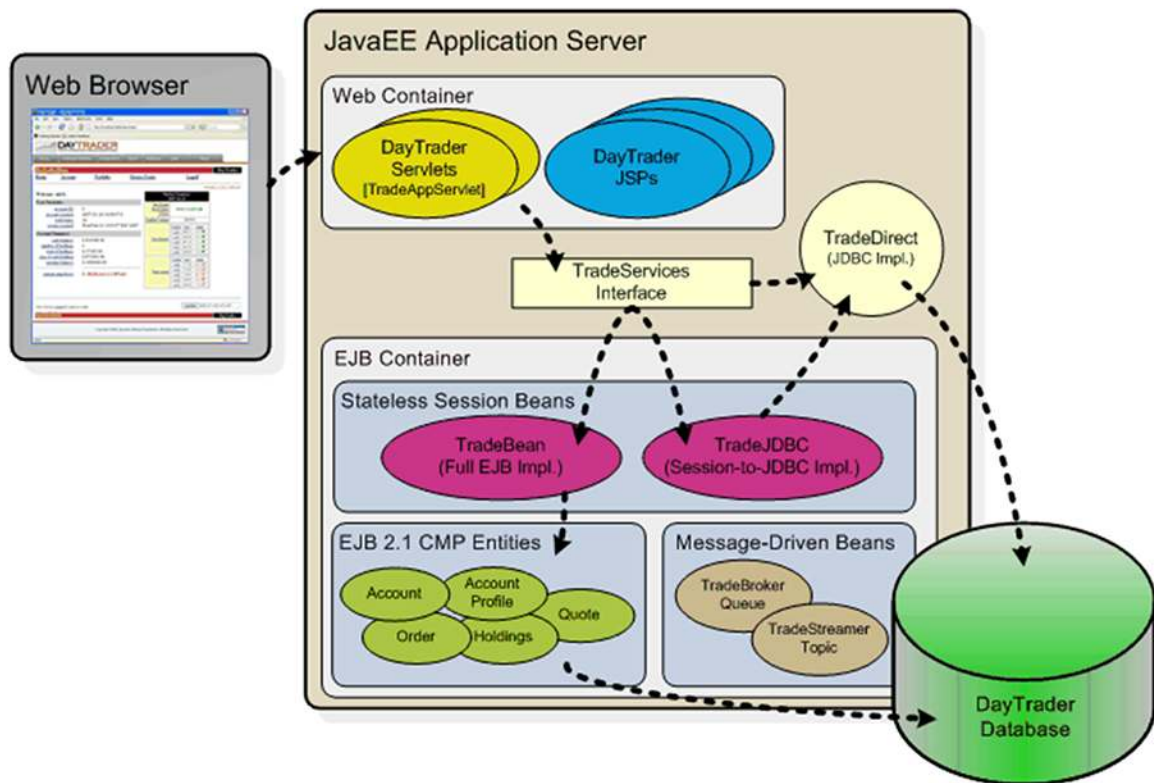
Figure: Citrix Cloudstack / XenServer – Topology

### DayTrader Application and Workload

DayTrader is a benchmark application built on the paradigm of an online stock trading system. This application allows users to log in, view their portfolio, look up stock quotes, and buy or sell stock shares. DayTrader is built on a core set of Java Enterprise Edition (EE) technologies that include Java Servlets and Java Server Pages (JSPs) for the presentation layer and Java database connectivity (JDBC), Java Message Service (JMS), Enterprise JavaBeans (EJBs), and Message-Driven Beans (MDBs) for the back-end business logic and persistence layer (see Figure 23).

The DayTrader application is typically used for analyzing the performance of three-tier application architecture as it provides a sample representative workload in an online web application paradigm. It has been widely used by a variety of vendors, including IBM, Red Hat and Oracle.

**Figure 23.** High-Level Overview of DayTrader Application and Various Components




### Performance Analysis

This performance study is designed to represent small, three-tier application deployments. The performance and scalability of the application were evaluated based on these important criteria:

- Throughput, such as transactions per second, achieved from the deployed application with the condition of acceptable application response time or saturation of available system resources
- Response time required to execute a deployed application transaction (such as buy and sell, buy only, and quote transactions)

Hewlett-Packard's LoadRunner tool is used to simulate the user workload as illustrated in Figure 24.

**Figure 24.** Test Bed Setup



#### Workload Mix

For a realistic load simulation, test scripts were modeled after DayTrader business processes (for example, transactions such as buy, portfolio, and quote). These transactions were simulated with independent load driver threads to ensure unique users for each type of transaction. The percentage of users dedicated to each transaction were:

- 37.5 percent of users executed buy and sell transactions
- 37.5 percent of users executed buy-only transactions
- 25 percent of users executed quote transactions

The run-time mode chosen for the DayTrader application configuration were: 1) Full EJB3; 2) Direct (JDBC); and 3) Session (EJB3) To Direct mode. The run-time mode for the executed benchmark described in subsequent result sections was configured with Full EJB3 using Hibernate as the Java Persistence API (JPA) layer.


The test was executed with a 320 concurrent user load with a database populated with approximately 5000 users and 2000 quotes.

#### Test Results

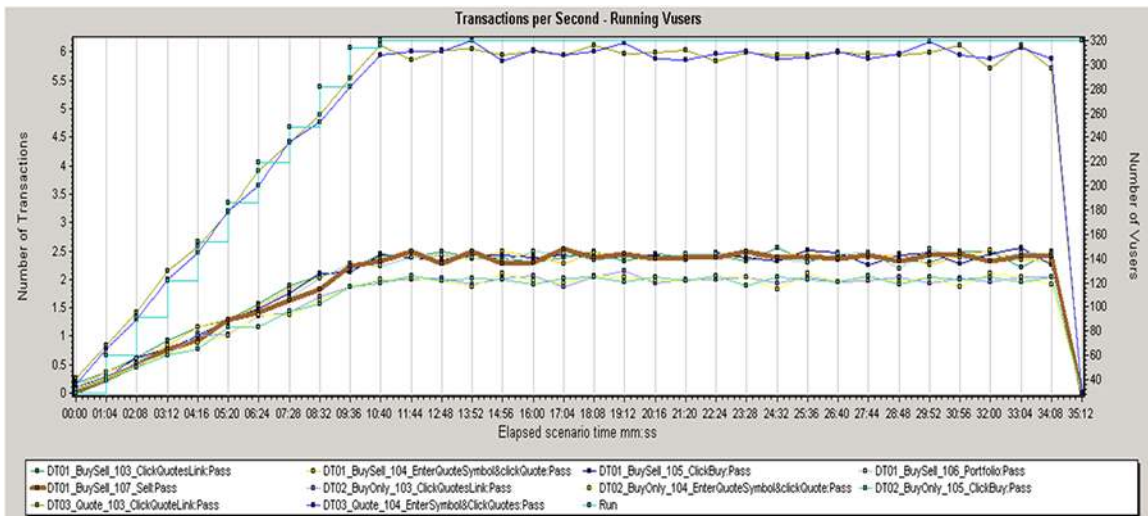
Each of the benchmark tests was executed for 30 minutes (steady state duration after ramp-up), making sure that the threshold of transaction response times did not exceed two seconds.

Figures 25 and 26 show the response times and transactions per second (throughput) for the entire test.

**Figure 25.** Application Response Time



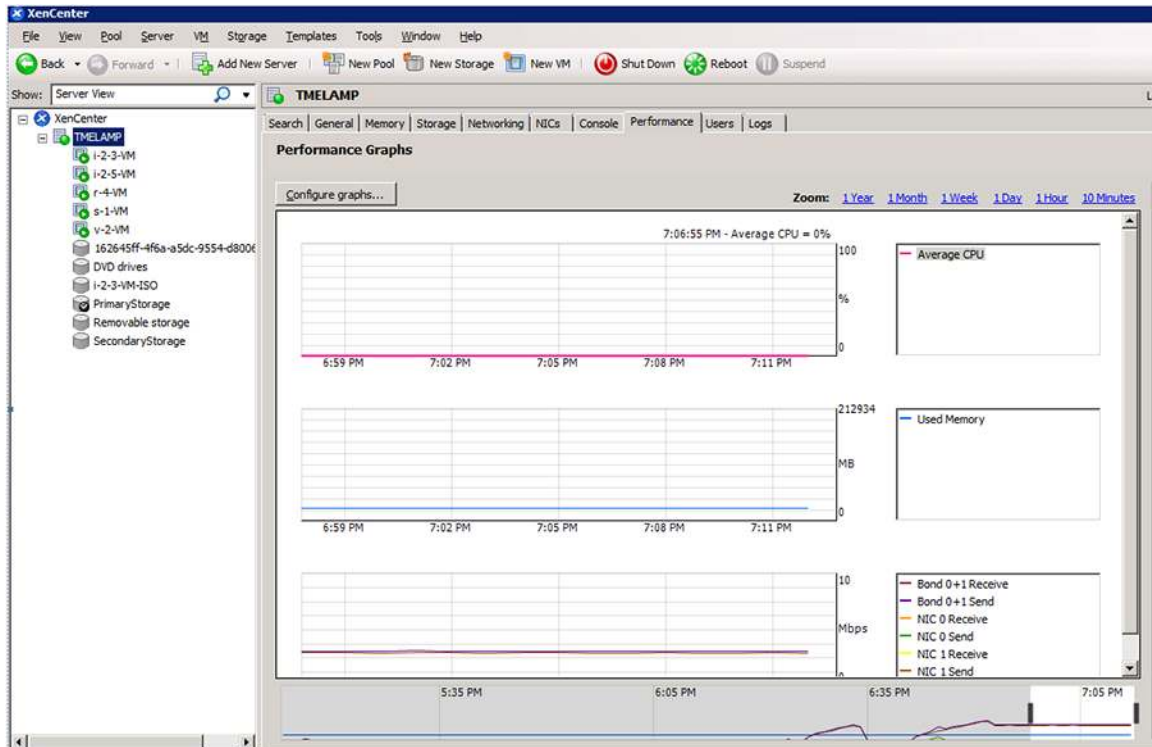
**Figure 26.** Application Transactions per Second (Throughput)



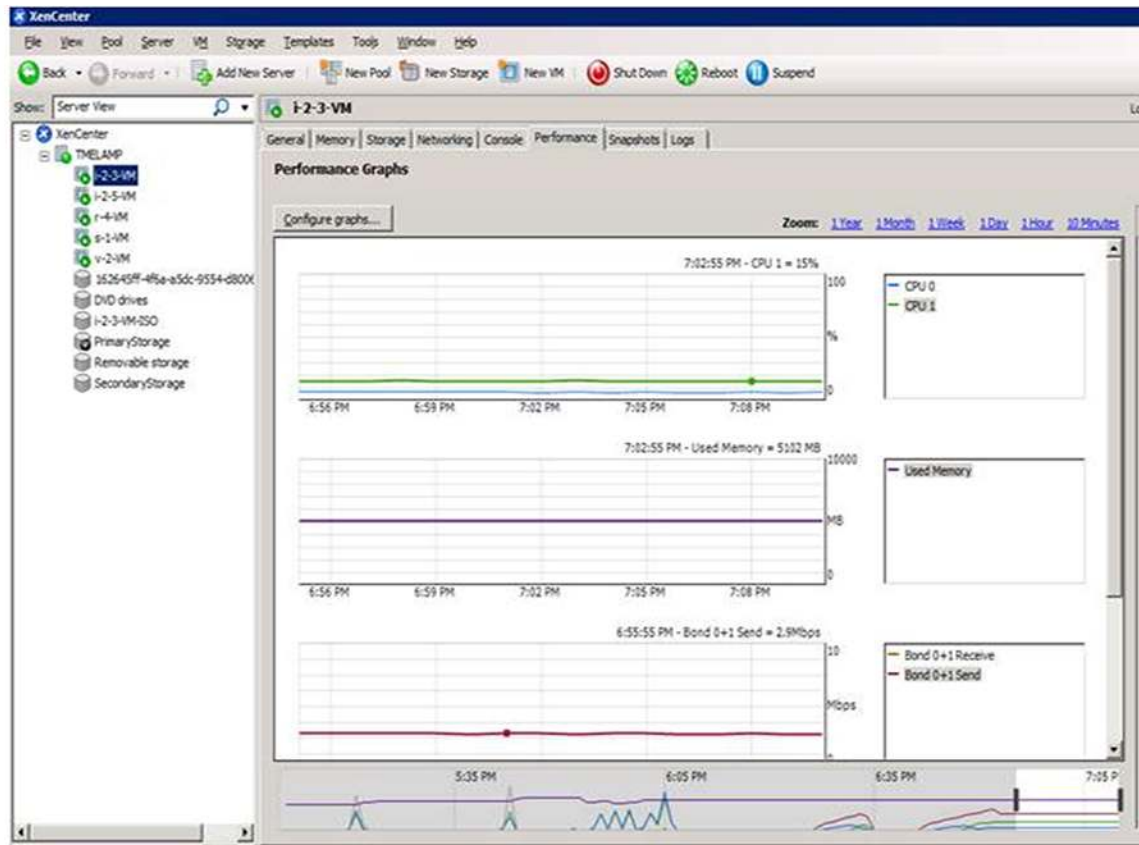
From the graphs shown in Figures 25 and 26, it can be concluded that the transactions were under one-second response times and consistent transactions per second were observed once all users were in a steady state.

For the resource usage at the VM level as well at the physical blade level, the XenServer performance graphs are captured in Figures 27 and 28.

Figure 27. Physical Server Resource Usage



**Figure 28.** VM Resource Usage



From Figures 27 and 28, it is evident that the allocated 2 vCPU/ 8-GB RAM can handle the 320-user workload in the DayTrader application without any resource contention or impact on response times for transactions per second. Using this data, the total count of VMs that can be hosted on a Cisco UCS B200 M3 Blade Server (the current XenServer setup) can be calculated as shown here:

- The CPU usage seen on VM on one core of the CPU allocated (in a non-oversubscribed scenario) = 15%
- With the assumption of 10% virtualization overhead, the physical CPU usage = 17%
- The number of VMs that can be created per CPU core (70% max CPU usage limit) =  $70/17 \sim 4$  VMs
- Total number of cores (B200 M3 has 16 cores and assuming HyperThreading gives 25% gain) =  $16 \times 1.25 = 20$  cores
- Total VMs in 20 core scenario =  $4 \times 20 = 80$  VMs
- Though theoretically 80 VMs can be created, available memory will become a limiting factor (maximum addressable memory in B200 M3 with 16-GB DIMMs is 384 GB).
- Hence the total number of VMs in this scenario would be (leaving 10% of memory for the host OS) = 40 VMs

---

## Conclusion

Cisco and Citrix are both committed to providing superior cloud solutions to a global clientele. Together, products from both companies enable the flexible and agile delivery of cloud-based services with simplified physical and virtual infrastructure management. Many established enterprises have selected the combination of Cisco and Citrix offerings after analyzing best-in-class cloud products and have built robust and scalable enterprise clouds based on Cisco UCS and Citrix CloudPlatform. The reference architecture in this document further simplifies the deployment of these products by providing best practices for infrastructure and software configurations.

## For More Information

[Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide, Release 5.0\(3\)N2\(1\)](#)

[Cisco Nexus 7000 Series Switches Data Sheet](#)

[Cisco Nexus 7000 Series Switches Configuration Guides](#)



**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).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)